

**FACULTY
OF MATHEMATICS
AND PHYSICS**
Charles University

DOCTORAL THESIS

Matěj Konečný

**Model theory and extremal
combinatorics**

Department of Applied Mathematics

Supervisor of the doctoral thesis: doc. Mgr. Jan Hubička, Ph.D.

Consultant of the doctoral thesis: prof. RNDr. Jaroslav Nešetřil,
DrSc.

Study programme: Computer Science - Theory of
Computing, Discrete Models and
Optimization (P0613D140007)

Study branch: Computer Science - Theory of
Computing, Discrete Models and
Optimization

Prague 2023

I declare that I carried out this doctoral thesis independently, and only with the cited sources, literature and other professional sources. It has not been used to obtain another or the same degree.

I understand that my work relates to the rights and obligations under the Act No. 121/2000 Sb., the Copyright Act, as amended, in particular the fact that the Charles University has the right to conclude a license agreement on the use of this work as a school work pursuant to Section 60 subsection 1 of the Copyright Act.

In date

Author's signature

I would like to thank Honza Hubička and Jarik Nešetřil for all the support and generosity that they showed me ever since, as a second year undergraduate, I started showing up to their Ramsey DocCourse lectures. They kindly included me in their research projects, took me to various conferences and workshops and helped me all along the way.

I am incredibly grateful to my wife Inka for all the encouragement and patience, especially when, at times, I was away for long periods of time. I thank my parents and all my friends and family for all the help that they gave me.

What I enjoy about Mathematics the most is its collaborative nature. Over the course of my journey, I have been incredibly fortunate to have many amazing coauthors and I would like to thank them for letting me work with them. Thank you, Andrés Aranda, Martin Balko, Adam Bartoš, Tristan Bice, David Bradley-Williams, Sam Braunfeld, Gregory Cherlin, David Chodounský, Noé de Rancourt, Natasha Dobrinen, David Evans, Eng Keat Hng, Colin Jahel, Miltos Karamanlis, Jamal Kawach, Michael Kompatscher, Standa Kučera, Yibei Li, Janča Novotná, Michal Opler, Micheal Pawliuk, Jakub Pekárek, Vašek Rozhoň, Marcin Sabok, Norbert Sauer, Martin Smolík, Jakub Sosnovec, Štěpán Šimsa, Kuba Tětek, Martin Töpfer, Lluís Vena, Martin Ziegler, and Andy Zucker.

I have also been incredibly fortunate to be able to talk or correspond about Mathematics with many other smart people to whom I am thankful for the time they gave me. Thank you, Miklós Abért, Dana Bartošová, Manuel Bodirsky, Peter Cameron, Honza Grebík, Geňa Hahn, Michael Hrušák, Dan Král, Gábor Kun, Ola Kwiatkowska, Dugald Macpherson, Dragan Mašulović, Nadav Meir, Justin Moore, Antoine Mottet, Lionel Nguyen Van Thé, Patrice Ossona de Mendez, Michael Pinsky, Márton Pósfai, Maurice Pouzet, Vojta Rödl, Pythagoras of Samos, Pierre Simon, Sławek Solecki, Balázs Szegedy, Stevo Todorčević, Honza Volec, Robert Woodrow, Dima Zhuk, and others to whom I apologize greatly for forgetting them.

Finally, I would like to thank all the people at the Department of Applied Mathematics, Computer Science Institute of Charles University, and the School of Computer Science of the Faculty of Mathematics and Physics in general for providing a friendly and creative environment to spend my student years in.

Title: Model theory and extremal combinatorics

Author: Matěj Konečný

Department: Department of Applied Mathematics

Supervisor: doc. Mgr. Jan Hubička, Ph.D., Department of Applied Mathematics

Consultant: prof. RNDr. Jaroslav Nešetřil, DrSc., Computer Science Institute of Charles University

Abstract: This thesis is concerned with combinatorial properties of homogeneous structures such as the Ramsey property, big Ramsey degrees, EPPA, and others. What these properties have in common is that, while being finitary problems on classes of finite structures, they are equivalent to various dynamical properties of automorphism groups of the corresponding homogeneous structures. This thesis consists of an extended introduction to these areas, a list of open problems, and ten papers of which the author is a co-author, seven of which have been published at the time of writing this thesis, the other three have been submitted. The goal is to demonstrate that, at least on the combinatorial side of things, there are many interplays of these properties which can be (and have been) exploited to further each of the areas.

Keywords: homogeneous structures, automorphism groups, Ramsey property, big Ramsey degrees, EPPA

Contents

Preface	5
1 Background	8
1.1 Structures and morphisms	8
1.2 Amalgamation	9
1.2.1 Fraïssé’s theorem	10
1.3 Automorphism groups	11
1.4 A modest bestiary of homogeneous structures and amalgamation classes	12
1.5 Stationary independence relations and simple automorphism groups	18
1.6 EPPA and ample generics	20
1.6.1 (Dis)proving EPPA	20
1.6.2 General constructions	23
1.6.3 Examples	27
1.6.4 Profinite topology	28
1.6.5 Coherent EPPA	30
1.6.6 Ample generics	30
1.7 Structural Ramsey theory	34
1.7.1 Expansions	37
1.7.2 The KPT correspondence	38
1.7.3 Proving the Ramsey property	40
1.7.4 How to apply Theorem 1.7.17 in practice?	41
1.7.5 Examples	46
1.7.6 Ramsey degrees	48
1.7.7 Category theory	48
1.7.8 Dual Ramsey theorem(s)	49
1.7.9 Approximate Ramsey property	50
1.7.10 Conclusion	50
1.8 Big Ramsey degrees	51
1.8.1 Big Ramsey degrees of (\mathbb{Q}, \leq)	53
1.8.2 Milliken’s tree theorem	55
1.8.3 Big Ramsey degrees of the Rado graph	57
1.8.4 Parameter spaces and the Carlson–Simpson theorem	59
1.8.5 The triangle-free graph and age changes	61
1.8.6 Diaries	63
1.8.7 3-uniform hypergraphs	65
1.8.8 A cookbook?	68
1.8.9 Examples	69
1.8.10 Big Ramsey structures and topological dynamics	71
1.8.11 From big Ramsey degrees to Ramsey classes	72
1.8.12 Even bigger Ramsey properties?	73
1.8.13 Metric big Ramsey degrees	73
2 Problems, questions and conjectures	75

3	Simplicity of the automorphism groups of generalised metric spaces	85
3.1	Introduction	85
3.1.1	Stationary independence relations	86
3.2	Geodesic sequences	89
3.3	Proof of Theorem 3.1.1	91
3.4	Corollaries	94
3.4.1	Semigroup-valued metric spaces	94
3.4.2	Metrically homogeneous graphs	97
3.5	Conclusion	97
4	All those EPPA classes (Strengthenings of the Herwig–Lascar theorem)	99
4.1	Introduction	99
4.1.1	Γ_L -structures	100
4.1.2	The main results	100
4.1.3	Applications of our results	103
4.1.4	EPPA and Ramsey	103
4.2	Background and notation	105
4.2.1	Maps between Γ_L -structures	106
4.2.2	Γ_L -structures as standard model-theoretic structures	107
4.2.3	Amalgamation classes	108
4.2.4	EPPA for Γ_L -structures	109
4.2.5	EPPA and automorphism groups	110
4.2.6	Coherence of EPPA-witnesses	113
4.3	Warm-up: new proof of EPPA for graphs	114
4.4	Coherent EPPA for relational structures	118
4.5	Infinite languages	122
4.5.1	Proof of Proposition 4.5.1	123
4.6	EPPA for structures with unary functions	125
4.7	Irreducible structure faithful EPPA	128
4.8	Unwinding induced cycles	135
4.9	Locally tree-like EPPA-witnesses: Proof of Theorem 4.1.2	142
4.10	A generalisation of the Herwig–Lascar theorem: Proof of Theorem 4.1.5	146
4.11	Connections to the structural Ramsey theory: Proof of Theorem 4.1.6	147
4.12	Applications	149
4.12.1	Free amalgamation classes	150
4.12.2	Metric spaces without large cliques	150
4.12.3	Structures with constants	152
4.12.4	EPPA for special non-unary functions	155
4.12.5	EPPA for k -orientations with d -closures	156
4.13	Conclusion	167
5	EPPA for two-graphs and antipodal metric spaces	169
5.1	Introduction	169
5.2	Notation and preliminaries	171
5.2.1	Coherent EPPA	172
5.3	EPPA for antipodal metric spaces	172

5.4	Proofs of the main results	175
5.5	Existence of a dense locally finite subgroup	176
5.6	Amalgamation property with automorphisms	177
5.7	Ramsey expansion of two-graphs	179
5.8	Remarks	181
6	Extending partial isometries of antipodal graphs	183
6.1	Introduction	183
6.2	Preliminaries	184
6.2.1	Γ_L -structures	185
6.2.2	EPPA for Γ_L -structures	187
6.2.3	Metrically homogeneous graphs	190
6.3	The odd diameter non-bipartite case	192
6.3.1	Motivation	192
6.3.2	The expanded language	194
6.3.3	The class \mathcal{K} and completion to it	196
6.3.4	Constructing the witness	197
6.3.5	Extending partial automorphisms	197
6.3.6	Remarks	198
6.4	The even diameter bipartite case	199
6.4.1	The expanded language	200
6.4.2	The class \mathcal{K} and completion to it	201
6.4.3	Constructing the witness	202
6.4.4	Extending partial automorphisms	202
6.5	Conclusion	202
6.6	Acknowledgements	203
7	Ramsey expansions of 3-hypertournaments	204
7.1	Introduction	204
7.2	Preliminaries	204
7.2.1	Homogeneous 4-constrained 3-hypertournaments	205
7.3	Positive Ramsey results	206
7.4	The \mathbf{H}_4 -free case	207
8	Big Ramsey degrees of 3-uniform hypergraphs are finite	210
8.1	Introduction	210
8.2	Preliminaries	211
8.3	Proof of Theorem 8.1.1	213
8.4	Concluding remarks	218
8.5	Acknowledgment	219
9	Big Ramsey degrees and infinite languages	220
9.1	Introduction	220
9.2	Preliminaries	221
9.2.1	Milliken's tree theorem	222
9.3	Valuation trees	223
9.4	L -structures on valuation trees	225
9.4.1	Envelopes	228
9.5	The main results	231

9.5.1	Unary relations	231
9.5.2	Non- L -hypergraphs	232
9.5.3	Forbidding structures	233
9.6	Infinite big Ramsey degrees?	234
9.7	Conclusion	238
9.7.1	Infinite big Ramsey degrees	238
9.7.2	Small Ramsey degrees and the partite lemma	240
10	Characterisation of the big Ramsey degrees of the generic partial order	242
10.1	Introduction	242
10.2	Preliminaries	247
10.2.1	Relational structures	247
10.2.2	Trees	248
10.2.3	Parameter words	248
10.3	Tree of 1-types	248
10.4	Poset-diaries and level structures	252
10.5	A poset-diary coding \mathbf{P}	254
10.6	Interesting levels and sub-diaries	255
10.7	Upper bounds	258
10.8	The lower bound	262
10.9	Concluding remarks	266
10.9.1	Comparison to big Ramsey degrees of the order of rationals	266
10.9.2	The triangle-free graph	266
11	Big Ramsey degrees and forbidden cycles	268
11.1	Introduction	268
11.2	Preliminaries	269
11.3	Proof of Theorem 11.1.1	271
12	Type-respecting amalgamation and big Ramsey degrees	274
12.1	Introduction	274
	Bibliography	281

Preface

Let \mathbf{A} be a structure. We say that \mathbf{A} is homogeneous if every isomorphism between finite substructures of \mathbf{A} extends to an automorphism of \mathbf{A} . Equivalently, the orbits of the action $\text{Aut}(\mathbf{A}) \curvearrowright A^n$ coincide with the isomorphism types of (enumerated) substructures of \mathbf{A} on n vertices for every $n \geq 1$.

A permutation group is an automorphism group of a countable structure with vertex set \mathbb{N} if and only if it is a closed subgroup of $\text{Sym}(\mathbb{N})$. To every such group one can assign a homogeneous structure by adding relations for the orbits of its action on n -tuples.

It turns out that various properties of automorphism groups of homogeneous structures are equivalent to combinatorial properties of their ages, that is, classes of all finite substructures of the given homogeneous structure. These correspondences, and mainly their combinatorial parts, have been in the centerpiece of my mathematical interests since my undergraduate studies.

The main substance of this PhD thesis consists of the following papers:

David M. Evans, Jan Hubička, Matěj Konečný, Yibei Li, and Martin Ziegler. Simplicity of the automorphism groups of generalised metric spaces. *Journal of Algebra*, 584:163–179, 2021.

Jan Hubička, Matěj Konečný, and Jaroslav Nešetřil. All those EPPA classes (strengthenings of the Herwig–Lascar theorem). *Transactions of the American Mathematical Society*, 375(11):7601–7667, 2022.

David M. Evans, Jan Hubička, Matěj Konečný, and Jaroslav Nešetřil. EPPA for two-graphs and antipodal metric spaces. *Proceedings of the American Mathematical Society*, 148:1901–1915, 2020.

Matěj Konečný. Extending partial isometries of antipodal graphs. *Discrete Mathematics*, 343(1):111633, 2020.

Gregory Cherlin, Jan Hubička, Matěj Konečný, and Jaroslav Nešetřil. Ramsey expansions of 3-hypertournaments. In *Extended Abstracts EuroComb 2021*, pages 696–701. Springer, 2021.

Martin Balko, David Chodounský, Jan Hubička, Matěj Konečný, and Lluís Vena. Big Ramsey degrees of 3-uniform hypergraphs are finite. *Combinatorica*, 42(2):659–672, 2022.

Samuel Braunfeld, David Chodounský, Noé de Rancourt, Jan Hubička, Jamal Kawach, and Matěj Konečný. Big Ramsey degrees and infinite languages. Submitted, *arXiv:2301.13116*, 2023.

Martin Balko, David Chodounský, Natasha Dobrinen, Jan Hubička, Matěj Konečný, Lluís Vena, and Andy Zucker. Characterisation of the big Ramsey degrees of the generic partial order. Submitted, *arXiv:2303.10088*, 2023.

Martin Balko, David Chodounský, Jan Hubička, Matěj Konečný, Jaroslav Nešetřil, and Lluís Vena. Big Ramsey degrees and forbidden cycles. In *Extended Abstracts EuroComb 2021*, pages 436–441. Springer, 2021.

Andres Aranda, Samuel Braunfeld, David Chodounský, Jan Hubička, Matěj Konečný, Jaroslav Nešetřil, and Andy Zucker. Type-respecting amalgamation and big Ramsey degrees. Extended abstract accepted to EuroComb 2023, *arXiv:2303.12679*, 2023.

All these papers satisfy the following conditions: I am their co-author, I believe that my contribution was substantial, they fit into the narrative of this thesis and, altogether, they still keep this thesis at least a remotely reasonable length. In total, over the course of my academic career I have so far co-authored the following papers, extended abstracts and preprints: [ABC⁺23, ABWH⁺17a, ABWH⁺21, ABWH⁺17b, BCD⁺21a, BCD⁺21b, BCD⁺23a, BCdR⁺23, BCH⁺21a, BCH⁺19, BCH⁺22, BdRHK23, CHKN21, EHK⁺21, EHKN20, HJKS19, HKK18, HKN19a, HKN19b, HKN21, HKN22, Kon20]

This thesis is organized into three parts:

1. It starts with Chapter 1 which provides some background, outlines the relevant areas on which the thesis focuses, and tries to hint at some of their interplays. It does not aspire to provide a full overview, to give a precise historical account, nor to describe things in full formality and precision. Instead, I tried to pretend that I am giving a series of seminar talks on the topics of my choice whose goal is to introduce the audience to my very specific interests and to try to transfer some of my intuition. The reader which I had in mind when writing the chapter is someone with certain familiarity with homogeneous structures who is interested in the more specific problems which I discuss. Nevertheless, I tried to keep it mostly self-contained and accessible.

It would be conceited to think that my first attempt at this kind of mixture of lecture notes and a survey will be overly successful. It turned out to be very useful to me for organizing my thoughts, and I will be happy if a part of it happens to be interesting also to someone else. I will be grateful for any remarks and comments which will certainly be useful if at some point I happen to decide to use parts of this introduction as a basis for a survey or lecture notes of some kind (even if the comment tells me to abandon such an idea).

The referees of this thesis can rest assured that Chapter 1 contains no original unpublished results. It is the third part which I consider to be the main scientific content of this thesis, with Chapters 1 and 2 being only supplementary.

2. In Chapter 2 I tried to collect various relevant questions, problems and conjectures from (not only) my papers, as well as adding a few previously unpublished ones. I plan to keep an occasionally updated version on my website.

3. The third part (Chapters 3–12) consists of the aforementioned collection of papers. The differences between the presentation here and in the versions that are available on arXiv at the time of writing this thesis are minor and mainly due to formatting requirements of the university. The bibliographies have been merged and moved to the Bibliography section at the end of the thesis. The reader is strongly advised to consult arXiv or the journal version for the most recent versions of the papers as these will be the first places where possible updates or errata will be published. For the same reason, I kindly ask the reader to primarily cite the paper versions of results or questions.

From now on, I switch from “I” to “we”, reserving the singular “I” for occasions where I want to emphasize that I am expressing my personal opinion (mostly in Chapter 2).

1. Background

We will use the set-theoretic convention that $n = \{0, \dots, n-1\}$ and $\omega = \{0, \dots\}$.

1.1 Structures and morphisms

We start by recalling some standard model-theoretic notions regarding structures with relations and functions with a small variation that our functions are partial.

A *language* L is a collection $L = L_{\mathcal{R}} \cup L_{\mathcal{F}}$ of relation symbols $R \in L_{\mathcal{R}}$ and function symbols $F \in L_{\mathcal{F}}$ each having associated arities. For relations, the arity is denoted by $a(R) > 0$, for functions $a(F)$ is the arity of the domain. In this chapter, the range will always have arity one. A language L is *relational* if $L_{\mathcal{F}} = \emptyset$.

An L -*structure* \mathbf{A} is then a tuple $(A, \{R_{\mathbf{A}}\}_{R \in L_{\mathcal{R}}}, \{F_{\mathbf{A}}\}_{F \in L_{\mathcal{F}}})$, where A is the *vertex set*, $R_{\mathbf{A}} \subseteq A^{a(R)}$ is an interpretation of R for each $R \in L_{\mathcal{R}}$ and $F_{\mathbf{A}}: A^{a(F)} \rightarrow A$ is a partial function for each $F \in L_{\mathcal{F}}$. We denote by $\text{Dom}(F_{\mathbf{A}})$ the domain of F (i.e. the set of tuples of vertices of \mathbf{A} for which F is defined). An L -structure is *finite* if its vertex set is finite. Notationally, we distinguish structures from unstructured sets by typesetting structures in bold font. When the language L is clear from the context, we will use it implicitly.

Let \mathbf{A} and \mathbf{B} be L -structures. A *homomorphism* $f: \mathbf{A} \rightarrow \mathbf{B}$ is a mapping $f: A \rightarrow B$ satisfying for every $R \in L_{\mathcal{R}}$ and for every $F \in L_{\mathcal{F}}$ the following three statements:

- (a) $(x_1, \dots, x_{a(R)}) \in R_{\mathbf{A}} \Rightarrow (f(x_1), \dots, f(x_{a(R)})) \in R_{\mathbf{B}}$;
- (b) $f(\text{Dom}(F_{\mathbf{A}})) \subseteq \text{Dom}(F_{\mathbf{B}})$; and
- (c) For every $(x_1, \dots, x_{a(F)}) \in \text{Dom}(F_{\mathbf{A}})$ we have that $f(F_{\mathbf{A}}(x_1, \dots, x_{a(F)})) = F_{\mathbf{B}}(f(x_1), \dots, f(x_{a(F)}))$.

For a subset $A' \subseteq A$ we denote by $f(A')$ the set $\{f(x); x \in A'\}$ and by $f(\mathbf{A})$ the homomorphic image of a structure \mathbf{A} .

If f is an injective homomorphism, it is a *monomorphism*. A monomorphism is an *embedding* if for every $R \in L_{\mathcal{R}}$ and $F \in L_{\mathcal{F}}$ the following holds:

- (a) $(x_1, \dots, x_{a(R)}) \in R_{\mathbf{A}} \iff (f(x_1), \dots, f(x_{a(R)})) \in R_{\mathbf{B}}$, and,
- (b) $(x_1, \dots, x_{a(F)}) \in \text{Dom}(F_{\mathbf{A}}) \iff (f(x_1), \dots, f(x_{a(F)})) \in \text{Dom}(F_{\mathbf{B}})$.

If f is a bijective embedding then it is an *isomorphism* and we say that \mathbf{A} and \mathbf{B} are *isomorphic*. An isomorphism $\mathbf{A} \rightarrow \mathbf{A}$ is called *automorphism*. If the inclusion $\mathbf{A} \subseteq \mathbf{B}$ is an embedding then \mathbf{A} is a *substructure* of \mathbf{B} . For \mathbf{A} and \mathbf{B} structures, we denote by $\binom{\mathbf{B}}{\mathbf{A}}$ the set of all embeddings of \mathbf{A} to \mathbf{B} . Note that while for relational languages every set $A \subseteq B$ gives a substructure of \mathbf{B} , it does not hold in general for languages with functions (we need A to be *closed* on functions).

1.2 Amalgamation

Given structures \mathbf{B}_1 and \mathbf{B}_2 , a structure \mathbf{C} is a *joint embedding* of \mathbf{B}_1 and \mathbf{B}_2 if there are embeddings $\beta_1: \mathbf{B}_1 \rightarrow \mathbf{C}$ and $\beta_2: \mathbf{B}_2 \rightarrow \mathbf{C}$.

Given structures \mathbf{A} , \mathbf{B}_1 , \mathbf{B}_2 , and embeddings $\alpha_1: \mathbf{A} \rightarrow \mathbf{B}_1$ and $\alpha_2: \mathbf{A} \rightarrow \mathbf{B}_2$, a structure \mathbf{C} is an *amalgamation of \mathbf{B}_1 and \mathbf{B}_2 over \mathbf{A}* (with respect to α_1 and α_2) if there are embeddings $\beta_1: \mathbf{B}_1 \rightarrow \mathbf{C}$ and $\beta_2: \mathbf{B}_2 \rightarrow \mathbf{C}$ such that $\beta_1 \circ \alpha_1 = \beta_2 \circ \alpha_2$. The amalgamation \mathbf{C} is *strong* if $\beta_1(B_1) \cap \beta_2(B_2) = \beta_1(\alpha_1(A))$. It is *free* if it is strong, $C = \beta_1(B_1) \cup \beta_2(B_2)$ and moreover whenever a tuple $\bar{x} \in C^n$ is in some relation of \mathbf{C} or some function is defined for \bar{x} then either $\bar{x} \in \beta_1(B_1)^n$ or $\bar{x} \in \beta_2(B_2)^n$.

A structure is *irreducible* if it is not a free amalgamation of its proper substructures. The following observation is immediate:

Observation 1.2.1. *If L is a relational language and \mathbf{A} is an L -structure then \mathbf{A} is irreducible if and only if for every $x, y \in A$ there exists a tuple $\bar{x} \in A^n$ and a relation $R \in L$ such that $x, y \in \bar{x}$ and $\bar{x} \in R_{\mathbf{A}}$. \square*

Given structures \mathbf{A} and \mathbf{B} , a homomorphism $f: \mathbf{A} \rightarrow \mathbf{B}$ is a *homomorphism-embedding* if whenever $\mathbf{C} \subseteq \mathbf{A}$ is irreducible, $f|_{\mathbf{C}}$ is an embedding. This notion has been first isolated by Hubička and Nešetřil [HN19]. For a set \mathcal{F} of L -structures, we denote by $\text{Forb}_{\text{he}}(\mathcal{F})$ the set of all finite and countable L -structures \mathbf{A} such that there is no $\mathbf{F} \in \mathcal{F}$ with a homomorphism-embedding $\mathbf{F} \rightarrow \mathbf{A}$.

Later in the combinatorial constructions, we will create structures by “gluing” many copies of an irreducible structure over its substructures in complex ways. For example, we might be gluing finite metric spaces over their subspaces, with the resulting structure being a graph with edges labelled by distances, but not every pair of vertices having a defined distance. In this setting, homomorphism-embeddings are the correct morphisms as they preserve all embeddings of the glued copies while being able to “fill in the holes” to obtain a metric spaces (if possible). The following folklore result has been used by Solecki [Sol05] and Nešetřil [Neš07] in their proofs of EPPA resp. Ramsey property for the class of metric spaces.

Proposition 1.2.2. *Let \mathbf{G} be a finite graph with edges labelled by positive real numbers. Then \mathbf{G} has a homomorphism-embedding into a metric space (understood as a complete graph with edges labelled by positive real numbers) if and only if no non-metric cycle has a homomorphism-embedding into \mathbf{G} .*

Here, a *non-metric cycle* is a graph cycle with edges labelled by positive real numbers such that one label is larger than the sum of all the other labels.

This proposition is proved by analysing the *shortest path completion* of \mathbf{G} , that is, the metric space on G where the distance between any two vertices is the length of the shortest path connecting them in \mathbf{G} (measured as the sum of the labels), or some large enough number if the vertices lie in different connected components of \mathbf{G} . Generalisations of the shortest path completion have been extensively studied by the author in his Bachelor and Master theses [Kon18, Kon19] as well as the author and co-authors in several papers [ABWH⁺17a, ABWH⁺21, ABWH⁺17b, HKN21, HKN18]. It also plays a key role in the paper [EHK⁺21] which is included as Chapter 3 of this thesis.

1.2.1 Fraïssé's theorem

We say that a class \mathcal{C} of finite L -structures has the *joint embedding property* (*JEP*) if for every $\mathbf{B}_1, \mathbf{B}_2 \in \mathcal{C}$ there exists $\mathbf{C} \in \mathcal{C}$ which is a joint embedding of \mathbf{B}_1 and \mathbf{B}_2 . It has the *amalgamation property* if for every $\mathbf{A}, \mathbf{B}_1, \mathbf{B}_2 \in \mathcal{C}$ and embeddings $\alpha_1: \mathbf{A} \rightarrow \mathbf{B}_1$ and $\alpha_2: \mathbf{A} \rightarrow \mathbf{B}_2$, there is $\mathbf{C} \in \mathcal{C}$ which is an amalgamation of \mathbf{B}_1 and \mathbf{B}_2 over \mathbf{A} with respect to α_1 and α_2 . We say that \mathcal{C} has the *strong amalgamation property* if one can find \mathbf{C} which is a strong amalgamation of \mathbf{B}_1 and \mathbf{B}_2 over \mathbf{A} with respect to α_1 and α_2 , and the *free amalgamation property* is defined analogously.

We call \mathcal{C} a *Fraïssé class* if it consists of finite structures, is hereditary (i.e. closed under substructures) and closed under isomorphisms, has the joint embedding property and the amalgamation property, and contains only countably many members up to isomorphism. A structure \mathbf{M} is *homogeneous* if for every finite $\mathbf{A}, \mathbf{B} \subseteq \mathbf{M}$ and every isomorphism $g: \mathbf{A} \rightarrow \mathbf{B}$ there is an automorphism f of \mathbf{M} with $g \subseteq f$. We say that \mathbf{M} is a *Fraïssé structure* if it is countable, homogeneous, and *locally finite* (that is, for every finite $X \subseteq M$ there is a finite substructure $\mathbf{Y} \subseteq \mathbf{M}$ with $X \subseteq Y$). Given a structure \mathbf{A} , its *age* $\text{Age}(\mathbf{A})$ is the class of all finite structures which embed into \mathbf{A} . Fraïssé classes and structures are called after Fraïssé who in the early 1950s proved the following theorem.

Theorem 1.2.3 (Fraïssé [Fra53]).

1. Let \mathbf{M} be a Fraïssé structure. Then $\text{Age}(\mathbf{M})$ is a Fraïssé class.
2. For every Fraïssé class \mathcal{C} there is a Fraïssé structure \mathbf{M} such that $\text{Age}(\mathbf{M}) = \mathcal{C}$. Furthermore, if \mathbf{N} is a countable homogeneous L -structure such that $\text{Age}(\mathbf{N}) = \mathcal{C}$, then \mathbf{M} and \mathbf{N} are isomorphic.

We call the structure \mathbf{M} from the second point the *Fraïssé limit* of \mathcal{C} .

Fraïssé structures (or, more generally, homogeneous structures) are important objects in model theory, see e.g. the survey on homogeneous structures by Macpherson [Mac11]. They are rather rare: For example, the only countably infinite homogeneous graphs are the random graph, the generic K_n -free graphs, disjoint unions of cliques of the same size, and their complements, as classified by Lachlan and Woodrow in 1980 [LW80]. The goal of the Lachlan–Cherlin classification programme of homogeneous structures is to provide other classifications like the Lachlan–Woodrow one. So far, there have been several successful attempts (e.g. [Che98, Che22]). In Section 1.4, we will see some examples of (families of) homogeneous structures with varying combinatorial properties.

1.3 Automorphism groups

Let \mathbf{M} be a countable (that is, finite or countably infinite) structure and let $G = \text{Aut}(\mathbf{M})$ be its automorphism group. By identifying the vertex set of \mathbf{M} with (a subset of) \mathbb{N} , we can view G as a subgroup of $\text{Sym}(\mathbb{N})$. As we shall see in this chapter, one can say much more than this. For simplicity, we will assume that the vertex set is always the set of natural numbers, since we are always working up-to-isomorphism, the contents of this chapter hold for any countable set.

The group $\text{Sym}(\mathbb{N})$ is a subset of the space $\mathbb{N}^{\mathbb{N}}$ of all functions $\mathbb{N} \rightarrow \mathbb{N}$, which can be endowed with the product topology (called, in this context, the *pointwise-convergence topology*). It is easy to see that $\text{Sym}(\mathbb{N})$ is closed in $\mathbb{N}^{\mathbb{N}}$ and that it is a topological group with the inherited topology. In fact, it is a *Polish group* (i.e. the topology is Hausdorff, separable and completely metrizable). It turns out that automorphism groups of Fraïssé structures with vertex set \mathbb{N} coincide with closed subgroups of $\text{Sym}(\mathbb{N})$:

Fact 1.3.1. *Let \mathbf{M} be a structure with vertex set \mathbb{N} . Then $\text{Aut}(\mathbf{M})$ is a closed subgroup of $\text{Sym}(\mathbb{N})$. Conversely, let $G \leq \text{Sym}(\mathbb{N})$ be closed. Then there exists a Fraïssé structure \mathbf{M} with vertex set \mathbb{N} such that $G = \text{Aut}(\mathbf{M})$.*

In light of this fact, it is perhaps not that surprising that various properties of closed subgroups of $\text{Sym}(\mathbb{N})$ can be equivalently phrased as combinatorial properties of Fraïssé structures. And the definition of pointwise-convergence topology (for example, pointwise stabilisers of finite sets form a system of neighbourhoods of the identity) justifies that, in fact, it often boils down to finite combinatorics on Fraïssé classes. We will see several examples of this in this thesis.

An important property in model theory is ω -categoricity. A complete theory is ω -categorical if it has only one countable model up to isomorphism. A closed $G \leq \text{Sym}(\mathbb{N})$ is *oligomorphic* if its coordinate-wise action on \mathbb{N}^k has only finitely many orbits for every k . The following is known as (a part of) the Ryll-Nardzewski theorem, although it has been proved independently by various authors (see e.g. [Mac11]):

Theorem 1.3.2. *Let \mathbf{M} be a structure with vertex set \mathbb{N} . Then the theory of \mathbf{M} is ω -categorical if and only if $\text{Aut}(\mathbf{M})$ is oligomorphic. In particular, every homogeneous structure in a finite relational language has an ω -categorical theory.*

1.4 A modest bestiary of homogeneous structures and amalgamation classes

Several of the included papers prove various combinatorial properties of various homogeneous structures. The added benefit of having them together in a thesis is that it allows us to give a more global view of the interplay (or lack thereof) between these properties. Hence, one needs to have good examples at hand to demonstrate this. In this section we will briefly introduce various (classes of) homogeneous structures. The aim is neither to give a comprehensive overview, nor to describe the structures in depth. Wherever it makes sense, we will try to give references to further material where the reader can learn more about the respective structure. Note that in this section as well as in the rest of the thesis we will speak interchangeably about Fraïssé classes and their respective Fraïssé limits whenever there is no risk of confusion (for example, we might say that a Fraïssé structure has some property which, in fact, is a property of its age, or vice versa).

The infinite set with no structure The simplest example of a homogeneous structure is the infinite set with no relations and functions (that is, $L = \emptyset$). It is in fact a Fraïssé structure, its age is the class of all finite sets, it has the free amalgamation property and its automorphism group is $\text{Sym}(\mathbb{N})$.

(\mathbb{Q}, \leq) The probably next simplest Fraïssé structure is the set of rationals with its order. Its language is $L = \{\leq\}$ consisting of one binary relation which is reflexive, antisymmetric and transitive, and its age is the class of all finite linear orders. It has the strong amalgamation property but not the free amalgamation property.

The Rado graph The *Rado graph*, also called the (*countable*) *random graph* is a Fraïssé structure in the language $L = \{E\}$ consisting of one binary relation which is symmetric and irreflexive. It is the Fraïssé limit of the class of all finite graphs, which has free amalgamation property. It is described (up to isomorphism) by the *extension property*: for every pair of disjoint finite sets of vertices, there exists a vertex connected to every member of the first set and no member of the second. For its history, see e.g. [Cam15]. From a certain point of view, this is the simplest “random” homogeneous structure (random meaning that there are many isomorphism types of finite substructures on n vertices).

Note that in fact every homogeneous structure \mathbf{M} satisfies the following general form of extension property: For every pair of structures $\mathbf{A} \subseteq \mathbf{B} \in \text{Age}(\mathbf{M})$ with \mathbf{A} and \mathbf{B} finite and for every embedding $f: \mathbf{A} \rightarrow \mathbf{M}$, there exists an embedding $g: \mathbf{B} \rightarrow \mathbf{M}$ such that $f \subseteq g$. If \mathbf{M} is ω -categorical, it describes it to isomorphism.

The ordered Rado graph This is the free interposition of the Rado graph and (\mathbb{Q}, \leq) , in other words, the Fraïssé limit of the class of all finite linearly ordered graph. Its language is $L = \{E, \leq\}$.

The (random) k -uniform hypergraph Generalising the Rado graph for higher uniformities, we get, for every $k \geq 2$, the Fraïssé limit of the class of all finite k -uniform hypergraphs.

Fraïssé limits of free amalgamation classes and their ordered variants

Generalising this further, if \mathcal{C} is a Fraïssé class with the free amalgamation property, one also has its Fraïssé limit. Examples include graphs, directed graphs, hypergraphs, K_n -free graphs for every $n \geq 2$, hypergraphs omitting a complete hypergraph, classes of all finite relational structures in a given language, but also, for example, the class of all finite structures in a language containing one partial unary function, or one unary function, or the class of all finite structures in a language containing one (partial) unary function such that the closure of every vertex has size at most some fixed n . In general, given a language L and a (not necessarily finite) family \mathcal{F} of finite irreducible L -structures, the class of all finite \mathcal{F} -free structures is a free amalgamation class. Similarly as for the Rado graph, one can consider linearly ordered variants of these – note that they no longer have the free amalgamation property.

Various metric spaces A metric space can equivalently be seen as a (possibly infinite) complete graph with edges labelled by positive real numbers which omits *non-metric triangles*, that is, triangles where the label of one edge is larger than the sum of the labels of the other two edges. In turn, edge-labelled (complete) graphs can be seen as relational structures where there is one binary relation for every label, all relations are symmetric, irreflexive, and every pair of vertices is in at most (exactly) one relation.

Given a subset $S \subseteq \mathbb{R}^+$, one can consider *S -metric spaces*, that is, metric spaces where all the distances come from $S \cup \{0\}$. If S is countable, there are only countably many finite S -metric spaces, and so the standard Fraïssé theory works for constructing homogeneous S -metric spaces. If S is uncountable, one can often approximate it by a countable dense subset, do Fraïssé theory on the subset and then take the completion of the countable homogeneous metric spaces.

Sauer [Sau13a, Sau13b] classified those sets S for which there exists a complete separable homogeneous S -metric space universal for all separable S -metric spaces, often called the *S -Urysohn space* after Urysohn who, in 1927, constructed the \mathbb{R}^+ -Urysohn space. [Ury27] (Note that this predates the Fraïssé theorem.) Sauer’s condition for the existence of an S -Urysohn space (the *four values condition*) is, in the case when $S \cup \{0\}$ is (topologically closed), equivalent to the following operation being associative: $a \oplus_S b = \sup\{c \in S : c \leq a + b\}$. Notable examples of such sets are:

1. $S = \mathbb{R}^+$, the \mathbb{R}^+ -Urysohn space is called simply the the *Urysohn space*.
2. $S = (0, 1]$, the $(0, 1]$ -Urysohn space is called the *Urysohn sphere*, because it is isomorphic to the substructure induced by the Urysohn space on the set of all points in distance $\frac{1}{2}$ from an arbitrary fixed point.
3. $S = \mathbb{Q}^+$ resp. $S = (0, 1] \cap \mathbb{Q}$ are the rational variants of the Urysohn space resp. Urysohn sphere and, in fact, one constructs the Urysohn space resp. sphere as the completion of the Fraïssé limit of the class of all finite S -metric spaces for $S = \mathbb{Q}^+$ resp. $S = (0, 1] \cap \mathbb{Q}$.

4. $S = \{1, \dots, n\}$, or $S = \{1, \dots\}$ which is sometimes called the *integer Urysohn space*.
5. $S = \{2^1 - 1, 2^2 - 1, \dots, 2^n - 1\}$ and $S = \{2^1 - 1, 2^2 - 1, \dots\}$ which are the *ultrametric spaces*. Note that every distance is larger than the sum of any two smaller distances. This means that distances simply describe refining equivalence relations.

Metrically homogeneous graphs In 1980, Lachlan and Woodrow [LW80] classified all countably infinite homogeneous graphs (they are the Rado graph, its K_n -free analogues, disjoint unions of cliques of the same size, and complements of all of these). Note that the connected components of all of these have diameter at most two. This has a good reason – a graph of diameter at least three cannot be non-edge transitive (because there are multiple types of non-edges based on the distance of the given two vertices), hence cannot be homogeneous. However, one can only consider isometric embeddings (i.e. embeddings preserving the graph distance), or equivalently, look at the path-metric spaces associated to the given graphs. Here, the distance of two vertices in the path-metric space is the number of edges of the shortest path connecting the two vertices (or infinity, but let us only consider connected graphs for now).

Extending the Lachlan–Woodrow classification in this direction, Cherlin provided a catalogue of metrically homogeneous graphs [Che11, Che22] which is conjectured to be complete with a recent purported (yet unpublished) proof.¹ This catalogue is rich and contains, in particular, all integer Urysohn spaces, integer Urysohn spaces omitting short odd cycles of edges labelled with 1, certain antipodal structures which we will see later in this section and many more. This catalogue has been the subject of the author’s Bachelor thesis [Kon18], it was an important motivation for the author’s Master thesis [Kon19] and it was also a subject of several of the author’s papers [ABWH⁺17b, ABWH⁺21, ABWH⁺17a, HKK18, EHK⁺21] including one included in this thesis in Chapter 3. We shall not introduce the catalogue here in detail, it suffices to know that it provides a rich family of examples of homogeneous structures in binary symmetric languages.

Generalised metric spaces Motivated by results on metrically homogeneous graphs, Hubička, Nešetřil, and the author axiomatized under which conditions do certain classes of generalised metric spaces, where distances come from a partially ordered commutative semigroup, have various combinatorial properties. (This was the topic of the author’s Master thesis [Kon19].) This framework covers most of the known amalgamation classes in a binary symmetric language. For this exposition we do not need to cover the exact definition, knowing that they generalise all metric spaces and all (primitive) metrically homogeneous graphs is enough. These structures have been also the centerpiece of several of the author’s papers [EHK⁺21, HKN21, HKN18] including one included in this thesis in Chapter 3.

The generic poset The class of all finite partial orders is a Fraïssé class whose limit is the generic poset. It is one of the natural examples that one needs to look

¹This is claimed on: <https://sites.math.rutgers.edu/~cherlin/Paper/inprep.html>

at to test various hypotheses as it often exhibits very different behavior compared to, for example, generalised metric spaces or free amalgamation classes. (One of the reasons is that there are, in fact, not that many known non-free strong amalgamation classes which are not generalised metric spaces, see also below.) We will see an instance of this in Chapter 10 where [BCD⁺23a] is included. (See also [BCD⁺21a].)

The generic tournament, n -partite tournament and semigeneric tournament In 1998, Cherlin provided a classification of homogeneous directed graphs [Che98], providing in particular several examples of non-free strong amalgamation classes which are not generalised metric spaces. In this thesis, we will be mainly concerned with three items from this list: The class of all finite tournaments, the class of all finite n -partite tournaments and the class of all finite semigeneric tournaments. A tournament is a structure with one irreflexive anti-symmetric total binary relation. An n -partite (for $n \in \{1, \dots, \infty\}$) tournament is a structure with one irreflexive antisymmetric binary relation such that there exists a partition into at most n independent sets such that whenever u and v are vertices from different parts, exactly one of uv and vu is in the relation. A semigeneric tournament is an ∞ -partite tournament such that whenever u, u' are from one part and v, v' from another part then there are an even number of edges going from $\{u, u'\}$ to $\{v, v'\}$. This implies that if one fixes one part, each part splits into two equivalence classes with the direction of edges to/from the fixed being determined just by the (finer) equivalence classes. These have been studied by the author in the context of EPPA in [HJKS19], the full version is in preparation [HJKS23].

The generic two-graph, and antipodal spaces A two-graph is a 3-uniform hypergraph such that there are an even number of hyperedges induced on every quadruple of vertices. Two-graphs have been introduced by G. Higman and studied extensively since the 1970's [Sei73, Cam99].

An antipodal metric space of diameter 3 is a $\{1, 2, 3\}$ -metric space such that edges of length 3 form a perfect matching and there is no triangle with distances 3, 2, 2. Equivalently, one can start with two copies of an arbitrary graph G with edges representing distance 1 and non-edges representing distance 2, connect every vertex with its copy by distance 3, and, if $u, v \in G$ and u' is the copy of u , putting $d(u', v) = 3 - d(u, v)$. Equivalently, antipodal metric spaces of diameter 3 describe *double covers* of complete graphs

In general, Cherlin has defined various classes of antipodal metrically homogeneous graphs (and hence metric spaces) of various diameters; their combinatorial properties are, however, given by combining the combinatorial properties of one *pode* (which is a metric space of diameter smaller by one from the catalogue of metrically homogeneous graphs) with the *switching* properties of antipodal metric spaces of diameter 3.

Given an antipodal metric space of diameter 3, one can define a two-graph whose vertices are the edges of length 3 such that three edges are in a hyperedge if and only if, if we look at distances 1 between the six endpoints of the edges, these distances form a graph 6-cycle. Otherwise they form two triangles, which corresponds to a non-hyperedge. Conversely, one can create an antipodal metric

space of diameter 3 from a two-graph by inverting this procedure.

These structures are also closely connected to graphs with *switching automorphisms* (two-graphs represent their isomorphism classes), and represent many interesting and exotic behaviors, making them important examples for our studies. These structures have been studied in the author’s papers [ABWH⁺17b, EHKN20, Kon20], where [EHKN20] is included in Chapter 5 and [Kon20] is included in Chapter 6.

Hypertournaments Answering a demand for non-free strong amalgamation classes in non-binary languages, Cherlin provided some classes of n -hypertournaments [Che] which then served as a very important example as well as source of an open problem. The paper [CHKN21] included in Chapter 7 deals with Ramsey properties of these structures.

An n -hypertournament is a structure with one irreflexive n -ary relation such that the automorphism group of every substructure on n vertices is the alternating group $\text{Alt}(n)$. For $n = 2$ we get tournaments. While at the first glance it may seem that graphs have all the symmetries on pairs and tournaments have none on n -substructures, hence hypertournaments should also have none, having “almost all” symmetries turns out to be combinatorially more convenient because one always has only two choices of a relation on a set of n vertices. That is, given distinct vertices v_1, \dots, v_n it always holds that exactly one of (v_1, v_2, \dots, v_n) and (v_2, v_1, \dots, v_n) is in the relation, and this choice determines all the other relations on $\{v_1, \dots, v_n\}$. This means that there is a 1-to-1 correspondence between linearly ordered n -hypertournaments and linearly ordered n -uniform hypergraphs: A hyperedge on $v_1 < \dots < v_n$ means that (v_1, v_2, \dots, v_n) are in the hypertournament relation. This correspondence is often useful for drawing concrete hypertournaments.

For $n = 3$, the hypertournament relation picks, for every triple, one of its two possible cyclic orientations. There are three isomorphism types of 3-hypertournaments on four vertices: C_4 , where the cyclic orientations of all four triples agree with one global cyclic orientation of the four vertices, O_4 which one gets from C_4 by reversing the orientation of one (arbitrary) triple, and H_4 , which is the unique homogeneous 3-hypertournament on four vertices; its automorphism group is $\text{Alt}(4)$. Equivalently, if we have four vertices v_1, v_2, v_3, v_4 ordered cyclically in this order, in H_4 either $v_1v_2v_3$ and $v_1v_3v_4$ agree with this orientation, or $v_1v_2v_4$ and $v_2v_3v_4$ agree with this orientation.

One can then ask which subsets of $\{C_4, O_4, H_4\}$ one can forbid in order to obtain a Fraïssé class. A simple Ramsey argument shows that every large enough 3-hypertournament contains a copy of C_4 ; and all the four remaining subsets of $\{C_4, O_4, H_4\}$ give rise to a Fraïssé class. Particularly interesting are two of them: The O_4 -free (or *even*) 3-hypertournaments are, when ordered, bi-interpretable with ordered two-graphs (more precisely, restricting the 1-to-1 correspondence between linearly ordered 3-hypertournaments and linearly ordered 3-uniform hypergraphs to the O_4 -free case, one gets the class of linearly ordered two-graphs).

The other particularly interesting case are the H_4 -free 3-hypertournaments. They are interesting for two reasons. The first one is that one can analogously define H_{n+1} -free n -hypertournaments and these always form a Fraïssé class. In the $n = 2$ case, H_3 is the cyclic tournament, and so H_3 -free tournaments are simply

the linear orders. Hence, H_4 -free 3-hypertournaments can be seen as some higher-order linear orders. And, indeed, they have been discovered independently in this setting by Bergfalk [Ber21] in his study of set theoretical higher-order Todorćević walks. The other reason is that the H_4 -free 3-hypertournaments are the only class of 3-hypertournaments from Cherlin's list for which we have been unable to find a Ramsey expansion in [CHK21], so it is possible that they can show some limitations of the current methods for proving Ramseyness (see Section 1.7 for more details and discussions).

Hall's universal locally finite group To give an example of very different flavour, note that the class of all finite groups is a Fraïssé class: To prove the amalgamation property, one can embed groups into a large enough symmetric group and take a quotient (see e.g. [Sin17]). The Fraïssé limit of this is called Hall's universal locally finite group.

1.5 Stationary independence relations and simple automorphism groups

The first group property which is, for automorphism groups of Fraïssé structures, closely connected to finite combinatorics, is simplicity (recall that a group is simple if it has no proper normal subgroups). This area has been initiated by Truss [Tru85] in 1985 who proved that the automorphism group of the countable random graph is simple. In 2011, Macpherson and Tent [MT11] proved that the automorphism groups of Fraïssé limits of free amalgamation classes are simple. This was followed by two papers of Tent and Ziegler [TZ13b, TZ13a] where they proved that the isometry group of the Urysohn space modulo bounded isometries (i.e. isometries f with a finite bound on the distance between x and $f(x)$) is simple, and that the isometry group of the Urysohn sphere is simple. For more details, see the PhD thesis of Li on this topic [Li20b].

The following definition of a relation on finite substructures of a Fraïssé structure is key for the proofs. Note that here it is purposefully stated in a weaker form for Fraïssé structures only in order to avoid having to give more definitions. A proper statement is given in Chapter 3 which contains the paper [EHK⁺21].

Definition 1.5.1 (Stationary Independence Relation). Let \mathbf{M} be a relational Fraïssé structure. A ternary relation \perp on finite subsets of \mathbf{M} is called a *stationary independence relation* (SIR, with $A \perp_C B$ being pronounced “ A is independent from B over C ”) if the following conditions are satisfied:

SIR1 (*Invariance*). For every automorphism f of \mathbf{M} , we have $A \perp_C B$ if and only if $f(A) \perp_{f(C)} f(B)$.

SIR2 (*Symmetry*). If $A \perp_C B$, then $B \perp_C A$.

SIR3 (*Monotonicity*). If $A \perp_C (B \cup D)$, then $A \perp_C B$ and $A \perp_{B \cup C} D$.

SIR4 (*Existence*). For every A, B and C , there is A' with $A' \perp_C B$ such that there is an automorphism of \mathbf{M} fixing C pointwise sending A to A' .

SIR5 (*Transitivity*) If $A \perp_C B$ and $A \perp_{B \cup C} B'$, then $A \perp_C B'$.

SIR6 (*Stationarity*) If A and A' are both independent over C from some set B and there is an automorphism of \mathbf{M} fixing C pointwise sending A to A' then there exists an automorphism of \mathbf{M} fixing $B \cup C$ pointwise sending A to A' .

If a Fraïssé structure has a stationary independence relation, this means that one can amalgamate *canonically* in its age. This has been made precise by Aranda, Bradley-Williams, Hubička, Karamanlis, Kompatscher, Pawliuk, and the author in [ABWH⁺17b] where (non-)existence of SIR’s has been determined for Cherlin’s catalogue of metrically homogeneous graphs [Che22]. In particular, while a SIR lives on an infinite structure, it is equivalent to a combinatorial property of a class of finite structures.

Example 1.5.1. Let \mathbf{M} be the Fraïssé limit of some relational free amalgamation class \mathcal{C} . Define a relation \perp on finite subsets of \mathbf{M} by $A \perp_C B$ if and only if whenever \bar{x} is a tuple of vertices from $A \cup B \cup C$ such that \bar{x} is in a relation of \mathbf{M}

then either \bar{x} contains no vertices from $A \setminus C$, or \bar{x} contains no vertices from $B \setminus C$. One can verify that \downarrow is a SIR and the corresponding canonical amalgamation is simply the free amalgamation.

Example 1.5.2. Let \mathbf{M} be the $\{1, \dots, n\}$ -Urysohn space. Define a relation \downarrow on finite subsets of \mathbf{M} by $A \downarrow_C B$ if and only if for every $a \in A$ and every $b \in B$ we have that $d(a, b) = \min(\{n\} \cup \{d(a, c) + d(b, c) : c \in C\})$. One can verify that \downarrow is a SIR. The corresponding canonical amalgamation is given by the *shortest path completion*: Among all possible amalgamations we pick the one which maximizes distances.

Tent and Ziegler [TZ13b] proved that the existence of a SIR together with the existence of the so-called *almost maximally moving automorphisms* is sufficient to prove simplicity of the automorphism group:

Theorem 1.5.1 (Corollary 5.4, [TZ13b]). *Let \mathbf{M} be a countable structure with a stationary independence relation and let g be an automorphism of \mathbf{M} which moves every type over every finite set almost maximally. Then every element of $\text{Aut}(\mathbf{M})$ is a product of sixteen conjugates of g .*

Here, $g \in \text{Aut}(\mathcal{F})$ moves a type p over a finite set A almost maximally if there is a realisation $x \models p$ such that $x \downarrow_A g(x)$.

In Chapter 3 we prove that under the assumption that \mathbf{M} is a homogeneous structure on which its automorphism group acts transitively and \downarrow is a stationary independence relation on \mathbf{M} coming from a sufficiently well-behaved shortest path completion with respect to a partially ordered commutative semigroup, then the automorphism group of \mathbf{M} is simple.

One of the conditions is that, intuitively, the semigroup needs to have a maximum. The reason is that for example the automorphism group of the \mathbb{N}^+ -Urysohn space is not simple as it contains one normal subgroup consisting of all automorphisms g for which there is some $n \in \mathbb{N}$ such that the distance between x and $g(x)$ is at most n for every vertex x . It is likely that, as in the case of the \mathbb{R}^+ -Urysohn space, the quotient by this normal subgroup will be simple.

There are several open problems in this area which the author considers to be interesting and, at the same time, realistic, see Chapter 2.

1.6 EPPA and ample generics

Let \mathbf{A} be a finite structure. We say that a finite structure \mathbf{B} is an EPPA-witness for \mathbf{A} if it contains \mathbf{A} as a substructure and every isomorphism of substructures of \mathbf{A} (a *partial automorphism* of \mathbf{A}) extends to an automorphism of \mathbf{B} . Class \mathcal{C} of finite structures has the *extension property for partial automorphisms* (EPPA, also called the *Hrushovski property*) if it contains an EPPA-witness for every structure in \mathcal{C} . A self-contained reference for topics from this chapter is the PhD thesis of Siniora [Sin17].

Suppose now that \mathcal{C} has EPPA and pick an arbitrary $\mathbf{B}_0 \in \mathcal{C}$. There is an EPPA-witness $\mathbf{B}_1 \in \mathcal{C}$ for \mathbf{B}_0 . Next, we get $\mathbf{B}_2 \in \mathcal{C}$, an EPPA-witness for \mathbf{B}_1 . Continuing this, we get a chain $\mathbf{B}_0 \subseteq \mathbf{B}_1 \subseteq \dots$ of finite structures from \mathcal{C} such that each of them is an EPPA-witness for the previous one (and therefore for all the previous ones). Put $\mathbf{M} = \bigcup \mathbf{B}_i$. Clearly, \mathbf{M} is a countable structure.

Let \mathbf{A} and \mathbf{B} be finite substructures of \mathbf{M} with an isomorphism $g: \mathbf{A} \rightarrow \mathbf{B}$. There is n such that $A \cup B \in \mathbf{B}_n$, hence g is a partial automorphism of \mathbf{B}_n and as such it extends to an automorphism f_{n+1} of \mathbf{B}_{n+1} . In turn, f_{n+1} is a (partial) automorphism of \mathbf{B}_{n+1} and as such it extends to an automorphism f_{n+2} of \mathbf{B}_{n+2} . Consequently, we get a chain $f_{n+1} \subseteq f_{n+2} \subseteq \dots$ of partial automorphisms of \mathbf{M} all of which extend g . It is easy to see that $f = \bigcup f_i$ is an automorphism of \mathbf{M} which extends g . Consequently, \mathbf{M} is a homogeneous structure. If \mathcal{C} has the joint embedding property and at most countably many members up to isomorphism, we can construct the sequence \mathbf{B}_i more carefully to ensure that every member of \mathcal{C} embeds into \mathbf{M} , thereby getting the following folklore observation:

Observation 1.6.1. *Let \mathcal{C} be a hereditary isomorphism-closed class of finite structures containing at most countably many members up to isomorphism. Assume that \mathcal{C} has EPPA and the joint embedding property. Then \mathcal{C} is a Fraïssé class.* \square

In the spirit of this thesis, one can ask what EPPA for \mathcal{C} means for the automorphism group its Fraïssé limit \mathbf{M} . Having the exhaustion $\mathbf{M} = \bigcup \mathbf{B}_i$, it seems that one can approximate automorphisms of \mathbf{M} by (finite) automorphisms of \mathbf{B}_i 's for larger and larger i 's. In other words, “local behavior stays local” – contrast this, for example, with the automorphism group of the oriented doubly infinite path where an automorphism is fully determined by the image of one arbitrary vertex.

Theorem 1.6.2 (Kechris–Rosendal [KR07]). *Let \mathcal{C} be a Fraïssé class with Fraïssé limit \mathbf{M} . Then \mathcal{C} has EPPA if and only if $\text{Aut}(\mathbf{M})$ can be written as the closure of a countable chain of compact subgroups.*

This means that EPPA, a property of a Fraïssé class, has an exact reflection in the automorphism group of its Fraïssé limit (and vice versa).

1.6.1 (Dis)proving EPPA

It turns out that the dynamical formulation of EPPA is not particularly useful for being proved and that one needs to do finite combinatorics. For structures with unary relations only, proving EPPA is simple as each such structure is actually

homogeneous and thus an EPPA-witness for itself. (Consequently, $\text{Sym}(\mathbb{N})$ can be written as the closure of a countable chain of compact subgroups, for example the setwise-stabilisers of larger and larger initial segments of \mathbb{N} .)

Finite linear orders do not have EPPA for the simple reason that they are rigid (only the identity is an automorphism) while certainly having non-trivial partial automorphisms (for example, any vertex can be sent to any other vertex). In general, whenever one can find some linear orders in their structures (e.g. posets) then such structures cannot have EPPA: Consider, in a hypothetical EPPA-witness, some maximal chain containing at least two vertices of the original structure and send one vertex to the other – a hypothetical extending automorphism needs to shift the maximal chain, but then where does its minimal/maximal element go? In fact, finding linear orders is one of the very few methods for disproving EPPA.

EPPA for the class of all finite graphs has been proved by Hrushovski in 1992 [Hru92] (see also Section 1.6.6). By now, there are several constructions of EPPA-witnesses of graphs. Below we present sketches of three of them.

Theorem 1.6.3 (Hrushovski [Hru92]). *The class of all finite graphs has EPPA.*

The (not so) naive attempt Suppose that we have a finite graph \mathbf{A} and want to construct an EPPA-witness \mathbf{B} for \mathbf{A} . The first natural idea might be to simplify the situation and first consider only one partial automorphism φ of \mathbf{A} which we want to extend. Put $D = \text{Dom}(\varphi)$, $R = \text{Range}(\varphi)$ and let \mathbf{D} and \mathbf{R} be the graphs induced by \mathbf{A} on D and R respectively. In order for our hypothetical witness to extend φ , there needs to be a copy of \mathbf{A} from whose point of view R looks like D . The freest structure satisfying this is the free amalgamation of two copies of \mathbf{A} over \mathbf{D} with respect to the identity and φ and extend φ so that the original copy of \mathbf{A} moves to the second one. But then we need to find an image for the second copy, and so we freely amalgamate another copy of \mathbf{A} and so on. Similarly, we need to have a pre-image of the original copy of \mathbf{A} et cetera. In the end, we get a doubly infinite path of copies of \mathbf{A} amalgamated over the respective copies of \mathbf{D} and \mathbf{R} (see Figure 1.1 (a) where \mathbf{A} is the path on three vertices and \mathbf{D} is an edge). This would be an EPPA-witness if it was not for the fact that it is infinite. In this case, one can clearly “wrap the path around” to a finite cycle to obtain an EPPA-witness (see Figure 1.1 (b)).




Figure 1.1: Construction of an EPPA-witness for one partial automorphism

The situation, however, becomes much more complicated when we consider at least two partial automorphism, as instead of a doubly-infinite path, we are gluing

the copies of \mathbf{A} along the free group on multiple generators (see Figure 1.2).




Figure 1.2: Trying to extend multiple partial automorphisms

Wrapping around the free group amalgam into a finite structure while preserving the extending automorphism is a much more complex problem. It is closely connected to the profinite topology on free groups, see Section 1.6.4. Let us just add here that with the correct tools, it is possible to make this approach work, see e.g. [HL00].

Now we present sketches of two combinatorial proofs of Hrushovski's theorem.

Proof of Theorem 1.6.3 (Herwig–Lascar [HL00]). Fix a finite graph $\mathbf{G} = (V, E)$. We will construct its EPPA-witness \mathbf{H} . Let Δ be the maximum degree of \mathbf{G} and put $X = \{(v, i) : v \in V, 0 \leq i < \Delta - \deg(v)\}$. Assuming that E and X are disjoint, put $Y = E \cup X$ and let $H = \binom{Y}{\Delta}$ be the set of all subsets of Y of size Δ . This will be the vertex set of \mathbf{H} . Given $A, B \in H$, let AB be an edge of \mathbf{H} if and only if $A \cap B \neq \emptyset$. Embed $\mathbf{G} \rightarrow \mathbf{H}$ sending $v \mapsto \{e \in E : v \in e\} \cup \{(v, i) : 0 \leq i < \Delta - \deg(v)\}$. This will be the copy of \mathbf{G} in \mathbf{H} whose partial automorphisms we will extend. It is easy to check that we have indeed constructed an embedding.

A partial automorphism φ of \mathbf{G} induces a partial permutation of E . It is easy to see that one can extend this into a permutation of Y such that this permutation induces a permutation of H which respects φ . (We need to make sure that vertices from the domain of φ are sent to their prescribed images. This prescribes some conditions on how we can extend the partial permutation of E to a permutation of Y but these conditions can always be satisfied). Finally, from the definition of edges of \mathbf{H} it follows that we have constructed an automorphism of \mathbf{H} . \square

Note that the number of vertices of \mathbf{H} is $(\Delta|V|)^\Delta$, so in particular, for graphs with bounded degree we have polynomial-size EPPA-witnesses. The study of *EPPA numbers* (cf. Ramsey numbers) of graphs is in its infancy, see some discussion in Chapter 2.

The following construction appears in full detail in Section 4.3.

Proof of Theorem 1.6.3 (Hubička–Konečný–Nešetřil [HKN22]). Let n be a natural number. Define a graph \mathbf{H}_n with the vertex set H_n being

$$H_n = \{(i, f) : 0 \leq i < n, f: \{0, \dots, n-1\} \rightarrow \{0, 1\}\}$$

and (i, f) forming an edge with (j, g) if and only if $f(j) \neq g(i)$. Note that an arbitrary permutation of $\{0, \dots, n-1\}$ induces an automorphism of \mathbf{H}_n by permuting the first coordinates and the function values accordingly. Note also that, given $0 \leq a < b < n$, the function $H_n \rightarrow H_n$ sending $(a, f) \mapsto (a, f')$ where $f'(b) = 1 - f(b)$ and $f'(x) = f(x)$ for $x \neq b$, $(b, f) \mapsto (b, f')$ where $f'(a) = 1 - f(a)$ and $f'(x) = f(x)$ for $x \neq a$, and $(i, f) \mapsto (i, f)$ if $i \notin \{a, b\}$, is an automorphism of \mathbf{H}_n . We call such an automorphism a *flip of a and b* .

Let \mathbf{G} be a graph on n vertices and assume without loss of generality that the vertex set of \mathbf{G} is $\{0, \dots, n-1\}$. Observe that the map $i \mapsto (i, f)$, where $f \in H_n$ is chosen such that $f(j) = 1$ if and only if $j < i$ and ij is an edge of \mathbf{G} , is an embedding $\mathbf{G} \rightarrow \mathbf{H}_n$.

Let φ be a partial automorphism of \mathbf{G} . First, extend it to a permutation of $\{0, \dots, n-1\}$ arbitrarily. This permutation induces an automorphism of \mathbf{H}_n , but this automorphism need not yet extend φ (acting on the copy of \mathbf{G} in \mathbf{H}_n). It can be shown that one can fix this by applying several flips. (In showing this one uses the fact that φ is an automorphism: If both u and v are in its domain then they both agree whether u and v should be flipped or not.) \square

In the Herwig–Lascar construction, the EPPA-witnesses only depended on the number of vertices and the maximum degree. The EPPA-witnesses in this construction only depend on the number of vertices and have always $n2^n$ vertices. It has been shown by Hrushovski [Hru92] that any EPPA-witness for the half-graph (also called ladder graph) with both parts of size n needs to have at least 2^n vertices. Hence the situation is, on the surface, very similar to Ramsey numbers: We have an exponential upper bound, an exponential lower bound, the bases of the exponents do not match and there are polynomial constructions for some special classes of graphs, see some discussion in Chapter 2. Note also that both constructions generalise naturally to higher arities and to general relational structures with no constraints (see the respective papers for the generalisations).

Look at the graphs \mathbf{H}_n from the second proof of Hrushovski’s theorem. They admit a natural perfect matching, namely between vertices (i, f) and $(i, 1 - f)$. If, instead of a graph, one defines a metric spaces with distances $\{1, 2, 3\}$ on \mathbf{H}_n by putting graph-adjacent vertices into distance 1, pairs (i, f) and $(i, 1 - f)$ into distance 3 and all the remaining pairs into distance 2, one actually gets an antipodal metric space of diameter 3 (see Section 1.4). In fact, this is how a minor variant of the second construction was first discovered – as an EPPA-witness for antipodal metric spaces of diameter 3 [EHKN20]. It is inspired by the valuation function constructions of Evans, Hubička and Nešetřil [EHN21] and Hubička, Nešetřil and the author [HKN19a] which are in turn inspired by the multiple-valued-logic EPPA construction of Hodkinson and Otto [HO03].

1.6.2 General constructions

Once there are some concrete examples of classes with some property, one should ask themselves two related questions: what do the proofs have in common and what are the underlying conditions making the property true, with the hope of unraveling some more general theory. In 2000, Herwig and Lascar [HL00] gave a very strong general condition for a class to have EPPA (in practice it boils down to being relational, having JEP and the amalgamation property and being described

by a finite family of forbidden homomorphisms). Hodkinson and Otto [HO03] later provided a construction of irreducible structure faithful EPPA-witnesses which implies EPPA for relational free amalgamation classes (see [Sin17]). This was generalized by Evans, Hubička and Nešetřil to free amalgamation classes where all functions are unary [EHN21]. All these results have been subsumed by the paper [HKN22] presented in Chapter 4 which proves the following theorem (in fact, it proves a significantly stronger result, the statement below is optimised for needing as little extra definitions as possible).

Theorem 1.6.4 (Hubička–Konečný–Nešetřil [HKN22]). *Let L be a language consisting of relations and unary functions, let \mathbf{A} be a finite irreducible L -structure, let \mathbf{B}_0 be a finite EPPA-witness for \mathbf{A} and let $n \geq 1$ be an integer. There is a finite L -structure \mathbf{B} satisfying the following.*

1. \mathbf{B} is an EPPA-witness for \mathbf{A} .
2. There is a homomorphism-embedding $\mathbf{B} \rightarrow \mathbf{B}_0$.
3. For every substructure \mathbf{C} of \mathbf{B} on at most n vertices there is a tree amalgamation \mathbf{D} of copies of \mathbf{A} and a homomorphism-embedding $f: \mathbf{C} \rightarrow \mathbf{D}$.
4. Every irreducible substructure of \mathbf{B} embeds into \mathbf{A} .

Here, a *tree amalgamation of copies of \mathbf{A}* is any structure which can be created by a series of free amalgamations of copies of \mathbf{A} over its substructures.

Intuitively, this theorem says the following: Start with a finite L -structure \mathbf{A} and its EPPA-witness \mathbf{B}_0 (obtained, for example, using [EHN21], or the generalised version of the valuation function construction of EPPA for graphs which we have seen several paragraphs above) and pick some parameter n . Then one can use \mathbf{B}_0 as a template for constructing a larger EPPA-witness \mathbf{B} which has a *projection* to \mathbf{B}_0 such that locally it looks tree-like. (Remember that one approach for constructing EPPA-witnesses is to glue the copies along a free group and then factorize, see also Section 1.6.4. In a way, the resulting structures have exactly the property that locally they look tree-like.) Note that in the theorem we only promise a homomorphism-embedding to a tree amalgamation and the cycle of length 4 has a homomorphism to the edge which surely is a tree amalgamation of edges, the existence of a homomorphism-embedding essentially means that the local structure looks like a tree amalgamation after we are allowed to blow up vertices. In the very limited understanding of the author, Otto’s result [Ott20] is able to get a more refined local structure in some cases.

For the full strength of Theorem 1.6.4 see Theorem 4.1.2. In particular, it works in a more general category of structures where the images of functions are not necessarily singletons and the language is equipped by a permutation group. This generalises Herwig’s *permorphisms* [Her95]. Theorem 4.1.1 then provides the basic EPPA-witnesses without any conditions in this category. Note that equipping the language by a permutation group is a very strong generalisation as it allows, for example, to prove EPPA for certain classes with functions of higher arity whose ranges are in different sorts than their domains, see e.g. Theorems 4.12.4 and 4.12.7.

The original formulation of the Herwig–Lascar theorem is different from Theorem 1.6.4, it is however its simple consequence (firstly, the Herwig–Lascar theorem

only considers relational structures, secondly, it talks about forbidden homomorphisms while the formulation below forbids homomorphism-embeddings):

Theorem 1.6.5 (Herwig–Lascar [HL00]). *Let L be a language consisting of relations and unary functions. Let \mathcal{F} be a finite family of finite L -structures and let $\mathbf{A} \in \text{Forb}_{\text{he}}(\mathcal{F})$ be a finite L -structure. If there exists a (not necessarily finite) structure $\mathbf{M} \in \text{Forb}_{\text{he}}(\mathcal{F})$ containing \mathbf{A} as a substructure such that each partial automorphism of \mathbf{A} extends to an automorphism of \mathbf{M} , then there exists a finite structure $\mathbf{B} \in \text{Forb}_{\text{he}}(\mathcal{F})$ which is an EPPA-witness for \mathbf{A} .*

Again, the true formulation in Chapter 4 is significantly stronger, see Theorem 4.1.5.

There is one more consequence of Theorem 1.6.4 which mimics the main result of [HN19] for Ramsey classes (see Section 1.7). For its true formulation, see Theorem 4.1.6:

Theorem 1.6.6. *Let L be a language consisting of relations and unary functions and let \mathcal{E} be a class of finite L -structures which has EPPA. Let \mathcal{K} be a hereditary locally finite automorphism-preserving subclass of \mathcal{E} with the strong amalgamation property which consists of irreducible structures. Then \mathcal{K} has EPPA.*

The following subsection defines what a locally finite automorphism-preserving subclass is:

Completions

As we alluded to in Section 1.1, there are general theorems which produce complicated structures from simple ones by means of complicated *multiamalgamations*, an example of this being Theorem 1.6.4. However, one usually wants to apply such a theorem to prove EPPA (or some other property) for a whole class. In the case of Theorem 1.6.4, this involves the following steps:

1. Choose a class \mathcal{K} for which we want to prove EPPA.
2. Find a super-class \mathcal{E} of \mathcal{K} which has EPPA and whose EPPA-witnesses will serve as the base structures \mathbf{B}_0 for an eventual application of Theorem 1.6.4.
3. Given $\mathbf{A} \in \mathcal{K}$ and its EPPA-witness $\mathbf{B}_0 \in \mathcal{E}$, prove that there exists a finite n such that if an EPPA-witness \mathbf{B} for \mathbf{A} has a homomorphism-embedding to \mathbf{B}_0 and every finite substructure of \mathbf{B} on n vertices has a homomorphism-embedding into a tree amalgam of copies of \mathbf{A} then one can use it to produce an EPPA-witness for \mathbf{A} in \mathcal{K} .

Obviously, the third step looks complicated. In practice, one usually wants to create the EPPA-witness in \mathcal{K} from \mathbf{B} just by “filling-in the holes”: For example, let \mathcal{K} be the class of all finite metric spaces with distances from $\{1, \dots, n\}$ (this example is also used in Section 4.12.2). Then \mathcal{E} can be the class of all finite graphs with edges labelled by $\{1, \dots, n\}$. This in particular means that the structures \mathbf{B}_0 and \mathbf{B} are finite graphs with edges labelled by $\{1, \dots, n\}$. In order to turn \mathbf{B} into a metric space, we need to fill in the missing distances, and we need to preserve all automorphisms of \mathbf{B} while doing it (more precisely, we only need to

preserve enough automorphisms to be able to extend all partial automorphisms of \mathbf{A} , but the theorem does not give us any control over what these are).

Recall the *shortest path completion* which was mentioned already several times in this paper: Given a graph \mathbf{B} with edges labelled by $\{1, \dots, n\}$ (and $d(u, v)$ being the label of the edge u, v if it exists, zero if $u = v$ and undefined if $u \neq v$), we can define a metric d' by putting

$$d'(u, v) = \min(n, \min_{\mathbf{P} \text{ a path } u \rightarrow v \text{ in } \mathbf{B}} \|\mathbf{P}\|),$$

where $\|\mathbf{P}\|$ is the sum of the labels of edges of \mathbf{P} . It is easy to see that d' is a metric on B^2 and that $d \subseteq d'$ if and only if \mathbf{B} contains no *non-metric cycle*, that is, a cycle with one label larger than the sum of the others. Moreover, if $d \subseteq d'$ then $\text{Aut}(\mathbf{B}) = \text{Aut}((B, d'))$. Finally, note that the largest non-metric cycle has n vertices. It is easy to see that no tree amalgam of copies of \mathbf{A} contains a non-metric cycle and that the existence of a non-metric cycle is preserved by homomorphism-embeddings. This means that we can use n as the parameter for Theorem 1.6.4 and use the shortest path completion of \mathbf{B} to obtain the desired EPPA-witness in \mathcal{K} .

One can try to abstract this argument which leads to the notion of *completions* and *locally finite subclasses* which were introduced by Hubička and Nešetřil in [HN19] (we again remark that in Chapter 4 these definitions are given for languages equipped with a permutation group):

Definition 1.6.1. Let \mathbf{C} be a structure. An irreducible structure \mathbf{C}' is a *completion* of \mathbf{C} if there is a homomorphism-embedding $\mathbf{C} \rightarrow \mathbf{C}'$. It is a *strong completion* if the homomorphism-embedding is injective. A completion is *automorphism-preserving* if it is strong and for every $\alpha \in \text{Aut}(\mathbf{C})$ there is $\alpha' \in \text{Aut}(\mathbf{C}')$ such that $\alpha \subseteq \alpha'$ and moreover the map $\alpha \mapsto \alpha'$ is a group homomorphism $\text{Aut}(\mathbf{C}) \rightarrow \text{Aut}(\mathbf{C}')$.

Note that that completion is a strengthening of amalgamation: Let \mathcal{K} be a class of irreducible structures. The amalgamation property for \mathcal{K} can be equivalently formulated as follows: For $\mathbf{A}, \mathbf{B}_1, \mathbf{B}_2 \in \mathcal{K}$ embeddings $\alpha_1: \mathbf{A} \rightarrow \mathbf{B}_1$ and $\alpha_2: \mathbf{A} \rightarrow \mathbf{B}_2$, there is $\mathbf{C} \in \mathcal{K}$ which is a completion of the free amalgamation of \mathbf{B}_1 and \mathbf{B}_2 over \mathbf{A} with respect to α_1 and α_2 (which itself need not be in \mathcal{K}). In the same way, strong completion strengthens strong amalgamation and automorphism-preserving completion strengthens the so-called *amalgamation property with automorphisms* (see Section 1.6.6). See also Question 2.0.12 for a related question.

Definition 1.6.2. Let L be a language, let \mathcal{E} be a class of finite L -structures and let \mathcal{K} be a subclass of \mathcal{E} consisting of irreducible structures. We say that \mathcal{K} is a *locally finite subclass* of \mathcal{E} if for every $\mathbf{A} \in \mathcal{K}$ and every $\mathbf{B}_0 \in \mathcal{E}$ there is a finite integer $n = n(\mathbf{A}, \mathbf{B}_0)$ such that every L -structure \mathbf{B} has a strong completion $\mathbf{B}' \in \mathcal{K}$, provided that it satisfies the following:

1. Every irreducible substructure of \mathbf{B} has an embedding to \mathbf{A} ,
2. there is a homomorphism-embedding from \mathbf{B} to \mathbf{B}_0 , and
3. every substructure of \mathbf{B} on at most n vertices has a completion in \mathcal{K} .

We say that \mathcal{K} is a *locally finite automorphism-preserving subclass* of \mathcal{E} if \mathbf{B}' can always be chosen to be automorphism-preserving.

Note that if \mathcal{K} is hereditary, point 1 implies that every irreducible substructure of \mathbf{B} is in \mathcal{K} . Note also that we are only promised that every substructure on at most n vertices has a completion in \mathcal{K} , even though we are asking for a strong one.

Luckily, for languages where all functions are unary, one can prove that if a structure has a completion in a strong amalgamation class then it has in fact a strong completion, which makes verifying local finiteness much easier. This was first proved in [HN19] as Proposition 2.6. With all these definitions together, it is easy to see that Theorem 1.6.4 implies Theorem 1.6.6.

1.6.3 Examples

By now there are many amalgamation classes for which EPPA is known. The aim of the list below is not to give an exhaustive list nor to faithfully attribute all partial results, instead the aim is to point out the interesting cases which one should keep in mind when thinking about EPPA.

- EPPA for sets with no structure is immediate as each such set is in fact homogeneous (one can extend a partial permutation arbitrarily).
- EPPA for free amalgamation classes where all functions are unary has been proved by Evans, Hubička and Nešetřil [EHN21]. Above we have seen simple proofs for graphs which easily generalise to hypergraphs.
- EPPA for finite metric spaces has been proved by Solecki [Sol05] using the Herwig–Lascar theorem and independently by Vershik [Ver08]. By now there are several other proofs of this result, let us mention a self-contained combinatorial proof by Hubička, Nešetřil, and the author [HKN19a] and Sabok’s proof using profinite topology on free groups [Sab17] (see also Section 1.6.4).
- The most general result on EPPA for generalised metric spaces is in the Master thesis of the author [Kon19] (see also [HKN18]) which in particular implies EPPA for S -metric spaces as well as Λ -ultrametric spaces which correspond to families of cross-cutting equivalence relations and have been introduced by Braunfeld [Bra16].
- EPPA holds for all metrically homogeneous graphs which are not tree-like [ABWH⁺17b, EHKN20, Kon20].
- EPPA also holds for the class of all finite two-graphs which was proved by Evans, Hubička, Nešetřil, and the author [EHKN20]. This is an important example because it exhibits unusual behavior: It does have EPPA but no APA (see Section 1.6.6), ample generics (see Section 1.6.6), or SIR. Coherent EPPA (see Section 1.6.5) is open for this class but likely does not follow from the (only) existing proof. It is moreover one of the few known classes where their EPPA and Ramsey expansions do not differ just by linear orders – the Ramsey expansion of two-graphs is the class of all finite ordered graphs (see Section 1.7).

- The class of all finite groups has EPPA. This was proved by Siniora [Sin17] and the proof of this is very different from the other proofs of EPPA, the EPPA-witness for a finite group G is simply $\text{Sym}(G)$ where we understand G as a subgroup of $\text{Sym}(G)$ using the standard representation by left multiplication.
- The class of all finite semigeneric tournaments and classes of all finite n -partite tournaments for all $n \geq 2$ have EPPA by a result of Hubička, Jahel, Sabok, and the author [HJKS19, HJKS23]. These are interesting because, similarly to two-graphs, ad-hoc valuation constructions were necessary as opposed to a straightforward application of Theorem 1.6.6.
- In Chapter 4 we show that certain classes with constants and certain classes with two unary relations and certain non-unary functions whose domains have one unary relations while ranges have the other have EPPA. On the other hand, EPPA for the class of all partial binary functions remains open (see Question 4.13.2).
- In the same chapter we also prove EPPA for the class of all finite k -orientations with d -closures. These are important structures arising from the Hrushovski constructions which were used by Evans, Hubička, and Nešetřil to provide a counterexample for an important question in the structural Ramsey theory (see Section 1.7), and they also contain non-unary functions.
- There are various finite homogeneous structures which, naturally, give rise to various classes with EPPA, for example finite vector spaces, or the class of all finite skew-symmetric bilinear forms. These have a very different flavour from the rest of this list and, from the point of view of studying general constructions and sufficient conditions, are not overly interesting. On the other hand, they are very important examples to keep in mind when posing general questions about EPPA (e.g. Conjecture 4.1.7).

There is one very natural class missing in this list: The class of all finite tournaments. The reason for this is that EPPA is open for this class (see Question 2.0.3 for more discussion). This problem has been posed already in 2000 by Herwig and Lascar [HL00] who proved that it is equivalent to a problem from the area of profinite group theory.

1.6.4 Profinite topology

Let G be a group. The *profinite topology* on G is given by the following basis of open sets:

$$\{gH : g \in G, H \text{ is a subgroup of } G \text{ of finite index}\}.$$

There are two relevant important theorems in this area:

Theorem 1.6.7 (M. Hall [Hal50]). *Let n be finite, let F_n be the free group on n generators and let $H \leq F_n$ be finitely generated. Then H is closed in the profinite topology on F_n .*

An equivalent formulation is the following: Given $H \leq F_n$ finitely generated, we have that

$$H = \bigcap \{K : H \leq K \leq F_n \text{ and } K \text{ has finite index}\}.$$

Theorem 1.6.8 (Ribes–Zalesskii [RZ93]). *Let H_1, \dots, H_m be finitely generated subgroups of F_n . Then $H_1 \cdots H_m = \{h_1 \cdots h_m : (\forall i) h_i \in H_i\}$ is closed in the profinite topology on F_n .*

These results allow us to make precise the intuition about factoring the free group in order to get a finite EPPA-witness from Section 1.6.1 and sketch the Mackey(-like) construction [Mac66]: Given a graph \mathbf{A} , let P be the set of all partial automorphisms of \mathbf{A} and consider the free group $F(P)$. Eventually, we want to find a finite index subgroup of $F(P)$ and define a graph structure on the cosets such that the left multiplication by any member of $F(P)$ induces an automorphism of the coset-graph. Moreover, we want this graph to contain a copy of \mathbf{A} (the natural candidates for its vertices are the identities on single vertices) and extend its partial automorphisms (the natural candidate for the extension of $p \in P$ is the automorphism induced by the left multiplication by p). In fact, the graph structure on the cosets will then simply be chosen to be the minimal graph satisfying these conditions. Finally, these conditions can be written as identities on $F(P)$ and the aforementioned theorems then ensure that these identities can indeed be enveloped by a finite index subgroup. See Section 2 of [HL00] for a proper proof of EPPA for graphs using this technique, see also Section 8 of [Sab17] for a proof of EPPA for metric spaces by profinite-topological methods.

Herwig and Lascar proved that EPPA for tournaments (see Question 2.0.3) is equivalent to a question in *pro-odd* (or *odd-adic*) topology: Here, the pro-odd topology on a free group F_n is given by a basis of open sets

$$\{gH : g \in F_n, H \text{ is a normal subgroup of } F_n \text{ of odd index}\}.$$

Theorem 1.6.9 (Herwig–Lascar [HL00]). *The following statements are equivalent:*

1. *The class of all finite tournaments has EPPA.*
2. *Let H be a finitely generated subgroup of F_n . Then H is closed in the pro-odd topology if and only if, for every $a \in F_n$, if $a^2 \in H$ then $a \in H$.*

Let us remark that Huang, Pawliuk, Sabok, and Wise [HPSW19] disproved EPPA for a certain version of hypertournaments (different from the one defined in Section 1.4 and used in Chapter 7) defined specifically to obtain a variant of the equivalence above for which the topological statement is false.

Remark 1.6.1. Rosendal [Ros11b] proved that the Ribes–Zalesskii theorem [RZ93] is equivalent to EPPA for the class of all finite metric spaces. Hubička, Nešetřil, and the author [HKN19a] gave a short combinatorial proof of EPPA for this class (using a variant of the valuation function construction), thereby obtaining a simple proof of the Ribes–Zalesskii theorem.

1.6.5 Coherent EPPA

Given a finite structure \mathbf{A} , the set of its partial automorphism is naturally equipped with a partial composition operator such that $p_2 \circ p_1$ is defined if the range of p_1 is equal to the domain of p_2 . Bhattacharjee and Macpherson [BM05] proved that graphs have a “composition-respecting” EPPA and that this implies that the automorphism group of the random graph contains a dense locally finite subgroup. (Here, a group is locally finite if finite sets generate finite subgroups.)

Solecki and Sinióra [Sol09, SS19] generalised their methods: Let \mathbf{B} be an EPPA-witness for \mathbf{A} and let P be the set of all partial automorphisms of \mathbf{A} . We say that \mathbf{B} is a *coherent* EPPA-witness for \mathbf{A} if there is a map $\phi: P \rightarrow \text{Aut}(\mathbf{B})$ such that $p \subseteq \phi(p)$ for every p and, moreover, if $p_2 \circ p_1$ is defined then $\phi(p_2 \circ p_1) = \phi(p_2) \circ \phi(p_1)$. A class has *coherent EPPA* if it has EPPA and all witnesses can be chosen to be coherent. Solecki and Sinióra strengthened the Herwig–Lascar theorem to prove coherent EPPA [SS19] and also proved that if a Fraïssé class \mathcal{C} has EPPA then the automorphism group of its Fraïssé limit contains a dense locally finite subgroup.

Let us remark that all results from Chapter 4 (presented in Section 1.6.2) actually prove coherent EPPA and that essentially all classes for which EPPA is known have even coherent EPPA. The odd one out is the class of all finite two-graphs for which EPPA is proved by proving EPPA for the class of all antipodal metric spaces of diameter 3 and then considering the actions on antipodal pairs. It turns out that this step does not preserve coherence (or at least we have been unable to make it preserve coherence), see Remark 5.4.1. On the other hand, the existence of a dense locally finite subgroup transfers to the Fraïssé limit of the class of all finite two-graphs.

1.6.6 Ample generics

In mathematics, one always wants to understand the “typical” objects. For example, the “typical” finite graph has diameter 2, the “typical” countable graph is isomorphic to the random graph, the “typical” polynomial does not have a root at π . In topological spaces, a “typical” (usually called *generic*) property is one which holds on a comeagre set (a set is *comeagre* if it is the intersection of countably many sets such that the interior of each of them is dense). In this section we will be concerned with generic automorphisms.

Definition 1.6.3 ([Tru92, HHLS93, KR07]). Let L be a language, let \mathbf{M} be a countable locally finite homogeneous L -structure and let $n \geq 1$ be an integer. We say that \mathbf{M} has *n -generic automorphisms* (or simply *n -generics*) if G has a comeagre orbit on G^n in its action by diagonal conjugation. We say that \mathbf{M} has *ample generics* if it has n -generic automorphisms for every $n \geq 1$.

Here, the action by diagonal conjugation is defined by

$$g \cdot (h_1, \dots, h_n) = (gh_1g^{-1}, \dots, gh_ng^{-1}).$$

For $n = 1$ we just say that \mathbf{M} has generic automorphisms. This was first studied by Truss [Tru92] who for example proved that the generic member of $\text{Sym}(\mathbb{N})$ has no infinite cycles and infinitely many k -cycles for every $k \in \mathbb{N}$.

He also proved the existence of generic automorphisms of (\mathbb{Q}, \leq) or the random graph. Hodges, Hodkinson, Lascar, and Shelah [HHLS93] proved that the random graph has ample generics (and used Hrushovski's theorem about EPPA for graphs in their proof). Their methods were abstracted by Kechris and Rosendal [KR07]:

Definition 1.6.4. Let L be a language, let \mathcal{C} be a class of finite L -structures and let $n \geq 1$ be an integer. An n -system over \mathcal{C} is a tuple $(\mathbf{A}, p_1, \dots, p_n)$, where $\mathbf{A} \in \mathcal{C}$ and p_1, \dots, p_n are partial automorphisms of \mathbf{A} . We denote by \mathcal{C}^n the class of all n -systems over \mathcal{C} .

If $P = (\mathbf{A}, p_1, \dots, p_n)$ and $Q = (\mathbf{B}, q_1, \dots, q_n)$ are both n -systems over \mathcal{C} and $f: \mathbf{A} \rightarrow \mathbf{B}$ is an embedding of L -structures, we say that f is an *embedding of n -systems* $P \rightarrow Q$ if for every $1 \leq i \leq n$ it holds that $f \circ p_i \subseteq q_i \circ f$ (in particular, $f(\text{Dom}(p_i)) \subseteq \text{Dom}(q_i)$ and $f(\text{Range}(p_i)) \subseteq \text{Range}(q_i)$).

Definition 1.6.5. Let L be a language, let \mathcal{C} be a class of finite L -structures and let $n \geq 1$ be an integer. We say that \mathcal{C}^n has the *joint embedding property* if for every $P, Q \in \mathcal{C}^n$ there exists $S \in \mathcal{C}^n$ with embeddings of n -systems $f: P \rightarrow S$ and $g: Q \rightarrow S$. We say that \mathcal{C}^n has the *weak amalgamation property* if for every $T \in \mathcal{C}^n$ there exists $\hat{T} \in \mathcal{C}^n$ and an embedding of n -systems $\iota: T \rightarrow \hat{T}$ such that for every pair of n -systems $P, Q \in \mathcal{C}^n$ and embeddings of n -systems $\alpha_1: \hat{T} \rightarrow P$ and $\alpha_2: \hat{T} \rightarrow Q$ there exists $S \in \mathcal{C}^n$ with embeddings on n -systems $\beta_1: P \rightarrow S$ and $\beta_2: Q \rightarrow S$ such that $\beta_1 \alpha_1 \iota = \beta_2 \alpha_2 \iota$.

Theorem 1.6.10 (Kechris–Rosendal [KR07]). *Let L be a language, let \mathbf{M} be a countable locally finite homogeneous L -structure, put $\mathcal{C} = \text{Age}(\mathbf{M})$ and fix $n \geq 1$. Then \mathbf{M} has n -generic automorphisms if and only if \mathcal{C}^n has the joint embedding property and the weak amalgamation property.*

In order to explain the connection between EPPA and ample generics, we need one more definition

Definition 1.6.6. Let L be a language and let \mathcal{C} be a class of finite L -structures. We say that \mathcal{C} has the *amalgamation property with automorphisms* (abbreviated as *APA*) if for every $\mathbf{A}, \mathbf{B}_1, \mathbf{B}_2 \in \mathcal{C}$ and embeddings $\alpha_1: \mathbf{A} \rightarrow \mathbf{B}_1$, $\alpha_2: \mathbf{A} \rightarrow \mathbf{B}_2$ there exists $\mathbf{C} \in \mathcal{C}$ with embeddings $\beta_1: \mathbf{B}_1 \rightarrow \mathbf{C}$ and $\beta_2: \mathbf{B}_2 \rightarrow \mathbf{C}$ such that $\beta_1 \circ \alpha_1 = \beta_2 \circ \alpha_2$ (i.e. \mathbf{C} is an amalgamation of \mathbf{B}_1 and \mathbf{B}_2 over \mathbf{A} with respect to α_1 and α_2) and moreover whenever we have $f \in \text{Aut}(\mathbf{B}_1)$ and $g \in \text{Aut}(\mathbf{B}_2)$ such that $f(\alpha_1(A)) = \alpha_1(A)$, $g(\alpha_2(A)) = \alpha_2(A)$ and for every $a \in A$ it holds that $\alpha_1^{-1}(f(\alpha_1(a))) = \alpha_2^{-1}(g(\alpha_2(a)))$ (that is, f and g agree on the copy of \mathbf{A} we are amalgamating over), then there is $h \in \text{Aut}(\mathbf{C})$ which extends $\beta_1 f \beta_1^{-1} \cup \beta_2 g \beta_2^{-1}$. We call such \mathbf{C} with embeddings β_1 and β_2 an *APA-witness* for \mathbf{B}_1 and \mathbf{B}_2 over \mathbf{A} with respect to α_1 and α_2 .

Proposition 1.6.11 (Kechris–Rosendal [KR07]). *Let L be a language and let \mathcal{C} be a class of finite L -structures. If \mathcal{C} has EPPA and APA then \mathcal{C}^n has the weak amalgamation property for every $n \geq 1$.*

Proof. Fix $n \geq 1$. If $S = (\mathbf{S}, s_1, \dots, s_n) \in \mathcal{C}^n$ is an n -system, we denote by $\hat{S} = (\hat{\mathbf{S}}, \hat{s}_1, \dots, \hat{s}_n) \in \mathcal{C}^n$ the n -system where $\hat{\mathbf{S}}$ is an EPPA-witness for \mathbf{S} (with respect to the inclusion embedding) and for every $1 \leq i \leq n$ it holds that \hat{s}_i is an automorphism of $\hat{\mathbf{S}}$ extending s_i .

We now prove that \mathcal{C}^n has the weak amalgamation property. Towards that, fix some $T = (\mathbf{T}, t_1, \dots, t_n) \in \mathcal{C}^n$. Let $P = (\mathbf{P}, p_1, \dots, p_n), Q = (\mathbf{Q}, q_1, \dots, q_n) \in \mathcal{C}^n$ be arbitrary n -systems with embeddings $\alpha_1: \hat{T} \rightarrow P$ and $\alpha_2: \hat{T} \rightarrow Q$.

Use APA for \mathcal{C} to get $\mathbf{S} \in \mathcal{C}$ and embeddings $\beta_1: \hat{\mathbf{P}} \rightarrow \mathbf{S}$ and $\beta_2: \hat{\mathbf{Q}} \rightarrow \mathbf{S}$ such that \mathbf{S} with β_1 and β_2 form an APA-witness for $\hat{\mathbf{P}}$ and $\hat{\mathbf{Q}}$ over $\hat{\mathbf{T}}$ with respect to α_1 and α_2 . Let $S = (\mathbf{S}, s_1, \dots, s_n) \in \mathcal{C}^n$ be some n -system such that for every $1 \leq i \leq n$ we have that s_i extends $\beta_1 \hat{p}_i \beta_1^{-1} \cup \beta_2 \hat{q}_i \beta_2^{-1}$. It is straightforward to verify that S is the desired n -system witnessing the weak amalgamation property for P, Q and T . \square

Example 1.6.1. Consider the class \mathcal{C} of all finite graphs. By a theorem of Hrushovski [Hru92] (or by Section 4.3) we know that \mathcal{C} has EPPA. APA for \mathcal{C} is an easy exercise (in general, APA for free amalgamation classes is always true). Hence, by the above proposition, \mathcal{C}^n has the weak amalgamation property for every $n \geq 1$. To prove ample generics for the countable random graph it thus remains to prove the joint embedding property for \mathcal{C}^n . However, it is again an easy exercise (simply take the disjoint union of the graphs and the partial automorphisms).

Consequences of ample generics

Ample generics are an important property because their existence has many consequences. Here we mention a few, see also [Sin17].

Theorem 1.6.12 (Kechris–Rosendal [KR07]). *If G is a Polish group with ample generics then it has the small index property.*

Here, a Polish group has ample generics if it has a comeagre orbit on G^n in its action by diagonal conjugation for every $n \geq 1$. A group has the *small index property* if every subgroup of index $< 2^{\aleph_0}$ is open. The small index property for the automorphism group of the random graph was proved by Hodges, Hodkinson, Lascar, and Shelah in 1992 [HHLS93] and their paper can be seen as the beginning of this whole theory.

Note that the small index property is not equivalent to the existence of ample generics. For example, (\mathcal{Q}, \leq) does not have even 2-generics but it has the small index property [Tru89].

Theorem 1.6.13 (Kechris–Rosendal [KR07]). *If \mathbf{M} is a countable ω -categorical structure with ample generics then $\text{Aut}(\mathbf{M})$ has uncountable cofinality.*

Here, the *cofinality* of a group G is the smallest cardinal κ such that G can be written as the union of a chain of proper subgroups of size κ .

Ample generics also imply having the 21-Bergman property or the property (FA), see [Sin17].

Examples

- If a class has the strong amalgamation property and admits automorphism-preserving completions then it has APA. This means that whenever EPPA is proved using theorem 1.6.6, ample generics are an immediate corollary.

- The class of all finite tournaments has APA (just orient all missing edges in one direction) and EPPA is an open problem (Question 2.0.3). Siniara proved [Sin17] that EPPA for tournaments is in fact equivalent to ample generics.
- The class of all finite groups has both EPPA and APA (see [Sin17]) and, consequently, Hall's universal locally finite group has ample generics.
- The class of all finite two-graphs has EPPA but it does not have APA. In the published version of the paper from Chapter 5 we ask whether the generic two-graph has ample generics. Since then this question has been answered negatively by Evans, Hubička, Nešetřil, Simon, and the author in an unpublished note: Two-graphs do have 1-generics (one can simply pick a graph from the represented switching classes according to the partial automorphisms and pretend we are working with graphs), but they do not have 2-generics, in fact, for \mathbb{T} being the Fraïssé limit of finite two-graphs and $G = \text{Aut}(\mathbb{T})$, there is no dense G -orbit on G^2 . If there was one, then in particular any two partial automorphisms g_1, g_2 can be extended to automorphisms of \mathbb{T} with a common fixed point. However, there is a finite non-complete, non-empty two-graph \mathbf{A} with its automorphism group being 2-generated and acting 2-transitively on A . Assume $\mathbf{A} \subseteq \mathbb{T}$ and let g_1, g_2 be generators of $\text{Aut}(\mathbf{A})$. If there are automorphisms of \mathbb{T} extending g_1 and g_2 with a common fixed point b , it follows that the two-graph on $A \cup \{b\}$ has either all or no hyperedges of the form $\{b, a, a'\}$ for $a, a' \in A$, consequently, \mathbf{A} is empty or complete, a contradiction.
- Various order-like classes have 1-generics but not 2-generics. Since they do not have EPPA, they are not that relevant for this thesis.
- Using the yet-unpublished result of Hubička, Jahel, Sabok, and the author [HJKS19, HJKS23], the semigeneric tournament and all n -partite tournaments have ample generics (APA follows by an analogous argument as for tournaments, one can orient the added edges in a canonical way).

Note that the list above contains no example with n -generics but not $(n + 1)$ -generics for any $n \geq 2$, see also Question 2.0.10.

1.7 Structural Ramsey theory

For more background than given in this section, see the introductions to Hubička's habilitation thesis [Hub20b] and the author's Master thesis [Kon19], and the surveys by Bodirsky [Bod15] and Nguyen Van Thé [NVT15].

In 1930, Ramsey published a paper where he proves the following theorem (here, $\binom{A}{p}$ denotes the set of all p -element subsets of A , recall also that we identify $N = \{0, \dots, N - 1\}$):

Theorem 1.7.1 (Ramsey's theorem [Ram30]). *For every triple of natural numbers n, p, k with $k \geq 1$ there is N such that the following holds: For every colouring $c: \binom{N}{p} \rightarrow k$ there is an n -element subset $H \in \binom{N}{n}$ such that c restricted to $\binom{H}{p}$ is constant.*

Let us start with an easy observation.

Observation 1.7.2. *Ramsey's theorem for $k = 2$ implies Ramsey's theorem with an arbitrary k .*

Proof. Start with an arbitrary k -colouring c and define a 2-colouring c' such that $c'(f) = 0$ if $c(f) = 0$ and $c'(f) = 1$ otherwise. Apply Ramsey's theorem for $k = 2$, thereby either finding a monochromatic set in colour 0, or reducing the number of colours to $k - 1$. Use induction. \square

This observation holds in general for these Ramsey-type statements which justifies that we will state some definitions and results for two colours only.

Note that Ramsey's theorem can be phrased in terms of finite linear orders and embeddings between them:

Theorem 1.7.3 (Ramsey's theorem for finite linear orders [Ram30]). *For every pair of finite linear orders \mathbf{A} and \mathbf{B} and an integer $k \geq 1$ there exists a finite linear order \mathbf{C} such that for every colouring $c: \binom{\mathbf{C}}{\mathbf{A}} \rightarrow k$ there exists an embedding $f: \mathbf{B} \rightarrow \mathbf{C}$ such that c is constant on $\binom{f(\mathbf{B})}{\mathbf{A}}$.*

This motivates the following notation: Given L -structures \mathbf{A} , \mathbf{B} and \mathbf{C} , we write

$$\mathbf{C} \longrightarrow (\mathbf{B})_{k,\ell}^{\mathbf{A}}$$

to denote the following statement:

For every colouring $c: \binom{\mathbf{C}}{\mathbf{A}} \rightarrow k$ there exists an embedding $f: \mathbf{B} \rightarrow \mathbf{C}$ such that c takes at most ℓ values on $\binom{f(\mathbf{B})}{\mathbf{A}}$.

If $\ell = 1$ we often omit it and write $\mathbf{C} \longrightarrow (\mathbf{B})_k^{\mathbf{A}}$ instead. We say that \mathbf{C} is a *Ramsey-witness for colouring copies of \mathbf{A} in \mathbf{B}* (or simply a *Ramsey-witness for \mathbf{A} and \mathbf{B}* because which one is being coloured is usually obvious from the context) if $\mathbf{C} \longrightarrow (\mathbf{B})_2^{\mathbf{A}}$. A class \mathcal{C} of finite L -structures is *Ramsey* (or *has the Ramsey property*) if for every $\mathbf{A}, \mathbf{B} \in \mathcal{C}$ there exists $\mathbf{C} \in \mathcal{C}$ which is a Ramsey witness for \mathbf{A} and \mathbf{B} .

In 1977, Nešetřil and Rödl, and independently in 1978, Abramson and Harrington proved the following:

Theorem 1.7.4 (Nešetřil–Rödl [NR77a, NR77b], Abramson–Harrington [AH78]). *The class of all finite linearly ordered graphs is Ramsey.*

The natural question after seeing this statement is “How about the class of all graphs without an order?” The answer is that the class of all graphs without an order is not Ramsey, and what follows are three explanations ordered increasingly by quality but also complexity.

1. Notice that, in the definition of a Ramsey-witness, we are colouring embeddings. A graph edge has two embeddings into itself – the identity and the reverse, and both these embeddings are always present, hence we can colour each of them by a different colour.
2. The explanation above can be summarised by saying that a Ramsey class needs to consist of rigid structures (i.e. no non-identity automorphisms), and adding a linear order is one way of ensuring rigidity. Another possible solution to this is to colour *copies* instead of embeddings (a copy is an equivalence class of embeddings which have the same image). However, it turns out that this solution does not work: Assume that \mathbf{A} is the path on three vertices a_1, a_2, a_3 such that a_1a_2 and a_2a_3 are the only edges of \mathbf{A} and \mathbf{B} is a five-cycle. Let \mathbf{C} be an arbitrary hypothetical Ramsey-witness for \mathbf{A} and \mathbf{B} , pick an arbitrary linear order on \mathbf{C} and colour copies of \mathbf{A} by colour 0 if and only if the middle vertex of the copy is also the middle vertex in the induced order, otherwise colour them by colour 1. It is easy to see that any linear ordering of \mathbf{B} will contain copies of \mathbf{A} of both colours thereby disproving the Ramsey property for copies. In fact, one can only obtain Ramsey-witnesses for copies when colouring complete graphs and independent sets, see [GGL95, Ch. 25, Sec. 5].
3. The previous explanation shows that the role of linear order is not only ensure rigidity. In fact, we will see in Proposition 1.7.13 that if a Fraïssé class has the Ramsey property then the automorphism group of its Fraïssé limit preserves a linear order on the Fraïssé limit. Morally, this says that Ramsey classes consist of linearly ordered structures.

In the third explanation we only talked about Fraïssé classes. Similarly to EPPA, if a Ramsey class has the joint embedding property, it has also the amalgamation property:

Observation 1.7.5 (Nešetřil [Neš89, Neš05]). *Let \mathcal{C} be a Ramsey class of finite structures with the joint embedding property. Then \mathcal{C} has the amalgamation property.*

Proof. We need to show that for every $\mathbf{A}, \mathbf{B}_1, \mathbf{B}_2 \in \mathcal{C}$ and embeddings $\alpha_1: \mathbf{A} \rightarrow \mathbf{B}_1$ and $\alpha_2: \mathbf{A} \rightarrow \mathbf{B}_2$ there is $\mathbf{C} \in \mathcal{C}$ and embeddings $\beta_1: \mathbf{B}_1 \rightarrow \mathbf{C}$ and $\beta_2: \mathbf{B}_2 \rightarrow \mathbf{C}$ such that $\beta_1 \circ \alpha_1 = \beta_2 \circ \alpha_2$.

Let \mathbf{B} be a joint embedding of \mathbf{B}_1 and \mathbf{B}_2 and take $\mathbf{C} \in \mathcal{C}$ such that $\mathbf{C} \rightarrow (\mathbf{B})_2^{\mathbf{A}}$. We will prove that \mathbf{C} is the amalgam we are looking for. Assume the contrary which means that there is no embedding $\alpha: \mathbf{A} \rightarrow \mathbf{C}$ with the property that there are embeddings $\beta_1: \mathbf{B}_1 \rightarrow \mathbf{C}$ and $\beta_2: \mathbf{B}_2 \rightarrow \mathbf{C}$ such that $\beta_i \circ \alpha_i = \alpha$

for $i \in \{1, 2\}$. Hence, for every $\alpha: \mathbf{A} \rightarrow \mathbf{C}$ there is at most one such embedding $\beta_i: \mathbf{B}_i \rightarrow \mathbf{C}$. Define the colouring $c: \binom{\mathbf{C}}{\mathbf{A}} \rightarrow \{0, 1\}$ by letting

$$c(\alpha) = \begin{cases} 0 & \text{if there is } \beta_1: \mathbf{B}_1 \rightarrow \mathbf{C} \text{ such that } \alpha = \beta_1 \circ \alpha_1 \\ 1 & \text{otherwise.} \end{cases}$$

For an illustration, see Figure 1.3.




Figure 1.3: Illustration of the proof of Observation 1.7.5. Copies of \mathbf{A} from \mathbf{B}_1 are coloured black, copies of \mathbf{A} from \mathbf{B}_2 are coloured white.

As $\mathbf{C} \rightarrow (\mathbf{B})_2^{\mathbf{A}}$, there is an embedding $\beta: \mathbf{B} \rightarrow \mathbf{C}$ such that c is constant on $\binom{\beta(\mathbf{B})}{\mathbf{A}}$. But there are at least two embeddings of \mathbf{A} into β — one is given by α_1 and the other is given by α_2 . And α_1 can be extended to an embedding of \mathbf{B}_1 , while α_2 can be extended to an embedding of \mathbf{B}_2 , hence they got different colours, which is a contradiction. \square

Observation 1.7.5 gives rise to the question whether every amalgamation class is a Ramsey class, to which the answer is negative: As we have seen, the class of all finite graphs is not Ramsey. A follow-up question might be whether every amalgamation class, when enriched by all possible linear orders (for each structure with n vertices in the original class there will be $n!$ structures in the new, ordered class), is Ramsey. And the answer is, again, negative:

Proposition 1.7.6. *Let \mathbf{M} be the disjoint union of two infinite cliques K_ω (it is clearly homogeneous). Let $\mathcal{C} = \text{Age}(\mathbf{M})$ be the class of all finite graphs such that they are either a clique or the disjoint union of two cliques. And let \mathcal{C}^\leq be the class of all possible linear orderings of members of \mathcal{C} . Then \mathcal{C}^\leq has the amalgamation property, but does not have the Ramsey property.*

Proof. The amalgamation property is easy: since the order and the graph structure are independent, we can use the amalgamation procedure for the order and the graph structure independently using the fact that both structures have the strong amalgamation property. Now let \mathbf{A} be a vertex and \mathbf{B} be a pair of vertices not connected by an edge. Any $\mathbf{C} \in \mathcal{C}^\leq$ which contains a copy of \mathbf{B} must consist of two cliques. Then one can colour the vertices of one of the cliques red and the vertices of the other clique blue and there will be no monochromatic non-edge in this colouring. \square

Such a situation happens often and the following section introduces a way to deal with it.

1.7.1 Expansions

Notice that our colouring was based on the fact that the edge relation in the structure \mathbf{M} is actually an equivalence relation with two equivalence classes. If one could, for example, expand the language by a unary relation and distinguish the equivalence classes by putting the unary relation on all vertices from one of them, then such a class equipped with an order would be Ramsey. In order to formalize this observation and state another variant of the question whether all classes are Ramsey, we need to give a model-theoretic definition.

Definition 1.7.1 (Expansion and reduct). Let L be a language and let L^+ be another language such that $L \subseteq L^+$ (i.e. L^+ contains all symbols that L contains and they have the same arities). Then we call L^+ an *expansion* of L and we call L a *reduct* of L^+ .

Let \mathbf{M} be an L -structure and let \mathbf{M}^+ be an L^+ -structure such that $\mathbf{M}^+ \upharpoonright_L = \mathbf{M}$ (by this we mean that \mathbf{M} and \mathbf{M}^+ have the same sets of vertices and the interpretations of symbols from L are exactly the same in both structures). Then we call \mathbf{M}^+ an *expansion* of \mathbf{M} and we call \mathbf{M} a *reduct* of \mathbf{M}^+ .

If \mathcal{C} is a class of finite L -structures, we say that \mathcal{C}^+ , a class of finite L^+ -structures, is its *expansion* if for every $\mathbf{A} \in \mathcal{C}$ there is $\mathbf{A}^+ \in \mathcal{C}^+$ which is its expansion and for every $\mathbf{A}^+ \in \mathcal{C}^+$ there is $\mathbf{A} \in \mathcal{C}$ which is its reduct.

Note that reduct and expansion often mean something more general, but for our purposes this definition is sufficient. We say that a class has a *Ramsey expansion* if it has an expansion which is Ramsey.

So far we have only been adding all linear orders, which is clearly a special expansion (and corresponds to adding to the Fraïssé limit the dense linear order with no endpoints which is independent from the rest of the relations). But we have also seen that sometimes it isn't enough.

In 2005, Nešetřil [Neš05] started the *classification programme of Ramsey classes*. Its goal is to classify all possible Ramsey classes, with the classification programme of homogeneous structures offering lists of possible Ramsey classes, or rather base classes for expansions. The first natural question is thus the following: “Does every amalgamation class have a Ramsey expansion?”

The answer to this question is positive, but it is not satisfactory: One can add infinitely many unary predicates to the language and let each vertex have its own predicate. Then every structure has at most one embedding to any other and the Ramsey question becomes trivial. Two different ways of avoiding this cheat have been offered:

Question 1.7.7 (Bodirsky–Pinsker–Tsankov [BPT11]). *Does every amalgamation class in a finite language have a Ramsey expansion in a finite language?*

This question, which is motivated by the area of constraint satisfaction problems, still remains open. The other possible fix is motivated by topological dynamics (see Section 1.7.2). An amalgamation class \mathcal{C} of finite L -structures is *ω -categorical* if it is the age of an ω -categorical structure, or equivalently, if for every n there are only finitely many non-isomorphic structures in \mathcal{C} on n vertices.

Let \mathcal{C} be a class of L -structures and let \mathcal{C}^+ be its expansion. We say that \mathcal{C}^+ is a *precompact* expansion of \mathcal{C} if for every $\mathbf{A} \in \mathcal{C}$ there are only finitely many non-isomorphic $\mathbf{A}^+ \in \mathcal{C}^+$ which are expansions of \mathbf{A} .

Question 1.7.8 (Melleray–Nguyen Van Thé–Tsankov [MNVTT15]). *Does every ω -categorical amalgamation class have a precompact Ramsey expansion?*

This question has been answered negatively by Evans, Hubička, and Nešetřil who proved that a certain ω -categorical structure coming from the Hrushovski construction does not have a precompact Ramsey expansion [EHN19].

1.7.2 The KPT correspondence

The references for this short section are a survey by Nguyen Van Thé [NVT15] and a paper of Zucker [Zuc16]. We will implicitly assume that every topology is Hausdorff.

For a topological group G , a G -flow is a continuous action of G on a compact space X , often denoted as $G \curvearrowright X$.

Definition 1.7.2 (Extremely amenable group). Let G be a topological group. We say that G is *extremely amenable* if every G -flow has a fixed point (i.e. $x \in X$ such that $g \cdot x = x$ for every $g \in G$).

Now we can state the theorem of Kechris, Pestov and Todorčević:

Theorem 1.7.9 (KPT correspondence [KPT05]). *Let \mathbf{M} be a countable homogeneous structure. Then the following are equivalent:*

1. $\text{Aut}(\mathbf{M})$ is extremely amenable; and
2. $\text{Age}(\mathbf{M})$ has the Ramsey property.

We have defined Ramseyness as a property of an amalgamation class, but thanks to the KPT correspondence the Ramsey property is now witnessed directly by the automorphism group of its Fraïssé limit (if it exists), similarly as EPPA. Also in analogy to EPPA, the topological counterpart has not yet proved to be very amenable to being proved directly.

The flow $G \curvearrowright X$ is *minimal* if every orbit is dense, and it is *universal* if for every flow $G \curvearrowright Y$ there is a continuous map $f: X \rightarrow Y$ which respects the G -action (a G -map). It is a fact that every topological group G admits a *universal minimal flow* $M(G)$ (i.e. a minimal G -flow which maps onto every other minimal G -flow) which is unique up to isomorphism. An equivalent definition of extreme amenability is that the universal minimal flow is a singleton, the KPT correspondence thus allows us to compute universal minimal flows for homogeneous structures whose ages are Ramsey.

Remark 1.7.1. More generally, a topological group G is *amenable* if every G -flow admits a left-invariant probability measure. In this setting, extreme amenability amounts to the measure being Dirac. Note that EPPA implies amenability [KR07] and that from having a tame-enough Ramsey expansion, one can often prove amenability by a simple counting argument [AKL14]. Also note that in the locally compact setting, this definition of amenability is equivalent to the standard one. However, in our case, they are no longer equivalent and automorphism groups typically do not satisfy the standard definition.

Let \mathbf{R} be the Rado graph, assume that its vertex set is \mathbb{N} , and put $G = \text{Aut}(\mathbf{R})$. Let $LO(\mathbb{N})$ be the space of all linear orders understood as a subspace of $\{0, 1\}^{\mathbb{N} \times \mathbb{N}}$, the space of all binary relations on \mathbb{N} . Note that $\{0, 1\}^{\mathbb{N} \times \mathbb{N}}$ is compact and that $LO(\mathbb{N})$ is its closed subspace, hence also compact. By Theorem 1.7.4 we know that the class of all finite ordered graphs is Ramsey. Its Fraïssé limit is (\mathbf{R}, \leq) , the Rado graph with a dense linear order \leq with no endpoints. In particular, $\leq \in LO(\mathbb{N})$. Put $X^* = \overline{G \cdot \leq}$, this is again a compact space. Moreover, it is easy to see that both $LO(\mathbb{N})$ and X^* are G -invariant. Kechris, Pestov, and Todorcević [KPT05] proved (for arbitrary homogeneous structures expanded by linear orders, not only for the Rado graph) that the flow $G \curvearrowright X^*$ is minimal if and only if $\text{Age}(\mathbf{R}, \leq)$ has the ordering property with respect to $\text{Age}(\mathbf{R})$.

Nguyen Van Thé later [NVT13] built on the ideas used in the proof of the KPT correspondence and introduced a way of computing the universal minimal flows using Ramsey expansions which are in certain sense *optimal*:

Definition 1.7.3 (Expansion property [NVT13]). Let \mathcal{C} be a class of finite structures and let \mathcal{C}^+ be its expansion. We say that \mathcal{C}^+ has the *expansion property* (with respect to \mathcal{C}) if for every $\mathbf{B} \in \mathcal{C}$ there is $\mathbf{C} \in \mathcal{C}$ such that for every $\mathbf{B}^+ \in \mathcal{C}^+$ and for every $\mathbf{C}^+ \in \mathcal{C}^+$ such that \mathbf{B}^+ is an expansion of \mathbf{B} and \mathbf{C}^+ is an expansion of \mathbf{C} it holds that there is an embedding $\mathbf{B}^+ \rightarrow \mathbf{C}^+$.

The expansion property says that for every small structure \mathbf{B} in the non-expanded class there is a large structure \mathbf{C} in the non-expanded class such that every expansion of \mathbf{C} contains every expansion of \mathbf{B} . The expansion property is a generalization of the *ordering property* studied by Nešetřil and Rödl in the 1970's and 80's [NR75].

Given Fraïssé classes \mathcal{C} and \mathcal{C}^+ such that \mathcal{C}^+ is a precompact expansion of \mathcal{C} , consider their Fraïssé limits \mathbf{M} and \mathbf{M}^+ , put $G = \text{Aut}(\mathbf{M})$, and assume that $\mathbf{M}^+ = (\mathbf{M}, \bar{R}_{\mathbf{M}^+})$, that is, \mathbf{M}^+ adds the relations from the tuple $\bar{R}_{\mathbf{M}^+}$. Moreover, assume that all languages are relational (this is only a technical assumption, every function can be understood as a special relation). Generalising the space $\{0, 1\}^{\mathbb{N} \times \mathbb{N}}$ of all binary relations, we can consider the space of all tuples of relations of the corresponding arity such that $\bar{R}_{\mathbf{M}^+}$ is its member, and we can put $X^* = \overline{G \cdot \bar{R}_{\mathbf{M}^+}}$ which is a compact G -invariant space. Nguyen Van Thé proved the following theorems:

Theorem 1.7.10 (Nguyen Van Thé [NVT13]). *Let \mathbf{M} be a Fraïssé structure, let $\mathbf{M}^+ = (\mathbf{M}, \bar{R}_{\mathbf{M}^+})$ be its precompact relational expansion (not necessarily Fraïssé), and put $X^* = \overline{\text{Aut}(\mathbf{M}) \cdot \bar{R}_{\mathbf{M}^+}}$. Then $\text{Aut}(\mathbf{M}) \curvearrowright X^*$ is minimal if and only if $\text{Age}(\mathbf{M}^+)$ has the expansion property with respect to $\text{Age}(\mathbf{M})$.*

Theorem 1.7.11 (Nguyen Van Thé [NVT13]). *Let \mathbf{M} be a Fraïssé structure, let $\mathbf{M}^+ = (\mathbf{M}, \bar{R}_{\mathbf{M}^+})$ be a Fraïssé precompact relational expansion of \mathbf{M} , and put $X^* = \overline{\text{Aut}(\mathbf{M}) \cdot \bar{R}_{\mathbf{M}^+}}$. Then $\text{Aut}(\mathbf{M}) \curvearrowright X^*$ is the universal minimal flow of $\text{Aut}(\mathbf{M})$ if and only if $\text{Age}(\mathbf{M}^+)$ is Ramsey and has the expansion property with respect to $\text{Age}(\mathbf{M})$.*

Corollary 1.7.12. *Up to bi-definability, there is at most one precompact Ramsey expansion with the expansion property.*

To sum it up, the expansion property is the correct notion of optimality of a Ramsey expansion, and understanding the optimal Ramsey expansion of a Fraïssé structure leads to understanding the universal minimal flow of its automorphism group. (See [EHN19] for an example where no Ramsey expansion is precompact, in that case one can still recover large parts of the theory even for the non-precompact expansion.)

We conclude this section with a sketch of an application of the KPT correspondence. For finite relational languages this proposition can be proved combinatorially (see Bodirsky’s proof [Bod15, Proposition 2.22]) and in a stronger setting where the order will be definable:

Proposition 1.7.13 (Kechris–Pestov–Todorčević [KPT05]). *Let \mathbf{M} be a Fraïssé structure whose age has the Ramsey property. Then $\text{Aut}(\mathbf{M})$ fixes a linear order, that is, there exists a linear order \preceq on the vertices of \mathbf{M} such that for every $g \in \text{Aut}(\mathbf{M})$ and every $x, y \in M$ it holds that $x \preceq y$ if and only if $g(x) \preceq g(y)$.*

Proof. Let $LO(M)$ be the compact space of all linear orders on M . It is easy to see that $\text{Aut}(\mathbf{M})$ acts continuously on it by its standard action: For $L \in LO(M)$ and $g \in \text{Aut}(\mathbf{M})$ we define $g \cdot L$ by $(x, y) \in g \cdot L$ if and only if $(g^{-1}(x), g^{-1}(y)) \in L$. Therefore, as \mathbf{M} is Ramsey, $\text{Aut}(\mathbf{M})$ is extremely amenable and thus this action has a fixed point, which is an order L such that $g \cdot L = L$ for every $g \in \text{Aut}(\mathbf{M})$. \square

1.7.3 Proving the Ramsey property

In Section 1.6.1 we promoted the study of EPPA numbers and alluded to a similarity with Ramsey numbers. While one can indeed draw many parallels, it is important to note that Ramsey numbers of graphs only talk about variants of the original Ramsey theorem with $p = 2$. Even with $p = 3$, there are super-exponential lower bounds, much less when one is concerned with the Ramsey property for classes of structures (contrast this with the lack of any examples where EPPA-witnesses need to be superexponential, see Question 2.0.9). This in particular means that we are unable to show any simple proofs or any (not so) naive attempts. While there are relatively simple proofs of the Ramsey property for linearly ordered graphs (see e.g. [Prö13]), they still build upon difficult Ramsey-like theorems such as the Graham–Rothschild theorem for parameter spaces [GR71].

The presently most effective technique for proving the Ramsey property is the *partite construction* developed by Nešetřil and Rödl in a series of papers since the 1970’s [NR76b, NR77a, NR79, NR81, NR82, NR83, NR84, NR87, NR89, NR90, Neš07]. In particular, Nešetřil and Rödl proved the following theorem.

Theorem 1.7.14 (Nešetřil–Rödl [NR77a, NR77b]). *Let L be a relational language and let \mathcal{C} be a free amalgamation class of L -structures. The class \mathcal{C}^+ consisting of all possible linear orderings of members of \mathcal{C} is Ramsey and has the expansion property with respect to \mathcal{C} .*

Hubička and Nešetřil, using a recursive variant of the partite construction, provided the following strengthening of Theorem 1.7.14 where the language can contain arbitrary functions.

Theorem 1.7.15 (Hubička–Nešetřil [HN19]). *Let L be a language and let \mathcal{C} be a free amalgamation class of L -structures. The class \mathcal{C}^+ consisting of all possible linear orderings of members of \mathcal{C} is Ramsey.*

In the same paper, they proved the following counterpart of Theorem 1.6.4.

Theorem 1.7.16 (Hubička–Nešetřil [HN19]). *Let L be a language, let \mathbf{A} and \mathbf{B} be finite irreducible L -structures, let \mathbf{C}_0 be a finite Ramsey-witness for \mathbf{A} and \mathbf{B} and let $n \geq 1$ be an integer. There is a finite L -structure \mathbf{C} satisfying the following.*

1. \mathbf{C} is a Ramsey-witness for \mathbf{A} and \mathbf{B} .
2. There is a homomorphism-embedding $\mathbf{C} \rightarrow \mathbf{C}_0$.
3. For every substructure \mathbf{D} of \mathbf{C} on at most n vertices there is a tree amalgamation \mathbf{T} of copies of \mathbf{B} and a homomorphism-embedding $f: \mathbf{D} \rightarrow \mathbf{T}$.
4. Every irreducible substructure of \mathbf{C} embeds into \mathbf{B} .

Remark 1.7.2. This formulation is not explicitly stated in [HN19], it was first used for the EPPA result in [HKN22] and afterwards Hubička realised that [HN19] actually does prove this formulation, see [Hub20b]. Note also that from the fact that every Ramsey class fixes a linear order it follows that it does not make sense to state Ramsey results with locally finite (i.e. finite closures of finite sets) permutation groups on the language as a Ramsey expansion would need to fix a linear order on the language as well, hence destroying all permutations. Moreover (and more importantly), one can simulate languages equipped with a permutation group using standard languages with functions or arbitrary arities, at least for finite languages, for infinite ones it becomes slightly more technical.

As we mentioned in Section 1.6.2, Theorem 1.6.6 was directly motivated by the main Ramsey result of [HN19] which we now state (see Section 1.6.2 for the definition of a locally finite subclass):

Theorem 1.7.17 (Hubička–Nešetřil [HN19]). *Let L be a language, let \mathcal{R} be a Ramsey class of finite irreducible L -structures and let \mathcal{K} be a hereditary locally finite subclass of \mathcal{R} with the strong amalgamation property. Then \mathcal{K} is Ramsey.*

1.7.4 How to apply Theorem 1.7.17 in practice?

In practice, Theorem 1.7.17 is the most used formulation. The cookbook for using it is as follows: Start with an amalgamation class \mathcal{K} of finite L -structures. If it does not have strong amalgamation, it is because of having non-trivial algebraic closures. In this case, expand it by functions to explicitly represent these closures, thereby turning it into a strong amalgamation class, so we can assume that \mathcal{K} has the strong amalgamation property.

Next, we need to identify a suitable Ramsey super-class \mathcal{R} , which most often can be chosen to be the class of all finite linearly ordered L -structures. Note that for this to work, L needs to already contain a binary symbol which is interpreted in \mathcal{K} as a linear order, but as we have seen in Proposition 1.7.13, this is more or less necessary.

Finally, we need to understand obstacles to completions. Let \mathbf{C} be a finite L -structure such that every irreducible substructure of \mathbf{C} embeds into \mathbf{B} , which in particular means that \mathbf{C} only contains those relations which appear in \mathbf{B} and that if, in \mathbf{B} , some relations are for example symmetric or some relations are mutually exclusive, then they are also like that in \mathbf{C} . Try to define some canonical completion \mathbf{C}' of \mathbf{C} in \mathcal{K} and, assuming that the definition failed, look at the failure point, derive the reason for it in \mathbf{C} and hope that there will be a finite bound n on the size of the reason.

For example, if \mathbf{B} is a metric space then \mathbf{C} is an edge-labelled graph which only uses those distances which appear in \mathbf{B} , let d be the smallest one and D the largest one. As we have seen in Section 1.6.2, \mathbf{C} has a completion to a metric space if and only if it contains no non-metric cycles, and the largest non-metric cycle has $n = \lceil \frac{D}{d} \rceil$ vertices.

In the previous two paragraphs, we did not tell the whole story because we ignored the fact that our structures need to be linearly ordered. Hence, \mathbf{B} is actually a linearly ordered metric space and, in \mathbf{C} , the relation \leq is antisymmetric and reflexive, but it need not be transitive. In particular, \mathbf{C} can contain arbitrarily long sequences of vertices c_0, c_1, \dots, c_{m-1} such that $c_i \leq c_{i+1}$ for every $0 \leq i < m-1$, $c_{m-1} \leq c_0$ and there are no other \leq relations between these vertices (in other words, oriented cycles in the binary relation \leq). Clearly, such structures have no completion to a linear order.

This is, however, where the other condition of a locally finite subclass comes into play, the existence of a homomorphism-embedding $\mathbf{C} \rightarrow \mathbf{C}_0$, because \mathbf{C}_0 is linearly ordered and thus in fact \mathbf{C} contains no such oriented cycles without a completion to a linear order (because the existence of such a cycle is preserved by a homomorphism-embedding). In general, one can often first study obstacles to completions for the unordered reducts and obtain results which are easily applicable in the ordered setting.

However, one needs to be a bit careful with choosing the correct unordered reducts. Let L be the language consisting of two binary relations \leq and E and let \mathcal{C} be the class of all finite structures on which \leq is a linear order and E is an equivalence relations. It is easy to see that \mathcal{C} is a strong amalgamation class consisting of irreducible structures. However, when we forget the order relation, many structures become reducible. This has serious consequences: If \mathbf{C} is a graph then it always has a completion to an equivalence: Just consider the complete graph on the same vertex set. On the other hand, if \mathbf{C} is an L -structure such that \leq is reflexive and antisymmetric and E is reflexive and symmetric then there might be some pairs of vertices which are in the \leq relation but not in the E relation, hence we can no longer simply complete \mathbf{C} to a complete graph. It is still usually a good heuristic to solve the unordered completion problem first, but substituting \leq with a complete binary symmetric relation in the original class – in the incomplete structures it then serves as a sort of “non-relation”.

Equivalences

In the case of equivalences, it is more convenient (but bi-definable) to do the following: Let L consist of two binary relations E and N and let \mathcal{K} consist of all finite structures on which E is an equivalence relation and N is its complement.

Clearly, \mathcal{K} is a strong amalgamation class which consists of irreducible structures. In the spirit of the previous paragraphs, let \mathbf{C} be a graph with edges labelled by E and N (formally, it should be a structure on which E is reflexive, N irreflexive, both relations are symmetric and mutually exclusive – reflexivity of E is, however, not interesting, because all vertices of all structures in \mathcal{K} are in E with themselves, and thus so are all vertices in \mathbf{C}). We want to see if we can complete \mathbf{C} to a complete graph on which the labels E describe an equivalence relation.

Let ℓ be the labelling of edges of \mathbf{C} . The natural attempt would be to define a new labelling ℓ' on pairs of vertices of \mathbf{C} such that $\ell'(u, v) = E$ if and only if there exists a path in \mathbf{C} between u and v such that all edges of the path are labelled by E , and $\ell'(u, v) = N$ otherwise. Clearly, whenever $\ell'(u, v) = E$ then it needs to be so in any completion of \mathbf{C} due to transitivity. But (\mathbf{C}, ℓ') is a completion of \mathbf{C} if and only if $\ell \subseteq \ell'$. If $\ell(u, v) = E$ then $\ell'(u, v) = E$ which is witnessed by the path uv . However, it might happen that $\ell(u, v) = N$ but there is a path from u to v with all edges labelled by E and hence $\ell'(u, v) = E$. Clearly, if this happens then \mathbf{C} does not have any completion into \mathcal{K} . Furthermore, the shortest witnessing path for u and v can be arbitrarily long, hence there is no bound on the size of the largest obstacle.

A few paragraphs above we have seen an analogous situation happen with the linear order and the solution there was to invoke another condition from local finiteness. This is, however, not applicable in this case. Instead, we need to sidestep the issue: Let L' be a language consisting of one unary function f and two unary relations O and I . Let \mathcal{D} be the class of all finite L' -structures such that O and I form a partition of the vertex set, the domain of f is O and the image of f is a subset of I . We will call the vertices in the O relation the *original* vertices and the vertices in the I relation the *imaginary* vertices. Notice that there is a simple correspondence between structures from \mathcal{D} and structures from \mathcal{K} : Given $\mathbf{A} \in \mathcal{D}$, define a new L -structure with vertex set O , putting $uv \in E$ if and only if $f(u) = f(v)$ and $uv \in N$ otherwise. On the other hand, given $\mathbf{A} \in \mathcal{K}$, let I be the set of all equivalence classes of \mathbf{A} and define a new L' -structure with vertex set $A \cup I$ such that all vertices from A are in the relation O , all vertices from I are in the relation I and for every $u \in A$, $f(u)$ is the equivalence class whose u is a member in \mathbf{A} .

Remark 1.7.3. The correspondence is rather strong because it preserves embeddings and their compositions, it is *functorial*.

Note that the completion problem is trivial for \mathcal{D} – any L' -structure where O and I form a partition of the vertex set, the domain of f is O and the image of f is a subset of I , which are conditions on irreducible substructures, is automatically a member of \mathcal{D} .

We have now seen two examples of a failure of local finiteness, both due to transitivity: One was because of orders, the other because of equivalences. The way to deal with orders was to use the existence of a homomorphism-embedding into an already ordered structure, while the way to deal with equivalences was to use unary functions to explicitly represent the equivalence classes. Both these methods are, by now, standard heuristics used when applying Theorem 1.7.17:

- For example, the class of all finite posets is Ramsey when equipped with a linear extension [NR84, PTW85] which can be proved using Theorem 1.7.17

(cf. also Chapter 10) by letting \mathcal{K} be the class of all finite posets with a linear extension and \mathcal{R} be the class of all linearly ordered directed graphs.

- Presence of definable equivalences is a hindrance for local finiteness. When applying Theorem 1.7.17, one needs to first understand these definable equivalences and then *eliminate imaginaries* – introduce imaginary vertices representing the equivalence classes, link them with the original vertices using functions and, possibly, add more structure on the imaginary vertices to obtain the strong amalgamation property. Several examples are shown in [HN19], see also the Mater thesis of the author for additional examples of various generalised metric spaces [Kon19]. Note that elimination of imaginaries is a model-theoretic concept with a lot of importance.

Remark 1.7.4. Note the similarity between the shortest path completion for metric spaces and the completion we defined for equivalences. This is not a coincidence, because if we consider metric spaces with distances $\{1, 3\}$ then distance 1 describes an equivalence relation and distance 3 its complement. In general, the *ultrametric spaces* are those where the allowed distances are so far apart that the triangle inequality can be equivalently stated as $\max(x, y) \geq z$ and they describe families of refining equivalences. See [Kon19] for more context and further generalisations.

Ordered equivalences?

For the third time in this section we are now going to contradict our own proposition that one should treat the order separately when applying Theorem 1.7.17 (and afterwards keep standing by it). Recall that several paragraphs ago we described a correspondence between the class \mathcal{K} of all finite equivalences and the class \mathcal{D} of structures with two unary relations and one unary function. This correspondence no longer works if we expand the classes by free linear orders: Given a linearly ordered structure with an equivalence relation, we need to find a way to define an order on the equivalence classes to produce a linearly ordered structure from \mathcal{D} . Moreover, we need to preserve embeddings: For example, ordering the equivalence classes based on their least member is not good enough because a substructure might not contain some of these least members and define the order on the equivalence classes differently, hence no longer being a substructure in \mathcal{D} .

The solution here is to only consider linear orders on \mathcal{K} which are *convex*, that is, the equivalence classes form intervals. In this way we get a pair of functors between the class \mathcal{K}^+ of all finite convexly ordered equivalences and \mathcal{D}^+ of all finite convexly ordered L' -structures and hence we can finish the proof of Ramsey property for \mathcal{K}^+ : Invoke Theorem 1.7.17 to prove the Ramsey property for \mathcal{D}^+ . Given $\mathbf{A}, \mathbf{B} \in \mathcal{K}^+$, use the Ramsey property for their images in \mathcal{D}^+ to obtain a Ramsey witness $\mathbf{C}^* \in \mathcal{D}^+$ and transfer it back to $\mathbf{C} \in \mathcal{K}^+$. Thanks to functoriality of the maps between \mathcal{K}^+ and \mathcal{D}^+ one can easily show that \mathbf{C} is a Ramsey-witness for \mathbf{A} and \mathbf{B} .

In general, such a situation always happens in the presence of equivalences, where one needs to ensure that the expansion is rich enough to also order the equivalence classes. Sometimes this means expanding by convex orders only, at other times it means that the expansion consists of several orders (see e.g. [Bra17])

or [Kon19]). And sometimes, when the equivalences are on tuples instead of singletons, one needs to expand by a relation describing a linear order of the tuples.

Remark 1.7.5. Note that in the above case the equivalence relation had infinitely many classes in the Fraïssé limit. If it only had finitely many then their respective imaginaries would be in the algebraic closure of the empty set and in order to have strong amalgamation one would have to add them as constants. In such a case it is more convenient to expand the language by one unary relation for every equivalence class and distinguish them by unary relations instead of unary functions.

The cookbook

The current cookbook method for obtaining a Ramsey expansion for a class \mathcal{C} of L -structures is to start by understanding the definable equivalence relations on the structures, then eliminate imaginaries, understand what kind of orders the structures contain and, with all this knowledge, try to understand obstacles to completions. If life is fair they will turn out to have bounded size. In this case, one possibly needs to expand the class by convex orders in order to be able to get a linear order on all original vertices as well as imaginaries, and with a bit of luck there will be functorial maps between the original class expanded by convex orders (and relations distinguishing equivalence classes of finite order) and the class further enriched by imaginary vertices. Theorem 1.7.17 will then do all the heavy lifting.

Except that life sometimes is not fair. Remember for example the class \mathcal{C} of all finite two-graphs, which was an interesting example already in Section 1.6. There are no orders and no relevant equivalences, but there are arbitrarily large obstacles to completion. It turns out that the correct Ramsey expansion of two-graphs with the expansion property adds not only a linear order but also a graph from the represented switching class (see Section 5.7). Since the two-graph structure can be defined from the graph structure, the resulting Ramsey expansion in the end turns out to be the class of all finite ordered graphs.

Remark 1.7.6. Showing that any Ramsey expansion of two-graphs needs to fix a graph from the switching class can be done purely combinatorially (and it is showed in Section 5.7), but it is also a nice example of an application of the KPT correspondence: Assume that \mathbf{T}^+ is the Fraïssé limit of a Ramsey class which has the expansion property with respect to the class of all finite two-graphs, and let \mathbf{T} be the Fraïssé limit of the class of all finite two-graphs. Assume that \mathbf{T}^+ and \mathbf{T} share the vertex set and the hyperedge relation. Let X be the space of all graphs on T which give rise to the hyperedge relation on \mathbf{T} (that is, the class of all graphs on T from the represented switching class). It is straightforward to show that X is compact and thus the natural action of $\text{Aut}(\mathbf{T}^+)$ on X has a fixed point by Section 1.7.2.

While the Ramsey expansions of the vast majority of examples can be found using the cookbook method, two-graphs are an example that sometimes the cookbook fails. And in that case there are no general methods better than playing with the class for long enough and hoping to think of a good expansion. For two-graphs, this is not that hard once one knows that the generic two-graph is a

reduct of the random graph. For the counterexample to Question 1.7.8 found by Evans, Hubička, and Nešetřil [EHN19] (and many related classes) the relevant combinatorial properties of Hrushovski constructions have been already known, but a priori the expansion is highly non-trivial. (Also note that, as the expansion is not relational, finding the optimal subclass of all linear orderings was much more involved.) Let us remark that in Section 1.8.11, we will see that methods used for studying big Ramsey degrees can be applied also to study the Ramsey property and that in the area of big Ramsey degrees, there are heuristics for computing the exact big Ramsey degrees which give hope that one might get another tool for finding Ramsey expansions, once these techniques are more developed (in particular, they reveal that a Ramsey expansion of two-graphs needs to fix a graph).

There are still some classes for which their Ramsey expansion (or the lack of a precompact one) is not known. One such example is the class of all finite groups (see Question 2.0.13). Another example are the H_4 -free 3-hypertournaments (see Question 7.4.4) and there are a couple of other classes where the current methods fail in one way or another (see e.g. [Hub20b]).

Remark 1.7.7. The selection of two-graphs and Hrushovski constructions as problematic examples was deliberate. There are a few other classes on which the cookbook method does not work but, unlike the two presented examples, there is no known proof of the Ramsey property of their expansion using Theorem 1.7.17 at all. One such example is the class of all finite Boolean algebras (which is Ramsey with a linear order on the atoms using the Graham–Rothschild theorem [GR71]), another example is the class of all finite C -relations [Bod15], yet another example are semilattices [Sok15]. One can argue that these examples are not as relevant as the examples we gave because the category of finite linear orders is not the natural “base category” for them. (For example, Ramseyness of Boolean algebras is essentially just a rephrasing of the dual Ramsey theorem.) See also Section 1.7.8.

Similarly as in the above remark, one can also argue that the open problem of Ramsey expansion of the class of all finite groups is less relevant than the class of all finite H_4 -free 3-hypertournaments, which seems to be a pure, combinatorial, relational class as opposed to the much more algebraic class of all finite groups.

1.7.5 Examples

While we have already seen many examples, it is still a good idea to have a (substantially incomplete) list of various examples:

- The class of all finite linear orders is Ramsey by Ramsey’s theorem and the class of all finite partial orders is Ramsey when expanded by a linear extension [NR84, PTW85].
- If \mathcal{C} is a free amalgamation class then the class of all finite linear orderings of \mathcal{C} is Ramsey by Theorem 1.7.15.
- The Ramsey property for finite linearly ordered metric spaces has been proved by Nešetřil [Neš07]. Optimal Ramsey expansions of S -Urysohn spaces for closed S have been proved by Hubička and Nešetřil [HN19] and

for non-closed S they have been obtained by Hubička, Nešetřil, Sauer, and the author in a yet unpublished draft [HKNS20]. See also [DR12].

- The most general result on Ramsey expansions for generalised metric spaces is in the Master thesis of the author [Kon19], see also [HKN21]. Note, in particular, that it implies (and is inspired by) Braunfeld’s Ramsey expansions of Λ -ultrametric spaces [Bra17].
- Ramsey expansions of metrically homogeneous graphs have been obtained by Aranda, Bradley-Williams, Hubička, Karamanlis, Kompatscher, Pawliuk, and the author [ABWH⁺17b, ABWH⁺21, ABWH⁺17a].
- The Ramsey expansion of the class of all finite two-graphs is the class of all finite linearly ordered graphs [EHKN20], see also Chapter 5.
- Ramsey expansions of all classes from Cherlin’s list of homogeneous directed graphs have been obtained by Jasiński, Laflamme, Nguyen Van Thé, and Woodrow [JLNVTW14]. This list in particular includes tournaments, semigeneric tournaments and n -partite tournaments.
- Evans, Hubička, and Nešetřil studied Ramsey expansions of certain Hrushovski constructions [EHN19]. They proved that there are no precompact Ramsey expansions but were still able to obtain Ramsey expansions with the expansion property. From the structural point of view, these classes exhibit quite an exotic behaviour.

Giving non-examples is complicated due to the nature of the question. The example of Evans, Hubička, and Nešetřil [EHN19] is the only presently known counterexample to Question 1.7.8 and there are no known counterexamples to Question 1.7.7.

There are several examples of classes for which their Ramsey expansion is known but there is no known proof using Theorem 1.7.17. These include Boolean algebras, semilattices, or C -relations. There are also various related Ramsey-type results which do not talk about classes of structures (see Section 1.7.8).

Finally, there are several classes for which finding their Ramsey expansions remains an open problem. They can be split into two families. One family consists of finite groups, skew-symmetric bilinear forms, Euclidean metric spaces, graphs of girth $\geq k$ for $k \geq 5$, or partial Steiner systems without short cycles, see e.g. [Hub20b] for more discussion. The classes from this family have complicated closures and while Theorem 1.7.17 does admit functions, it does not seem to be particularly amenable to complicated closure structures.

The other family consists of the H_{n+1} -free n -hypertournaments for $n \geq 3$ (with $n = 3$ likely being the key example). These classes have been identified more recently (see Chapter 7) and as such have not seen as much scrutiny as the other open cases. It seems that the standard cookbook method does not yield a Ramsey expansion (at least the author has not succeeded in finding it), and neither has the author been successful in conjuring an ingenious Ramsey expansion or crowdsourcing one at various workshops and conferences. While it is possible that there is such a Ramsey expansion and, after finding it, these classes will at best earn their spot in the naughty list of cookbook-exceptions,

they might also reveal a new cookbook heuristic, or even give an example where Theorem 1.7.17 fails and one cannot blame algebraicity for it. Or they might give a counterexample to Question 1.7.7.² See Question 7.4.4 and the discussion around it (and around its copy in Chapter 2) for more details.

1.7.6 Ramsey degrees

Let \mathcal{C} be a class and let \mathcal{C}^+ be its Ramsey expansion. Pick $\mathbf{A}, \mathbf{B} \in \mathcal{C}$ and an integer $k \geq 1$. If \mathcal{C} is not Ramsey, we know that we cannot hope, in general, for $\mathbf{C} \in \mathcal{C}$ such that $\mathbf{C} \rightarrow (\mathbf{B})_k^{\mathbf{A}}$. However, we can get something weaker: Let $\mathbf{B}_0^+ \in \mathcal{C}^+$ be an arbitrary expansion of \mathbf{B} and enumerate all expansions of \mathbf{A} as $\mathbf{A}_1^+, \dots, \mathbf{A}_\ell^+$. By induction, define a sequence $\mathbf{B}_1^+, \dots, \mathbf{B}_\ell^+$ such that $\mathbf{B}_i^+ \rightarrow (\mathbf{B}_{i-1}^+)_k^{\mathbf{A}_i^+}$ for every $1 \leq i \leq \ell$. Let $\mathbf{C} \in \mathcal{C}$ be the reduct of \mathbf{B}_ℓ^+ . We claim that

$$\mathbf{C} \rightarrow (\mathbf{B})_{k,\ell}^{\mathbf{A}}.$$

Indeed, let c be an arbitrary colouring of embeddings $\mathbf{A} \rightarrow \mathbf{C}$ by k colours. This, in particular, gives us a colouring $\binom{\mathbf{B}_\ell^+}{\mathbf{A}_\ell^+} \rightarrow k$, and by the property of \mathbf{B}_ℓ^+ we can find a copy of $\mathbf{B}_{\ell-1}^+$ on which all copies of \mathbf{A}_ℓ^+ have the same colour. Next use c to get a colouring $\binom{\mathbf{B}_{\ell-1}^+}{\mathbf{A}_{\ell-1}^+} \rightarrow k$, use the Ramsey property again and continue. In the end we will get a copy of \mathbf{B}_0^+ on which all copies of \mathbf{A}_i^+ are monochromatic for each i (but a copy of \mathbf{A}_i^+ can have a different colour from a copy of \mathbf{A}_j^+ if $i \neq j$). Clearly, after taking the reduct, we get a copy of \mathbf{B} in \mathbf{C} on which the colours of copies of \mathbf{A} only depend on the isomorphism type of their expansion in \mathbf{B}_0^+ . As there are only ℓ non-isomorphic expansions, we indeed get at most ℓ different colours, and this ℓ does not depend on \mathbf{B} or the number of colours k . If this happens, we say that the *Ramsey degree of \mathbf{A} in \mathcal{C}* is ℓ . Note that, in fact, one can equivalently develop the whole theory using Ramsey degrees instead of expansions, it turns out that if we have finite Ramsey degrees of all structures in a class then this is witnessed by an expansion with the expansion property. [NVT13, Zuc16, MZ22]

1.7.7 Category theory

The subject of Ramsey theory is rich and wide and this small subsection certainly does not aim to survey all of it, from Ramsey numbers, through various discrete geometrical partition theorems, density Ramsey theorems, to very infinite Erdős–Rado-type theorems or Ramsey ultrafilters. (See for example one of the several books on the topic [Neš95, NR12, Prö13, GRS90].) Instead, we briefly mention some of the variants of the Ramsey property which are relevant to the thesis contents or to the author’s interests.

Note that the partition arrow $C \rightarrow (B)_{k,\ell}^A$ is categorical in essence: A, B and C are objects of some category and the arrow denotes the statement that for an arbitrary colouring $c: \text{hom}(A, C) \rightarrow k$ there exists a morphism $f \in \text{hom}(B, C)$ such that $|f \circ \text{hom}(A, B)| \leq \ell$. Similarly, the joint embedding property and

²Or it might turn out that I have made a mistake somewhere and they are not interesting at all.

the amalgamation property are categorical properties and both Fraïssé theory and the KPT correspondence can also be phrased in a categorical language (see e.g. [Kub14, Maš21]). We talked about functors when transferring Ramsey properties and even the partite construction can be axiomatized for general categories. One can use this language or these ideas to formulate abstract theorems for Ramsey property transfer, see for example [Maš18].

1.7.8 Dual Ramsey theorem(s)

If we take the Ramsey theorem and reverse the arrows, we get an (equivalent) statement where the objects are still finite linear orders but instead of increasing injections, the maps are now *rigid surjections* where moreover the pre-image of each element is an interval. Here, if (A, \leq) and (B, \leq) are linear orders and $f: A \rightarrow B$ is a surjection, we say that it is a *rigid surjection* if for every $x < y \in B$ we have that $\min\{a \in A : f(a) = x\} < \min\{a \in A : f(a) = y\}$.

Note that rigidity of the surjections is a necessary condition for Ramseyness as otherwise one can colour surjections based on the permutation of the minimum elements of pre-images. However, it turns out that the condition on pre-images of elements being intervals is not necessary and hence one can strengthen the simple dualization to obtain the following theorem.

Theorem 1.7.18 (Dual Ramsey theorem, Graham–Rothschild [GR71]). *For every pair of finite linear orders \mathbf{A} and \mathbf{B} and integer $k \geq 1$ there exists a finite linear order \mathbf{C} such that for any colouring c of rigid surjections from \mathbf{C} to \mathbf{A} by k colours there exists a rigid surjection $f: \mathbf{C} \rightarrow \mathbf{B}$ such that all compositions $g \circ f$, where $g: \mathbf{B} \rightarrow \mathbf{A}$ is a rigid surjection, have the same colour.*

This theorem is in fact a special case of the Graham–Rothschild theorem for parameter spaces which we will discuss in Section 1.8.4.

Note that there is an alternative, again categorical, view on structural Ramsey theory, which uses *indexed categories*. Briefly, instead of working in a category where the objects are structures and morphisms are structural embeddings, we can consider the “base category” to be the category of finite linear orders with increasing injections. An n -ary relation on an object A can then be modeled as a subset of the arrows from O_n to A where O_n is the linear order on n vertices. Thus, a relational structure in a given language is a basic object A together with a subset of arrows from O_n to A for every n -ary relation, morphisms in the indexed category are standard morphisms $A \rightarrow B$ which moreover preserve the “relation arrows” after composition.

This point of view allows one to distinguish the *unstructured* Ramsey theorems (for example Ramsey’s theorem or the Dual Ramsey theorem) from their structural (sometimes called *induced*) variants and it also gives good notions of structural Ramsey theories with a different base theorem. See for example [Lee73, Prö85, FGR87, Sol10]. In particular, combined with the projective Fraïssé theory developed by Irwin and Solecki [IS06], these approaches have been very useful in the study of homeomorphism groups of various continua, see for example [BK17, BK19].

1.7.9 Approximate Ramsey property

Another possible unstructured Ramsey theorems are the approximate ones: The objects are usually some finite-dimensional Banach spaces and we are not looking for a monochromatic copy, but for a copy for which there exists a colour such that subspaces of the copy are ε -close to spaces of the chosen colour. This is closely connected to continuous logic, see for example [BLALM17, MT14], see also Section 1.8.13.

1.7.10 Conclusion

Even though this thesis only contains one paper which primarily is a structural Ramsey theory paper, structural Ramsey theory has been at the beginning of the author's mathematical journey and it still is the unifying centrepiece. The author has several papers from this area (e.g. [ABWH⁺17a, ABWH⁺21, ABWH⁺17b, HKN21, HKN18, HKNS20]) as well as both his Bachelor [Kon18] and Master [Kon19] theses. There has been mutual influence between the results from all the areas discussed in this thesis, and trying to study them together has been very fruitful. For example, the Hubička–Nešetřil theorem for Ramsey classes has influenced Theorem 1.6.4 which in turn suggested a new, more refined formulation for the Hubička–Nešetřil theorem itself. For another example, the method of valuation functions, which was developed in the EPPA context, has greatly influenced many of our results in the big Ramsey degrees area (see Section 1.8), and, in turn, there is a hope that the methods for studying big Ramsey degrees will lead to a new general theorem for studying Ramsey expansions (see Section 1.8.11).

1.8 Big Ramsey degrees

The epistemic status of this section is different from the previous ones. The area of big Ramsey degrees has been seeing rapid development in the past several years, there are still many open problems believed to be within reach and/or being actively worked on. Therefore, it is reasonable to expect that one will be able to develop better intuition for the area and, based on it, find a better narrative for an introduction of the kind that this section is trying to accomplish. Also, the terminology and notation might change.

This is of course possible also for the other sections (in fact, one should hope for such changes as they are exciting and move things forward), but for this one it is much more likely to happen in the near future. For these reasons, there are also no recent surveys or books which could serve as good references for this chapter with the notable exception of Dobrinen's recent ICM survey [Dob21].

In Section 1.7 we were interested in generalising Ramsey's theorem to for various amalgamation classes of finite structures. In this section, we will try to generalise the infinite Ramsey's theorem in the same spirit:

Theorem 1.8.1 (Infinite Ramsey's theorem [Ram30]). *For every finite linear order \mathbf{A} and every integer $k \geq 1$ it holds that*

$$(\omega, \leq) \longrightarrow ((\omega, \leq))_k^{\mathbf{A}},$$

where (ω, \leq) is the order of natural numbers.

The first example coming to mind is the set of rational numbers with its order (\mathbb{Q}, \leq) . This is closely related to (ω, \leq) , in particular, the age of both of them is the class of all finite linear orders. For colouring vertices, we do get a Ramsey result:

Observation 1.8.2. *Let \mathbf{A} be the linear order with one vertex. Then*

$$(\mathbb{Q}, \leq) \longrightarrow ((\mathbb{Q}, \leq))_2^{\mathbf{A}}.$$

Proof. Let us say that the colouring uses colours red and blue. If there is a blue interval then clearly we have found a copy of rationals in the blue colour. If no interval is fully blue then the red colour is dense and hence, again, we can embed rationals into the red colour. \square

If a structure \mathbf{M} has an infinite Ramsey theorem for colouring vertices, we say that it is *indivisible*. This proof of indivisibility works for many unconstrained structures: Try to embed the structure in one colour (in this case red), and if it fails it means that some extensions exist only in the other colour, and because the structure has no constraints, it will follow that we can embed the whole structure into the other colour.

Nevertheless, indivisibility for (\mathbb{Q}, \leq) is as far as one can go:

Observation 1.8.3 (Sierpiński [Sie33]).

$$(\mathbb{Q}, \leq) \not\rightarrow ((\mathbb{Q}, \leq))_2^{\mathbf{A}}$$

for \mathbf{A} being the 2-vertex linear order.

Proof. Let \preceq be an arbitrary order on \mathbb{Q} of order-type ω . Colour a pair of vertices blue if \preceq and \leq agree on the pair and red otherwise. Since \leq is dense it follows that there is no copy of \mathbb{Q} on which \leq and \preceq agree or are the inverse of each other. Consequently, every copy of \mathbb{Q} will attain both colours. \square

This motivates the following definition:

Definition 1.8.1. Let \mathbf{M} be a structure and \mathbf{A} be its finite substructure. A finite surjective colouring $c: \binom{\mathbf{M}}{\mathbf{A}} \rightarrow k$ is *persistent* (also called *unavoidable*) if for every embedding $f: \mathbf{M} \rightarrow \mathbf{M}$ it follows that $c \upharpoonright \binom{f(\mathbf{M})}{\mathbf{A}}$ is also surjective.

Complementing Sierpiński’s colouring, Galvin much later proved that this is, in a way, the only problem:

Theorem 1.8.4 (Galvin [Gal68, Gal69]).

$$(\mathbb{Q}, \leq) \longrightarrow ((\mathbb{Q}, \leq))_{k,2}^{\mathbf{A}}$$

for every $k \geq 1$ and for \mathbf{A} being the 2-vertex linear order.

In fact, Galvin proved something stronger: In any k -colouring of pairs of vertices of rationals one can find a copy of (\mathbb{Q}, \leq) on which the colouring “looks like” Sierpiński’s colouring. Formally, he proved that Sierpiński’s colouring is universal:

Definition 1.8.2. Let \mathbf{M} be a structure and \mathbf{A} its finite substructure. A finite surjective colouring $c: \binom{\mathbf{M}}{\mathbf{A}} \rightarrow k$ is *universal* if for every finite colouring $c': \binom{\mathbf{M}}{\mathbf{A}} \rightarrow m$ there exists an embedding $f: \mathbf{M} \rightarrow \mathbf{M}$ and a map $\pi: k \rightarrow m$ such that for every $e \in \binom{\mathbf{M}}{\mathbf{A}}$ it holds that $c'(f \circ e) = \pi(c(f \circ e))$.

Informally, for every colouring we can find a copy on which this colouring looks like the universal one after renaming (and potentially identifying) some colours. Note that identifying colours is necessary as we can, for example, let c' be the identity colouring.

Analogously to the notion of Ramsey degrees from Section 1.7.6 (which we will sometimes call *small Ramsey degrees* to avoid confusion), one can say that the *big Ramsey degree* of a finite structure \mathbf{A} in a structure \mathbf{M} is the least $\ell \in \omega \cup \{\omega\}$ such that $\mathbf{M} \longrightarrow (\mathbf{M})_{k,\ell}^{\mathbf{A}}$ for every $k \in \omega$. With this terminology, indivisibility means that the big Ramsey degree of a vertex is equal to 1 and Sierpiński’s and Galvin’s results together show that the big Ramsey degree of a pair of vertices in (\mathbb{Q}, \leq) is equal to 2. We say that \mathbf{M} *has finite big Ramsey degrees* if the big Ramsey degree of every finite substructure of \mathbf{M} is finite.

One can show that the big Ramsey degree of \mathbf{A} in \mathbf{M} is ℓ if and only if there is a colouring $c: \binom{\mathbf{M}}{\mathbf{A}} \rightarrow \ell$ which is both persistent and universal. We call such a colouring the *big Ramsey colouring of \mathbf{A} in \mathbf{M}* (it is clearly unique up to a permutation of the values). [Zuc19] Note that the same holds for small Ramsey degrees, but recall that, in fact, we have something stronger: Knowing that all small Ramsey degrees are finite, one gets a small Ramsey expansion, which means that the small Ramsey colouring can be chosen to cohere together, to all come from the expansion. This is open for big Ramsey degrees, see Section 1.8.10.

1.8.1 Big Ramsey degrees of (\mathbb{Q}, \leq)

We have already seen that the big Ramsey degree of a vertex of (\mathbb{Q}, \leq) is 1 and that the big Ramsey degree of a pair of vertices is 2. What about larger finite linear orders?

Let us first revisit Sierpiński's idea for colouring pairs in a slightly different light: Pick an enumeration \preceq of \mathbb{Q} . Given $a < b \in \mathbb{Q}$, then enumeration allows us to address the first vertex m in the enumeration such that $a \leq m \leq b$, and we can colour the pair $\{a, b\}$ based on the isomorphism type of the structure $(\{a, b, m\}, \leq, \preceq)$. There are four such types: If $a = m$ or $b = m$ then this structure has just two vertices a and b . We know that $a < b$, so there are two cases, either $a \preceq b$ or $b \preceq a$. Otherwise $a, b \neq m$. In this case we know that $m \preceq a, b$ and that $a < m < b$, hence there are again two cases: $a \preceq b$ or $b \preceq a$.

Let χ_2 be the colouring of pairs of vertices (\mathbb{Q}, \leq) which we described above. This is a universal colouring, and there is a copy of (\mathbb{Q}, \leq) inside (\mathbb{Q}, \leq) which omits the two isomorphism types where $a = m$ or $b = m$, which recovers Galvin's result.

Trees

Note that given an enumeration \preceq of (\mathbb{Q}, \leq) , we actually get a tree and that the isomorphism types correspond to shapes of subtrees of the tree. In order to formalize it, we need to introduce some terminology:

Definition 1.8.3 (Trees). A *tree* is a (possibly empty) partially ordered set $(T, <_T)$ such that, for every $t \in T$, the set $\{s \in T : s <_T t\}$ is finite and linearly ordered by $<_T$. All nonempty trees we consider are *rooted*, that is, they have a unique minimal element called the *root* of the tree. An element $t \in T$ of a tree T is called a *node* of T and its *level*, denoted by $|t|_T$, is the size of the set $\{s \in T : s <_T t\}$. (Note that the root has level 0.)

Given a tree T and nodes $s, t \in T$ we say that s is a *successor* of t in T if $t \leq_T s$. The node s is an *immediate successor* of t in T if $t <_T s$ and there is no $s' \in T$ such that $t <_T s' <_T s$. We denote the set of all successors of t in T by $\text{Succ}_T(t)$ and the set of immediate successors of t in T by $\text{ImmSucc}_T(t)$. We say that the tree T is *finitely branching* if $\text{ImmSucc}_T(t)$ is finite for every $t \in T$.

For $s, t \in T$, the *meet* $s \wedge_T t$ of s and t is the largest $s' \in T$ such that $s' \leq_T s$ and $s' \leq_T t$. A *subtree* of a tree T is a subset T' of T viewed as a tree equipped with the induced partial ordering such that $s \wedge_{T'} t = s \wedge_T t$ for each $s, t \in T'$.

Note that our notion of a subtree differs from the standard terminology, since we require the additional condition about preserving meets. An example of a tree is the tree $(\{0, 1\}^{<\omega}, \sqsubseteq)$ of all finite binary sequences ordered by end-extension. It is a binary tree, its root is the empty sequence \emptyset which has two immediate successors, the sequence (0) and the sequence (1).

Let us now return to Sierpiński's idea. Let \preceq be an enumeration of \mathbb{Q} . Let v_0 be its first (that is, \preceq -least) vertex. Then v_0 splits \mathbb{Q} into two *types* – those vertices which are $<$ -smaller than v_0 and those which are $<$ -larger than v_0 (technically, there is a third type consisting of v_0 itself but let us ignore that one). Let v_1 be the second vertex of \mathbb{Q} and let us assume that $v_0 < v_1$. It splits one of the two types into two and we have in total three types *over* $\{v_0, v_1\}$ – those vertices

which are smaller than v_0 , the open interval between v_0 and v_1 and those which are larger than v_1 . Next do the same with v_2 and so on.




Figure 1.4: The beginning of the tree of 1-types of (\mathbb{Q}, \leq) . In the tree, we labelled by v_i the type to which v_i belonged just before we extended the initial segment by v_i .

In the end we get a tree describing how types were refining as we were discovering larger and larger initial segments of \preceq . (See Figure 1.4.) This is called the *tree of 1-types*:

Definition 1.8.4 (Tree of 1-types). Let \mathbf{M} be a countable structure and let \preceq be an *enumeration* of \mathbf{M} (i.e. an order of vertices of \mathbf{M} which is either finite or has type ω). Let X be an initial segment of \preceq . Define the *1-type equivalence over X* , denoted by \sim_X , to be the equivalence on vertices of \mathbf{M} such that $u \sim_X v$ if and only if the map $f: X \cup \{u\} \rightarrow X \cup \{v\}$, such that $f(u) = v$ and $f(x) = x$ for every $x \in X$, is an isomorphism of substructures of \mathbf{M} . We call the equivalence classes of \sim_X the *1-types over X* . A *1-type* is an equivalence class over Y for some initial segment Y of \preceq . The *tree of 1-types of \mathbf{M}* is the tree whose nodes are the 1-types of \mathbf{M} and the tree order is given by being a superset. (In particular, M , being the unique 1-type over \emptyset , is the root of the tree of 1-types.)

Note that, in model-theoretic terms, these are realised quantifier-free types.

Remark 1.8.1. Trees of 1-types have been implicit in the proofs since the very beginning. The earliest explicit occurrence of these notions known to the author is in the paper of Pouzet and Sauer [PS96].

Now we can return to our original task, generalising Sierpiński's idea to colouring larger finite linear orders in (\mathbb{Q}, \leq) . Note that an embedding of a linear order into (\mathbb{Q}, \leq) is determined by its image, hence we can simply colour subsets of \mathbb{Q} of size n . We define a finite colouring χ_n of subsets of \mathbb{Q} of size n as follows: Fix an arbitrary enumeration \preceq of \mathbb{Q} . Given $S \subseteq \mathbb{Q}$ of size n , let $\bar{S} \subseteq \mathbb{Q}$ be its meet-closure in the tree of 1-types. Let $\chi_n(S)$ be the isomorphism type of the structure $(\bar{S}, \leq, \preceq, S)$, where we understand S as a unary relation (in the case when some members of S are comparable in the tree order we need to denote which of the elements of \bar{S} were the original vertices).

We have the following:

Theorem 1.8.5 (Laver (1969, unpublished, see Page 73 of [Dev79])). *For every n , χ_n is a universal colouring. Consequently, (\mathbb{Q}, \leq) has finite big Ramsey degrees,*

that is, there is a function $t: \omega \rightarrow \omega$ such that for every finite linear order \mathbf{A} and for every $k \in \omega$ we have

$$(\mathbb{Q}, \leq) \longrightarrow ((\mathbb{Q}, \leq))_{k, t(|A|)}^{\mathbf{A}}.$$

In his argument, Laver re-invented the Halpern–Läuchli theorem [HL66]. His technique was later formulated in a higher generality using Milliken’s tree theorem [Mil79] which we will present in the next section.

We conclude this section with a result of Devlin:

Theorem 1.8.6 (Devlin [Dev79]). *There exists an embedding of (\mathbb{Q}, \leq) into itself such that χ_n restricted to this embedding is persistent.*

Devlin’s embedding is quite easy to explain in terms of the tree of 1-types: The image of the embedding forms an antichain in the tree order. Curiously, the exact big Ramsey degree of the linear order on n vertices is $\tan^{(2n-1)}(0)$, where $\tan^{(k)}(x)$ denotes the k -th derivative of the tangent function.

1.8.2 Milliken’s tree theorem

Let T be a tree. For $D \subseteq T$, we write $L_T(D) = \{|t|_T : t \in D\}$ for the *level set* of D in T . We use $T(n)$ to denote the set of all nodes of T at level n , and by $T(<n)$ the set $\{t \in T : |t|_T < n\}$. Note that $T(0)$ consists of the root.

The *height* of T is the smallest natural number h such that $T(h) = \emptyset$. If there is no such number h , then we say that the height of T is ω . We denote the height of T by $h(T)$. A node $t \in T$ is *maximal* in T if it has no successors in T . The tree T is *balanced* if it either has infinite height and no maximal nodes, or all its maximal nodes are in $T(h-1)$, where h is the height of T .

Definition 1.8.5. A subtree S of a tree T is a *strong subtree* of T if either S is empty, or S is nonempty and satisfies the following three conditions.

1. The tree S is rooted and balanced.
2. Every level of S is a subset of some level of T , that is, for every $n < h(S)$ there exists $m \in \omega$ such that $S(n) \subseteq T(m)$.
3. For every non-maximal node $s \in S$ and every $t \in \text{ImmSucc}_T(s)$ the set $\text{ImmSucc}_S(s) \cap \text{Succ}_T(t)$ is a singleton.




Figure 1.5: An example of a strong subtree.

Observation 1.8.7. *If E is a subtree of a balanced tree T , then there exists a strong subtree $S \supseteq E$ of T such that $L_T(E) = L_T(S)$. \square*

For every $k \in \omega \cup \{\omega\}$ with $k \leq h(T)$, we use $\text{Str}_k(T)$ to denote the set of all strong subtrees of T of height k .

Theorem 1.8.8 (Milliken [Mil79]). *For every rooted, balanced and finitely branching tree T of infinite height, every non-negative integer k and every finite colouring of $\text{Str}_k(T)$ there is $S \in \text{Str}_\omega(T)$ such that the set $\text{Str}_k(S)$ is monochromatic.*

Note that Milliken's theorem in its true form talks about product trees and product subtrees. However, for our exposition it is sufficient to state its one-dimensional variant.

Proof of Laver's theorem

We are now ready to prove Laver's theorem. Let $T = \{0, 1\}^{<\omega}$ be the tree of finite binary sequences ordered by end-extension \sqsubseteq . Note that every strong subtree of T is isomorphic to an initial segment of T of the same height (T is binary). Given a finite set $S \subseteq T$, let \bar{S} be the minimal strong subtree of T containing S (it is an easy exercise to verify that it is well-defined and that there is a bound on the height of \bar{S} by a function of $|S|$), and let ι be the isomorphism of \bar{S} and the corresponding initial segment of T . By the *embedding type* of S we denote the set $\iota(S)$. For every $n \geq 1$, let ξ_n be the function assigning each subset of T of size n its embedding type. Observe that ξ_n is a finite colouring of subsets of T of size n . Let \leq be the lexicographic order on T .

Proposition 1.8.9. *ξ_n is a universal colouring of subsets of size n of (T, \leq) .*

Note that this implies Laver's theorem: First observe that, by universality, (T, \leq) embeds into (\mathbb{Q}, \leq) . Conversely, one can construct an embedding $\psi: (\mathbb{Q}, \leq) \rightarrow (T, \leq)$ such that there is exactly one vertex from the image of ψ on every level of T : Simply enumerate \mathbb{Q} arbitrarily and use density of (T, \leq) . This means that a colouring of subsets of (\mathbb{Q}, \leq) restricts to a colouring of subsets of (T, \leq) where we can use universality to obtain a nice copy of (T, \leq) which contains a nice copy of (\mathbb{Q}, \leq) inside. It remains to observe that if we restrict ξ_n to sets S with at most one node per level, it corresponds 1-to-1 to χ_n as defined earlier for (\mathbb{Q}, \leq) .

Proof of Proposition 1.8.9. Let c be an arbitrary finite colouring of subsets of size n of (T, \leq) and enumerate all embedding types of sets of size n in (T, \leq) as e_1, \dots, e_k . Let h_1 be the height of the strong subtree corresponding to e_1 and define a colouring of $\text{Str}_{h_1}(T)$ where we colour a strong subtree based on the c -colour of its subset corresponding to e_1 . Use Milliken's theorem to obtain a monochromatic infinite strong subtree of T (which is isomorphic to (T, \leq)) in which all subsets of size n of embedding type e_1 have the same colour. Do the same for e_2 inside of this subtree and continue by induction. \square

Let us recapitulate the scheme of the whole proof:

1. We started with a structure (\mathbb{Q}, \leq) for which we want to get universal colourings.

2. Then we identified the tree $T = \{0, 1\}^{<\omega}$ which satisfied three conditions simultaneously:

- We had a partition theorem for some notion of subtrees of T (Milliken's theorem),
- we were able to define a linear order on T which was rich enough to embed (\mathbb{Q}, \leq) such that the subtrees provided by the partition theorem preserved the order, and
- we were able to transfer every colouring of subsets of vertices of the copy of (\mathbb{Q}, \leq) of size n into boundedly many colourings of strong subtrees of T (based on the embedding types).

3. Having this, we could then use Milliken's theorem on each of the colourings of strong subtrees separately to obtain a copy where the colours only depended on the embedding type.

Note, in particular, that even though the end-goal was to prove finite big Ramsey degrees for (\mathbb{Q}, \leq) , the key ingredient was a proof of finite big Ramsey degrees for another structure. This is a recurring phenomenon, and as such, it makes sense to prove abstractly that one can transfer such results.

Observation 1.8.10. *Let \mathbf{M} and \mathbf{N} be structures with embeddings $f: \mathbf{M} \rightarrow \mathbf{N}$ and $g: \mathbf{N} \rightarrow \mathbf{M}$. If χ is a universal colouring of $\binom{\mathbf{M}}{\mathbf{A}}$ then its restriction is a universal colouring for $\binom{g(\mathbf{N})}{\mathbf{A}}$, and consequently (after applying g^{-1}) it is a universal colouring for $\binom{\mathbf{N}}{\mathbf{A}}$.*

Proof. Let $c: \binom{\mathbf{N}}{\mathbf{A}} \rightarrow k$ be an arbitrary finite colouring and let c' be its restriction to $\binom{f(\mathbf{M})}{\mathbf{A}}$, understood as a colouring of $\binom{\mathbf{M}}{\mathbf{A}}$. Since χ is universal, we get an embedding $\psi: \mathbf{M} \rightarrow \mathbf{M}$ such that c' looks like χ on $\psi(\mathbf{M})$. Hence c looks like χ on $f \circ \psi \circ g(\mathbf{N})$ \square

A similar observation holds for persistent colourings, and so we get the following corollary.

Corollary 1.8.11. *If structures \mathbf{M} and \mathbf{N} are bi-embeddable then they have the same big Ramsey degrees.*

1.8.3 Big Ramsey degrees of the Rado graph

Inspired by the arguments for (\mathbb{Q}, \leq) , we will now look at big Ramsey degrees of the countable random graph \mathbf{R} . As we have seen, it makes sense to split our work into two parts: First try to obtain a (finite) universal colouring using the tree of 1-types which would colour based on some variant of embedding types, and only after succeeding we can try to construct an embedding $\mathbf{R} \rightarrow \mathbf{R}$ which minimizes the number of embedding types. Also, similarly as for (\mathbb{Q}, \leq) , it turns out to be more convenient to colour all subsets of size n instead of substructures – if we prove a finite bound on the number of colours for colouring subsets, then all the more so we have one for colouring substructures.

As a warm-up, let us prove that the Rado graph is indivisible (this was already known to Erdős and other Hungarian mathematicians in the 1960's and the author is unaware of any particular good reference for this).

Observation 1.8.12. *The Rado graph is indivisible. That is, $\mathbf{R} \rightarrow (\mathbf{R})_2^{\mathbf{A}}$ for \mathbf{A} being the 1-vertex graph.*

Proof. Suppose that we have a red-blue colouring of vertices of \mathbf{R} . Pick an arbitrary enumeration \preceq of \mathbf{R} and start constructing an embedding $\mathbf{R} \rightarrow \mathbf{R}$ one-by-one in the red colour. That is, in each step we try to find a red image of the \preceq -first vertex which does not yet have an image. If such a red image exists, we pick an arbitrary one and continue. So suppose that, at some point, all possible images are coloured blue. This means that there are vertices $X = \{x_0, \dots, x_k\} \subseteq R$ (the already constructed images) and a 1-type over X whose all realisations are coloured blue. However, it is an easy fact about \mathbf{R} that it induces a copy of \mathbf{R} on the set of all realisations of any type over a finite set, hence we get a blue copy of \mathbf{R} . \square

Thus, pick an arbitrary enumeration \preceq of \mathbf{R} and look at the tree of 1-types. First we have just one type containing all vertices of \mathbf{R} . Then, once we discover the \preceq -first vertex v_0 , the type splits into two, depending on whether they have an edge to v_0 or not. If v_1 is the \preceq -second vertex, each type over v_0 splits into two upon discovering v_1 based on the connection to v_1 , and so on. Thus, the tree of 1-types of \mathbf{R} is isomorphic to $T = \{0, 1\}^{<\omega}$, where a binary string $w = (w_0, w_1, \dots, w_k)$ describes the type over $\{v_0, \dots, v_k\}$ which is connected precisely to those v_i 's for which $w_i = 1$.

Similarly as for \mathbb{Q} , one can use this to devise potentially bad colourings for \mathbf{R} . For example, given three vertices of \mathbf{R} , we can look at the meet of two of them (that is, the \preceq -least vertex to which they are not connected in the same way) and colour the triple based on, for example, whether the third vertex is connected to the meet or not.

The isomorphism between the tree of 1-types and T suggests that we can define an edge relation E on T putting $xy \in E$ if and only if $|x| \neq |y|$ and $y_{|x|} = 1$ (without loss of generality we can assume that $|x| < |y|$). The entry $y_{|x|} = 1$ is called the *passing number* of y at level $|x|$ and the whole graph relation is often called the *passing number representation* [Sau06]. Notice that this is an infinitary asymmetric variant of the simple construction of EPPA-witnesses for graphs using valuation functions from Section 1.6.1.

Since (T, E) is a countable graph, it embeds into \mathbf{R} . On the other hand, \mathbf{R} embeds into (T, E) – given an enumeration \preceq of \mathbf{R} , we can send the 0-th vertex of \mathbf{R} to the empty sequence, the 1-st vertex of \mathbf{R} to either (0) or (1) according to whether it has an edge to the 0-th vertex and so on.

Now we would like to construct a finite universal colouring of subsets of T of size n for every n . Let S be an arbitrary subset of T of size n and let \bar{S} be its meet-closure. Note that $|\bar{S}| \leq 2|S| - 1$, hence it has a bounded number of levels. By Observation 1.8.7 there is a strong subtree of T containing \bar{S} with the same level set. Since T is binary, any strong subtree of T is isomorphic to an initial segment of T , hence we can define the *embedding type* of S to be the image of S in the initial segment of T under the isomorphism. Note that the embedding

type does not depend on the choice of the strong subtree containing \bar{S} . Let χ_n be the colouring assigning each subset of T of size n its embedding type.

Theorem 1.8.13 (Sauer [Sau06]). *χ_n is universal. Consequently, the Rado graph has finite big Ramsey degrees.*

Proof. Similarly as in the proof of universality for (\mathbb{Q}, \leq) , enumerate all embedding types of sets of size n in (T, E) as e_1, \dots, e_k . Let c be an arbitrary finite colouring of subsets of size n of (T, E) . Then, one-by-one for every e_i , colour the strong subtrees of the corresponding height according to the c -colour of their subset given by e_i and use Milliken's theorem to stabilise this colouring. Note that strong subtrees respect passing numbers and hence they preserve the graph structure on (T, E) . (Formally, a strong subtree of (T, E) of height $k \in \omega \cup \{\omega\}$ is isomorphic, as a graph, to the initial segment of (T, E) of height k .)

The fact that the Rado graph has finite big Ramsey degrees now follows from Observation 1.8.10. \square

The exact big Ramsey degrees of the Rado graph have been characterised by Laflamme, Sauer, and Vuksanovic [LSV06]. While the proof itself is rather technical, the image of the optimal embedding $\mathbf{R} \rightarrow (T, E)$, intuitively, forms an antichain in the tree, on every level there is either exactly one meet of two vertices from the image, or one vertex from the image and no meets, and meets are *canonical* (that is, all types which do not split at that level have passing number 0).

Theorem 1.8.14 (Laflamme, Sauer, and Vuksanovic [LSV06]). *There is an explicitly described embedding $\mathbf{R} \rightarrow \mathbf{R}$ on which χ_n is both universal and persistent.*

1.8.4 Parameter spaces and the Carlson–Simpson theorem

Before presenting the more complicated examples of the triangle-free graph and the 3-uniform hypergraph, we will see another proof of Sauer's theorem using parameter spaces and the Carlson–Simpson theorem instead of Milliken's theorem. First, we need to introduce some terminology. (The main reference for this short section is [Hub20a] which first used parameter spaces in the big Ramsey context, see also Chapter 11 where we apply the same ideas to prove finiteness of big Ramsey degrees for a broad family of binary structures.)

Given a finite alphabet Σ and $k \in \omega \cup \{\omega\}$, a *k-parameter word* is a (possibly infinite) string W in the alphabet $\Sigma \cup \{\lambda_i : 0 \leq i < k\}$ containing all symbols $\lambda_i : 0 \leq i < k$ such that, for every $1 \leq i < k$, the first occurrence of λ_i appears after the first occurrence of λ_{i-1} . The symbols λ_i are called *parameters*. Given a parameter word W , we denote its *length* by $|W|$. The letter (or parameter) on index i with $0 \leq i < |W|$ is denoted by W_i . Note that the first letter of W has index 0. A 0-parameter word is simply a *word*.

Let W be an n -parameter word and let U be a parameter word of length $k \leq n$, where $k, n \in \omega \cup \{\omega\}$. Then $W(U)$ is the parameter word created by *substituting* U to W . More precisely, $W(U)$ is created from W by replacing each occurrence of λ_i , $0 \leq i < k$, by U_i and truncating it just before the first occurrence of λ_k in

W . Given an n -parameter word W and a set S of parameter words of length at most n , we define $W(S) := \{W(U) : U \in S\}$.

We let $[\Sigma] \binom{n}{k}$ be the set of all k -parameter words of length n , where $k \leq n \in \omega \cup \{\omega\}$. If k is finite, then we also define $[\Sigma]^* \binom{\omega}{k} := \bigcup_{k \leq i < \omega} [\Sigma] \binom{i}{k}$. For brevity, we put $\Sigma^* := [\Sigma]^* \binom{\omega}{0}$.

Our main tool will be the following infinitary dual Ramsey theorem, which is a special case of the Carlson–Simpson theorem [CS84, Tod10].

Theorem 1.8.15. *Let Σ be a finite alphabet and $k \geq 0$ a finite integer. If the set $[\Sigma]^* \binom{\omega}{k}$ is coloured by finitely many colours, then there exists an infinite-parameter word W such that $W \left([\Sigma]^* \binom{\omega}{k}\right)$ is monochromatic.*

Note that if $\Sigma = \{0, 1\}$ then $\Sigma^* = \{0, 1\}^{<\omega}$ is just the infinite binary tree. Also note that given $W \in [\Sigma]^* \binom{\omega}{k}$, the set $W(\Sigma^*)$ is a special strong subtree of Σ^* of height k . However, not every strong subtree can be obtained in this way: For example, the strong subtree of height 1 consisting of the words $\{\emptyset, 00, 11\}$ is described by the 1-parameter word $\lambda_0\lambda_0$, but the strong subtree of height 1 consisting of the words $\{\emptyset, 01, 10\}$ cannot be described by any parameter word. This property is not relevant in the proof of finite big Ramsey degrees for the Rado graph, but it will be crucial in Section 1.8.5 where it will allow us to prove finite big Ramsey degrees for the triangle-free graph.

Our proof of big Ramsey degrees of the Rado graph using parameter spaces will go in a full analogy to the proof using Milliken’s theorem: We will use the same tree $T = \{0, 1\}^{<\omega}$ and the same edge relation E on it. However, instead of enveloping subsets to strong subtrees, we will envelope them to parameter words which we will then colour. This motivates the following definition and fact.

Definition 1.8.6 ([Hub20a]). Given a finite alphabet Σ , a finite set $S \subseteq \Sigma^*$ and $d > 0$, we call $W \in [\emptyset]^* \binom{\omega}{d}$ a *d -parameter envelope* of S if there exists a set $S' \subseteq \Sigma^*$ satisfying $W(S') = S$. In such case the set S' is called the *embedding type* of S in W and is denoted by $\tau_W(S)$. If d is the minimal integer for which a d -parameter envelope W of S exists, then we call W a *minimal envelope*.

Fact 1.8.16 ([Hub20a]). *Let Σ be a finite alphabet and let $k \geq 0$ be a finite integer. Then there exists a finite $T = T(|\Sigma|, k)$ such that every set $S \subseteq \Sigma^*$, $|S| = k$, has a d -parameter envelope with $d \leq T$. Consequently, there are only finitely many embedding types of sets of size k within their corresponding minimal envelopes. Finally, for any two minimal envelopes W, W' of S , we have $\tau_W(S) = \tau_{W'}(S)$. We will thus also use $\tau(S)$ to denote the type $\tau_W(S)$ for some minimal W .*

The proof of this fact is constructive – one defines a canonical envelope, observes that it is minimal and that all other minimal envelopes differ from this one in a controlled way which does not affect embedding types.

Theorem 1.8.17. *The colouring ξ_n assigning each subsets $S \subseteq \Sigma^*$ of size n its embedding type $\tau(S)$ is universal. Consequently, the Rado graph has finite big Ramsey degrees.*

Proof. In a complete analogy to the proof of Theorem 1.8.13, enumerate all embedding types of sets of size n in $T = \Sigma^*$ as e_1, \dots, e_k . Let c be an arbitrary finite colouring of subsets of size n of (T, E) . Then, one-by-one for every e_i , colour the parameter words of the corresponding dimension according to the colour of their subset given by e_i and use Theorem 1.8.15 to stabilise this colouring. Note that substitution respects passing numbers and hence it preserves the graph structure on (T, E) . \square

1.8.5 The triangle-free graph and age changes

The triangle-free graph is the Fraïssé limit of the class of all finite triangle-free graphs, that is, the triangle-free analogue of the Rado graph. Indivisibility of the triangle-free graph was shown by Komjáth and Rödl in 1986 [KR86] and big Ramsey degrees for colouring edges were studied by Sauer in 1998 [Sau98]. However, it has long eluded attempts to prove finiteness of big Ramsey degrees which were finally shown to be finite by Dobrinen’s breakthrough around 2017 [Dob20a]. Her paper is long and complicated and uses a custom tree partition theorem for coding trees which was proved using the method of forcing. By now there is a simple proof using parameter spaces by Hubička [Hub20a] which we will sketch in this section.

But first, let us attempt to prove indivisibility of the triangle-free graph analogously as we did for the Random graph. Let \mathbf{H}_3 be the countable homogeneous triangle-free graph (\mathbf{H}_3 for *Henson graph*) and let \preceq be an enumeration of its vertices. Assume that the vertices of \mathbf{H}_3 are coloured red or blue and try to iteratively construct an embedding $\mathbf{H}_3 \rightarrow \mathbf{H}_3$ into the red colour. That is, given the \preceq -first not-yet-embedded vertex, look at the 1-type over the already embedded vertices which this new vertex should be in, and if there is a red vertex in the 1-type then we have found a new image. Otherwise all vertices from the 1-type are blue. In the Rado case, we have thus found a copy of the Rado graph. However, this need not be the case here: If the vertex we have been trying to embed has an edge to some previous vertex then all vertices of our blue 1-type have a common neighbour which implies that there are no edges between them (indeed, otherwise we would find a triangle).

This is a new phenomenon: Both in (\mathbb{Q}, \leq) and \mathbf{R} , every non-trivial 1-type realises the whole infinite structure while in \mathbf{H}_3 we have two kinds of 1-types: A 1-type over X which has no edges to X contains a copy of \mathbf{H}_3 while a 1-type over X with at least one edge to X contains no edges. Moreover, this phenomenon manifests itself throughout the big Ramsey theory of the triangle-free graph, not just as a nuisance in an indivisibility proof. Given an enumeration \preceq of \mathbf{H}_3 and a vertex x , we can address the first neighbour of x (if it exists) and use it in our colouring. Similarly, given two vertices $x \neq y$, we can address their first common neighbour (if it exists) and use it in our colouring.

The obvious question then is: “Why can we do it now and not in the Rado graph case?” And the difference is rather subtle. We, of course, can devise such colourings for the Rado graph, but it is possible to find copies of the Rado graph where all vertices of the copy are connected to some “external” vertex which precedes all of them, therefore stabilising this kind of colouring. For \mathbf{H}_3 , this is not possible and first neighbours and first common neighbours always have a rich

structure.

This phenomenon has been formalised in the notion of *age changes*. The ideas appear implicitly in the work of Sauer and co-authors (e.g. [EZS89, Sau98]), Dobrinen [Dob23] is the first person known to the author who considered age changes on tuples of types and Zucker [Zuc22] has formalised the concept in full generality, see also [BCD⁺21b].

Definition 1.8.7. Let \mathbf{M} be a countable structure with an enumeration \preceq and let t be a 1-type over X (recall that a 1-type is a subset of M). The *age* of t is the class of all finite structures realised by members from t . The collection of all ages of all 1-types of \mathbf{M} ordered by inclusion is the *age poset*.

For example, the age poset of the Rado graph has just one element, the class of all graphs. (Or rather this depends on the precise definition of a type, one could argue that it also contains the class consisting of one 1-vertex structure corresponding to the 1-types of vertices from X , and in some cases, potentially, also the empty age corresponding to unrealised 1-types.) Similarly, the age poset of (\mathbb{Q}, \leq) has just one element. The age poset of \mathbf{H}_3 has two elements – the class of all finite triangle-free graphs and the class of all finite edge-free graphs. In general, for \mathbf{H}_k being the generic K_k -free Henson graph, its age poset has $k - 1$ members, the classes of all finite K_i -free graphs for $2 \leq i \leq k$.

In all the above cases, the age poset was a chain. Consider now the class of all finite graphs with red and blue edges (that is, there are three types of pairs – red edge, blue edge, and no edge) containing no monochromatic triangles. Then the age poset consists of this class, the class of all finite triangle-free graphs with red edges (and no blue edges), the class of all finite triangle-free graphs with blue edges (and no red edges), and the class of all edge-free graphs. In particular, the age poset is not a chain.

We can now prove finiteness of big Ramsey degrees for the triangle-free graph. Let us first try to use Milliken’s theorem for our proof and see what happens. Put $T = \{0, 1\}^{<\omega}$. Next we will define an edge relation E on T as follows: We put $xy \in E$ if and only if the following conditions are satisfied:

1. $|x| \neq |y|$ (assume without loss of generality that $|x| < |y|$),
2. $y_{|x|} = 1$, and
3. there is no $0 \leq i < |x|$ such that $x_i = y_i = 1$.

Note that the third condition is an extra condition compared to the definition of the edge relation for the Rado graph. It says that if two vertices already have a common neighbour, we remove any potential edges between them. This makes the graph triangle-free. At the same time, by an analogous argument as before, (T, E) embeds \mathbf{H}_3 , hence we can try to find a universal colouring for (T, E) .

Assume for now that we were able to define embedding types and that we only have finitely many of them for each n . Then it remains to use Milliken’s theorem several times to stabilise the colouring on each embedding type. This means that we have a colouring of strong subtrees of finite height and Milliken’s theorem gives us an infinite strong subtree of T on which the colours are stabilised. However, for our proof we actually need this strong subtree to give us an isomorphic copy

of (T, E) inside (T, E) . Let r be the root of this infinite strong subtree. We know that r is a binary sequence and that all elements of the subtree extend r . In particular, if r is not the all-zero sequence then all elements of the subtree have a common neighbour somewhere below r , and consequently there are no edges between vertices of our subtree. This is the fundamental issue with using Milliken's theorem in the presence of age changes – it is oblivious to them and hence will typically give us subtrees with minimal ages in which we cannot embed our structures.

Dobrinen's original solution was to consider *coding trees* which are trees with special denoted coding nodes and their embeddings need to respect these coding nodes. Hubička's observation in [Hub20a] was that using a simple trick, parameter spaces can be made aware of age changes and parallel 1's. We will now sketch this proof, for a full (still very short) version see [Hub20a].

Let $\Sigma = \{0\}$ and put $T = [\Sigma]^* \binom{\omega}{1}$ to be the space of all finite 1-parameter words. For simplicity we will call the parameter λ . Note that T is almost the same as $\{0, 1\}^{<\omega}$ after renaming $\lambda \mapsto 1$ with the important exception that every member of T contains at least one λ . As such, T is not actually a tree but a disjoint union of infinitely many trees. (Equivalently, we can see T as $\{0, 1\}^{<\omega}$ after removing the constant-zero words.) This trick is not crucial, one can also use parameter spaces to prove the result on $\{0, 1\}^{<\omega}$, but it makes the arguments shorter. We define an edge relation E on T in the same way as above, putting $xy \in E$ if and only if the following conditions are satisfied:

1. $|x| \neq |y|$ (assume without loss of generality that $|x| < |y|$),
2. $y_{|x|} = \lambda$, and
3. there is no $0 \leq i < |x|$ such that $x_i = y_i = \lambda$.

Given $n \geq 1$, define a colouring χ_n such that, given $S \subseteq T$ of size n , we put $\chi_n(S) = \tau(S)$ (see Fact 1.8.16, an analogue of it holds also for $[\Sigma]^* \binom{\omega}{1}$ instead of Σ^*).

Theorem 1.8.18 (Hubička [Hub20a]). *χ_n is universal for every n . Consequently, \mathbf{H}_3 has finite big Ramsey degrees.*

Proof. As always we will enumerate all embedding types and use Theorem 1.8.15 for each of them separately. The only thing left to verify is that if W is an infinite-parameter word then (T, E) induces an isomorphic copy of itself on $W(T)$. As we have observed for the Rado graph, substitution preserves the order between lengths of the words as well as passing numbers. It only remains to see that, for $x, y \in T$ with $|x| < |y|$, there exists $0 \leq i < |x|$ such that $x_i = y_i = \lambda$ if and only if there exists $0 \leq j < |W(x)|$ such that $W(x)_j = W(y)_j = \lambda$. However, W is an infinite-parameter word in the alphabet $\{0\}$, hence $W(x)_j = \lambda$ if and only if $W_j = \lambda_k$ for some k such that $x_k = \lambda$. This means that $W(x)_j = W(y)_j = \lambda$ if and only if $W_j = \lambda_k$ such that $x_k = y_k = \lambda$. This concludes the proof. \square

1.8.6 Diaries

So far we have only proved finite upper bounds on big Ramsey degrees of \mathbf{H}_3 , or in other words, constructed finite universal colourings. The next natural question is to characterise the exact big Ramsey degrees, or in other words, construct

colourings which are universal and persistent. As we have seen with the Rado graph, this is usually done by starting with the universal colouring and constructing a copy of the countable structure which realises as few embedding types as possible. This, in turn, is done by understanding what the *interesting events* are (in the Rado case, they were *splitting* (also called *branching*) – that is, a meet of some pair of 1-types – and *coding* – an image of a “real” vertex). After understanding the interesting events, one can try to construct a copy which only has one interesting event on every level and everything happens as canonically as possible.

For the triangle-free graph, there is a new kind of interesting events, *age changes*. In fact, one can see it as two kinds of interesting events – an age change on a single type, corresponding to seeing a first neighbour of some type, and an age change on a pair of types, corresponding to seeing their first common neighbour. It is almost possible to construct a copy of the triangle-free graph where only one interesting event happens on every level, with one exception – if a 1-type wants to branch from the *generic 1-type* (that is, the type over X which has no edges to X), it needs to do it by introducing a vertex to which it is connected but the generic type is not, hence at the same time creating the first neighbour event. This was done by Balko, Chodounský, Dobrinen, Hubička, Vena, Zucker, and the author in [BCD⁺21b] where exact big Ramsey degrees are characterised for all free amalgamation classes in binary languages given by a finite set of forbidden structures. While both the results and the proofs are rather technical, the intuition is that there are the following three kinds of interesting events:

- Coding (that is, a new vertex of the nice copy appears),
- splitting (that is, a 1-type of the nice copy splits into two 1-types), and
- age changes.

Age changes are more complicated, they can happen on a single 1-type (such as the first neighbour in the triangle-free graph), but also on tuples of several 1-types (such as the first common neighbour). The conditions are that only one thing happens on every level except when it is not possible (cf. splitting and age change in the triangle-free graph), everything happens as canonically as possible (that is, if two events do the same then they are forced by the same gadgets, types always split into two, and upon coding, the vertex already has the minimal age which it achieved using a maximal path in the age poset). And there is one more thing to keep track of – if the age poset is not a chain then there are multiple maximal paths a coding node can choose from. In the final structure, there need to appear all of them and one can colour by the path the node took. Hence, “one needs to decide the path for each coding node at the very beginning”.

A useful intuition is the following (which the author learned from an unpublished draft by Hubička): Assume that Robinson Crusoe was cast away on a deserted island. To kill time, he watches trees grow. He is most fascinated by a tree \mathbf{R} to such a degree that, each time the tree grows, Robinson documents all of it in his *diary*.

At the beginning, the tree was just a small seed t_\emptyset . But at some point, a \preceq -first vertex has grown. Thereupon has Robinson written in his diary the following two lines:

“A new coding node v_0 grew from the type t_\emptyset .”

“The type t_\emptyset has split into two types, t_0 and t_1 such that t_1 is connected to v_0 .”

As you can imagine, when a new coding node grows in one of the types, Robinson again records it and splits t_0 into t_{00} and t_{01} and, similarly, splits t_1 into t_{10} and t_{11} and so on. If Robinson happened to be fascinated by tree \mathbf{H}_3 instead, he would be also recording all the age changes.

Note that if, at the end, we look at a subset of vertices, we can look at all the entries in the diary which only talk about subsets of these vertices (or their types). This gives us a *subdiary*. However, the events in this subdiary may have been caused by *external* vertices which are not in this subset. (Formally, in our proofs, the vertices are sequences, and when taking subsets, we keep the whole sequences instead of removing the levels from which no vertex has been selected.) And it is the existence of the external vertices that allows us to disintegrate the bundles of events into atomic pieces whose particular order then determines the embedding types.

1.8.7 3-uniform hypergraphs

On our voyage through simple examples of not-so-simple big Ramsey degrees, let us venture for the first time out of the realm of structures in a binary language and let us consider the generic 3-uniform hypergraph \mathbf{H} , that is, the Fraïssé limit of the class of all finite 3-uniform hypergraphs. As there are no forbidden cliques, the age poset is trivial and we can use the standard argument to obtain indivisibility: Try to embed \mathbf{H} vertex-by-vertex into the red colour, and if at some point we fail, we find a copy in the blue colour.

Next, let us look at the tree of 1-types. Fix an enumeration \preceq of \mathbf{H} such that $H = \{v_0 \preceq v_1 \preceq \dots\}$. We start with the type consisting of all vertices of \mathbf{H} . Upon discovering v_0 , nothing actually changes. Only upon discovering v_1 does the type split based on whether there is a hyperedge together with v_0v_1 . However, there are already 8 types over $\{v_0, v_1, v_2\}$ based on the (non-)existence of hyperedges with v_0v_1 , v_0v_2 and v_1v_2 . Upon discovering v_3 , each of these 8 types splits into 8 new types based on the (non-)existence of hyperedges with v_0v_3 , v_1v_3 and v_2v_3 . Then, each of these 64 types splits into 16 types and so on.

Similarly as the tree of 1-types for the Rado graph can be identified with the tree $\{0, 1\}^{<\omega}$ of all finite binary sequences, this tree can be identified with the tree of all finite lower-triangular binary square matrices: Let $M \in \{0, 1\}^{m \times m}$ be a matrix. We say that it is *lower-triangular* if $M_{i,j} = 0$ whenever $i \leq j$, and we denote $|M| = m$. Given two matrices $M \in \{0, 1\}^{m \times m}$, $N \in \{0, 1\}^{n \times n}$ with $m \leq n$, we put $M \sqsubseteq N$ if and only if for every $0 \leq i, j \leq m$ it holds that $M_{i,j} = N_{i,j}$. Then the set T of all finite lower-triangular binary square matrices together with \sqsubseteq forms a tree and it can be identified with the tree of 1-types of \mathbf{H} – given a node M , the entry $M_{i,j} = 1$ if and only if the type is in a hyperedge together with v_i and v_j with $i > j$. See Figure 1.6.

Analogously to the Rado graph case, we can define a hyperedge relation E on T , putting $XYZ \in E$ if and only if they have different dimensions and (without




Figure 1.6: The tree of lower-triangular binary square matrices, or equivalently, the tree of 1-types of \mathbf{H} .

loss of generality, $|X| < |Y| < |Z|$) $Z_{|Y|,|X|} = 1$. Again, we call this entry the *passing number* of Z at $|Y|, |X|$.

Note that the tree is very much not regularly branching, and so it does not seem like we can use parameter spaces to get a Ramsey theorem for it. On the other hand, Milliken’s theorem can handle such rapidly branching trees with no issues. Nevertheless, we run into another issue if we want to apply the standard arguments with Milliken’s theorem:

Look at, for example, strong subtrees of depth 2. They have two levels: One contains the root, and for every immediate successor of the root, the second level contains exactly one successor of the immediate successor. In particular, the higher the root is in the original tree, the more vertices will there be on the second level. This suggests that there will be issue proving a bound on the number of embedding types. On the other hand, Milliken’s theorem gives us an infinite strong subtree, but we really only want an isomorphic copy of the original structure. Since all trees of 1-types we have seen so far were regularly branching, strong subtrees were isomorphic to the original subtree, but this is no longer the case here – the higher we start in the tree, the more the strong subtrees branch compared to the original tree. So on one hand, wide trees lead to strong subtrees containing too many embedding types, on the other hand the theorem gives us much stronger monochromatic objects, so perhaps one could somehow leverage these two facts.

This turns out to be indeed possible but the best way to motivate it is by taking a detour. Remember that, before we even defined trees of 1-types, we observed that given two vertices of (\mathbb{Q}, \leq) (or \mathbf{R}), one can address the first vertex which lies between them (or to which they are not connected in the same way), their *meet*, and use it to define a richer colouring. In \mathbf{H} , one can do more than this: Given an enumeration \preceq and two pairs of vertices $x \prec y$ and $x' \prec y'$, we can address the first vertex m such that exactly one of mxy and $mx'y'$ is a hyperedge (sometimes such a vertex need not exist). This motivates the following definition:

Definition 1.8.8 (Tree of k -types). Let \mathbf{M} be a countable structure, let \preceq be an enumeration of \mathbf{M} and let $k \geq 1$ be an integer. Let X be an initial segment of \preceq .

Define the k -type equivalence over $X \sim_X^k$ to be the equivalence on decreasing k -tuples of vertices of \mathbf{M} such that $\bar{u} \sim_X \bar{v}$ if and only if the map $f: X \cup \bar{u} \rightarrow X \cup \bar{v}$, such that $f(u_i) = v_i$ for every $1 \leq i \leq k$ and $f(x) = x$ for every $x \in X$, is an isomorphism of substructures of \mathbf{M} . We call the equivalence classes of \sim_X^k the k -types over X . A k -type is an equivalence class over Y for some initial segment Y of \preceq . The tree of k -types of \mathbf{M} is the tree whose nodes are the k -types of \mathbf{M} and the tree order is given by the inclusion.

Note that, in model-theoretic terms, these are again the quantifier-free types. k -types for $k \geq 2$ have first been considered in the area by Balko, Chodounský, Hubička, Vena, and the author in [BCH⁺19], an announcement of the result which we present in Chapter 8 in which finiteness of big Ramsey degrees for the 3-uniform hypergraph is proved and only the tree of 1-types and the tree of 2-types are needed. In Chapter 9, these ideas are further extended to unrestricted relational structures with arbitrary arities.

Observe that for \mathbf{H} , the tree of 2-types is the infinite binary tree. This corresponds to the fact that if we fix one vertex of \mathbf{H} , we can define the *neighbourhood graph* on the remaining vertices according to hyperedges containing the fixed vertex. The tree of 2-types describes all the neighbourhood graphs together.




Figure 1.7: The trees of 1-types and 2-types of \mathbf{H}

Given a 1-type of \mathbf{M} and a level from below the type, we get a 2-type which we call the *passing type*. In terms of matrices, given $X \in T$ and a level n with $n < |X|$, the passing type of X at n is simply the row $X_{n,*}$ truncated to have length $n - 1$. In particular, it is a binary word. And vice versa, a 1-type X of level n and a 2-type w of level n together give a 1-type $X \hat{\ } w$ of level $n + 1$ which is an immediate successor of X – simply extend the matrix X by a new row which begins by w and is then padded by zeros:

$$\left(\begin{array}{c|c} X & \begin{array}{c} 0 \\ \vdots \\ 0 \end{array} \\ \hline \bar{w} & 0 \end{array} \right)$$

To recapitulate, we have the tree T of 1-types and a hyperedge relation on T . Moreover, we also have the tree T_2 of 2-types of \mathbf{H} . When stating Milliken’s theorem, we noted that the real Milliken’s theorem talks about product of trees. In our case, we have trees (T, T_2) . A strong product subtree is a pair T', T'_2 consisting of a strong subtree of T and a strong subtree of T_2 with the same level set. Given such a pair, we can use the observation from the previous paragraph

to define a (non-strong) subtree of T by pruning T' according to the passing types given by T'_2 to obtain what is called the *valuation tree* (see Figure 1.8). It turns out that valuation trees are isomorphic to initial segments of T and that they are the correct subtrees on which one can get both a Ramsey theorem (the product version of Milliken’s theorem) and a bounded number of embedding types, hence enabling us to prove the following theorem.

Theorem 1.8.19 (Balko, Chodounský, Hubička, Konečný, Vena [BCH⁺22]). *The big Ramsey degrees of the 3-uniform hypergraph are finite.*

The construction in detail is given in Chapter 8.




Figure 1.8: An example of a valuation tree

Let us remark that the exact big Ramsey degrees of the 3-uniform hypergraph have not yet been characterised, although there is such a characterisation in preparation. The interesting events are coding, splitting in the 1-type tree and splitting in the 2-type tree. Besides the standard conditions, there are extra conditions that all relevant splittings in the 2-type tree need to happen before splitting in the 1-type tree, and then an analogue of the condition on maximal paths through the age poset for Devlin types in the 2-type tree.

Let us also remark that in Chapter 9 these ideas are further strengthened to arbitrary unrestricted relational structures with finitely many relations of every arity.

1.8.8 A cookbook?

Compared to EPPA and Ramsey classes, big Ramsey degrees are still in the early stages of development, and so it would be overly conceited to try to give a cookbook algorithm for obtaining finite big Ramsey degrees. However, some properties seem to hold in all the known proofs and it is reasonable to expect that they will hold in a reasonable generality:

- Enumerating a countable structure gives rise to the tree of types (in some simple cases such as Ramsey’s theorem, this tree degenerates to a chain), and since one can always try to colour according to an enumeration, it is reasonable to expect that big Ramsey results will in general talk about trees. (The author is unaware of an analogue of the small Ramsey result that every Ramsey structure fixes a linear order.)

- It seems that if we want to study the big Ramsey degree for colouring one particular finite structure \mathbf{A} , we need to consider at least trees of k -types for k up to $\min(|A|, a - 1)$, where a is the maximum arity of a relation in the language (see Chapter 9).
- The quest of understanding big Ramsey degrees naturally splits into proving their finiteness and then finding a nice embedding on which the universal colouring attains the minimum number of colours in order to prove the exact bounds. It seems that having a stronger version of the upper bounds can be very helpful in proving the lower bounds (the situation is similar in the small Ramsey world – the expansion property is often proved by invoking the Ramsey property for colouring some special small structure); see Chapter 10 for an example of this. Thus it seems to make sense to first prove upper bounds before attempting to get lower bounds.
- At the same time, to contradict the previous point, trying to construct bad colourings turned out to be helpful in constructing universal colourings several times in this section.
- In order to obtain upper bounds, it is desirable to abstract the tree Ramsey business from the structural arguments (similarly as achieved by the general theorems for Ramsey and EPPA). At this point, it is not clear how far we are from such abstract theorems. In Chapter 12 we hint at a promising concept, but there are free amalgamation classes which cannot be handled by it, and it is open whether they have finite big Ramsey degrees or not. If not then there is hope to strengthen Nešetřil’s observation that Ramsey classes have the amalgamation property to some form of observation that big Ramsey structures have some form of type amalgamation property. If they do have finite big Ramsey degrees then we need to find stronger tree theorems.
- However, in the cases where we are lucky to have a tree theorem which works, we need to figure out how to define structures on the trees so that tree morphisms give rise to structural embeddings and, at the same time, there are only finitely many embedding types. This way we can aim to obtain universal colourings. To assist with figuring this out, one should look for interesting events and, so far, we have only seen three kinds of them: coding, splitting, and age changes.

1.8.9 Examples

- (ω, \leq) has all big Ramsey degrees equal to 1 by the infinite Ramsey theorem. Consequently, the countable set with no structure has finite big Ramsey degrees (the big Ramsey degree of a set of size n is $n!$).
- The big Ramsey degree of the linear order on n vertices in (\mathbb{Q}, \leq) is equal to $\tan^{(2n-1)}(0)$ [Dev79]. A finite upper bound follows easily from Milliken’s theorem.

- The Rado graph and its variant with finitely many edge colours have finite big Ramsey degrees, again by Milliken’s theorem. [Sau06] The exact bounds are also known. [LSV06]
- The generic K_n -free graphs all have finite big Ramsey degrees by a breakthrough result of Dobrinen. [Dob20a, Dob23]
- In general, free amalgamation classes in finite binary languages determined by finitely many forbidden structures have finite big Ramsey degrees by Zucker [Zuc22] and their exact values are also known [BCD⁺21b].
- The exact big Ramsey degrees of the generic poset are also known [Hub20a, BCD⁺23a], see Chapter 10.
- Finiteness of big Ramsey degrees for various metric spaces with finitely many distances as well as generalised metric spaces is known. In general, in Chapter 11 we prove finiteness of big Ramsey degrees for structures in binary languages described by forbidden induced cycles. Metric spaces, metrically homogeneous graphs with no Henson constraints, ultrametrics, Λ -ultrametric spaces and others fall into this category. Note that Sauer with co-authors has studied indivisibility of S -Urysohn spaces even for infinite sets S [DLPS08, Sau13b] and that this is connected to the distortion problem and oscillation stability (see Section 1.8.13).

We remark that metric spaces of infinite diameter have, generally, infinite big Ramsey degree even for colouring vertices. For example, consider the integer Urysohn space, fix an arbitrary vertex c and partition the vertices according to their distance from c . Observe that every copy of the integer Urysohn space intersects every partition except for finitely many: Clearly, since it has infinite diameter, it needs to contain vertices arbitrarily far from c . Suppose that it contains vertices x, y such that $d(c, x) < d(c, y)$. By universality, the copy also needs to contain the geodesic path from x to y with $d(x, y)$ edges of length 1. Consequently, it needs to contain a vertex of each distance from c between $d(c, x)$ and $d(c, y)$. This means that, for every $k \geq 1$, the colouring of vertices by their distance from c modulo k is persistent.

- 3-uniform hypergraphs and, in general, unrestricted relational structures in languages with finitely many relations of every arity have finite big Ramsey degrees, see Chapters 8 and 9.

Some tree theorems (such as Theorem 1.8.15) are self-productive which means that one gets finiteness of big Ramsey degrees of free superpositions if they can prove finiteness of big Ramsey degrees of each of the constituents by the theorem. This gives, for example, finite big Ramsey degrees for the densely order Rado graph etc.

There are simple constructions for adding certain unary functions or unary relations to a big Ramsey degree theorem, very similar to [EHN21] or the construction from Section 4.6, see [HN19].

Despite all this progress, the area has, from a certain point of view, not yet reached the state of small Ramsey classes from late 1970’s: The Nešetřil–Rödl

theorem (Theorem 1.7.14) states that free amalgamation classes are Ramsey when expanded by linear orders (and, consequently, have finite small Ramsey degrees). While one could argue that having infinite languages or infinitely many forbidden substructures can cause difficulties in the big Ramsey context which were not visible in the small Ramsey context, there are still very simple free amalgamation classes in a finite language with finitely many forbidden substructures for which finiteness of big Ramsey degrees is open (see for example Problem 12.1.3). This is connected to the lack of the type-respecting amalgamation property, which is one of the conditions for a general theorem in preparation to work (see Chapter 12). At this point we do not know whether this is a weakness of the general theorem, a surprising feature of the problem, or some easy oversight on our side. Resolving this problem (and, likely at the same time, related ones) will be a key step for further advancing the area.

1.8.10 Big Ramsey structures and topological dynamics

Similarly as EPPA and Ramsey classes reflect in the properties of the automorphism group of the corresponding Fraïssé limit (see Sections 1.6 and 1.7.2), finiteness of big Ramsey degrees also has its dynamical counterparts. However, once again, this area is still in its beginnings and it is presently being rapidly developed.

The first difference is that big Ramsey degrees are more a property of the self-embedding monoid than of the automorphism group. However, the paper of Zucker from 2017 [Zuc19] defines the *completion flows* and proves that having a *big Ramsey structure* gives rise to a *universal completion flow*. We refer the reader to Zucker’s paper for details (and, in the near future, to its follow-up with more refined notions motivated by the recent progresses), here we only give the definition of a big Ramsey structure.

As we remarked in Section 1.7.6, one can define small Ramsey degrees, but it turns out that having finite small Ramsey degrees is the same as having a precompact Ramsey expansion with the expansion property. This is not obvious from the very definitions: Having finite small Ramsey degrees means that for every \mathbf{A} there is some kind of universal and persistent colouring of \mathbf{A} (for formal definitions see e.g. [Zuc16]), while a Ramsey expansion with the expansion property means that all the colourings arise from one global expansion, so it gives some kind of coherence.

If \mathbf{M} has finite big Ramsey degrees, a *big Ramsey structure* is an expansion of \mathbf{M} such that, for every finite \mathbf{A} , the universal and persistent colouring of $\binom{\mathbf{M}}{\mathbf{A}}$ is determined by the expansion. For example, in all the examples we have seen, the corresponding big Ramsey structure is enumerated and has relations describing the internal (or external, this depends on the philosophy) vertices of the tree of types which are necessary for determining the embedding types. Note that a big Ramsey structure has all big Ramsey degrees equal to 1; for example, (ω, \leq) is a big Ramsey structure for itself but also for the countable set with no structure (see [Zuc16] for other examples). In all the known cases, structures with finite big Ramsey degrees in fact do admit a big Ramsey structure, it is however open whether this is true in general.

1.8.11 From big Ramsey degrees to Ramsey classes

In the previous section we defined the analogue of a Ramsey expansion for big Ramsey problems. The natural question is then: Why do we speak mostly about Ramsey expansions in the small Ramsey world, but about big Ramsey degrees in the big Ramsey world? And the reason is, mostly, convenience. The small Ramsey expansions we have encountered so far have been rather easily describable. Moreover, the general theorems have been, historically, phrased in such a way that they are easier to apply to prove the Ramsey property instead of Ramsey degrees.

On the other hand, in the big Ramsey world, big Ramsey structures typically encode special trees with some extra structure. While it is often very much not obvious from the relational description that they describe trees, the trees are the real combinatorial objects behind the arguments and they are what bears good intuition for understanding them. Moreover, the descriptions of the exact big Ramsey degrees / big Ramsey structures are usually very technical and long (which seems to be a feature of the area and not a bug of the present approaches). Nevertheless, it is possible that in the future when we understand the area better there will be good concepts and good terminology for stating theorems about big Ramsey structures in a concise way.

In any case, there is another natural question: Is the age of a big Ramsey structure a Ramsey class? And the answer is positive. If \mathbf{M} is a big Ramsey structure then for every finite $\mathbf{A} \subseteq \mathbf{M}$ it holds that $\mathbf{M} \rightarrow (\mathbf{M})_2^{\mathbf{A}}$, hence all the more so for every finite $\mathbf{B} \subseteq \mathbf{M}$ we have that $\mathbf{M} \rightarrow (\mathbf{B})_2^{\mathbf{A}}$. It then needs a non-trivial compactness argument to obtain that this is equivalent to there being a finite $\mathbf{C} \subseteq \mathbf{M}$ such that $\mathbf{C} \rightarrow (\mathbf{B})_2^{\mathbf{A}}$, in other words, $\text{Age}(\mathbf{M})$ being a Ramsey class.

Note, however, that the expansions which big Ramsey structures describe are typically much richer than the optimal small Ramsey expansions. For example, we know that $\text{Age}((\mathbb{Q}, \leq))$ is already a Ramsey class, but the big Ramsey structure for (\mathbb{Q}, \leq) describes the *Devlin trees*. In general, ages of big Ramsey structures typically consist of tree-like objects. Nevertheless, this correspondence is very useful in the other direction – knowing that the age of a big Ramsey structure needs to be richer than the Ramsey expansion, we for example immediately obtain the big Ramsey structure for two-graphs (it is the same as the big Ramsey structure for the Rado graph).

Not everything is lost in the other direction either. Hubička [Hub20a] has pioneered the usage of big Ramsey proof techniques for small Ramsey classes. It roughly goes as follows: Suppose that we have a proof of finite big Ramsey degrees of some structure \mathbf{M} using a tree theorem for tree T , and suppose that the tree theorem has a finitary version, that is, instead of finding a monochromatic copy of T inside T , we want to find a monochromatic copy of a finite initial segment of T in a larger initial segment of T . This already gives finite small Ramsey degrees (we look at the structures induced on the initial segments and, again, iterate this theorem). However, when working with finite initial segments, one can often significantly reduce the number of persistent embedding types. For example, we know from the beginning what the finite number of vertices of our structure will be, and thus it is possible to embed it so that first we do all splitting to prepare

all the necessary types and only after that we start coding real vertices. This is also the reason why, in the small Ramsey world, trees are not visible in the expansion: The trees are actually there, but they are hidden by the possibility to not mix different kinds of interesting events. On the other hand, if we need to realise infinitely many interesting events of different kinds, they will necessarily have to intertwine. This is a very promising direction and there is hope that it might lead to solving some open problems in the area of (small) Ramsey theory.

1.8.12 Even bigger Ramsey properties?

First we were colouring finite structures by finitely many colours, trying to find monochromatic (oligochromatic) finite structures in Section 1.7. Then we were colouring finite structures by finitely many colours, trying to find monochromatic (oligochromatic) infinite structures in this section. There are two natural generalisations:

If we want to colour by infinitely many colours, we need to somehow address the fact that it is possible to give each substructure its own colour. This is what the *canonical Ramsey theorems* do and it is out of the scope of this thesis.

If we want to colour infinite substructures then we run into problems very fast: There are too many of them and so the axiom of choice will have an easy job to produce bad colourings. However, one can restrict themselves to only nice colourings to recover positive theorems. This is what *topological Ramsey theory* is concerned with, see for example the book of Todorćević [Tod10]. This is a very natural extension of the study of big Ramsey degrees, and indeed, upon understanding the exact big Ramsey degrees one can try to understand the Ramsey behaviour of colouring copies of \mathbf{M} inside \mathbf{M} . This is an even newer and less developed direction where one probably should expect more development in the upcoming years. See for example [DZ23].

Let us conclude by remarking that “infinite” meant “countably infinite” in these paragraphs. If one is willing to work with large cardinals then there are many Ramsey-type theorems to have. For example, the Erdős–Rado theorem [ER56] is an extension of Ramsey’s theorem to uncountable cardinals and is actually one of the ingredients to the forcing proofs of some tree theorems used to prove finiteness of big Ramsey degrees [Dob23]. This is, however, also out of the scope of this thesis.

1.8.13 Metric big Ramsey degrees

Oscillation stability and the *distortion problem* are important concepts in Banach space theory which happen to be closely connected to Ramsey theory. For example, Gowers proved oscillation stability for the Banach space c_0 using his now famous Ramsey theorem for FIN_k as the key ingredient [Gow92]:

Theorem 1.8.20 (Gowers [Gow92]). *Let K be a compact metric space and $\chi: S_{c_0} \rightarrow K$ a Lipschitz map. For every $\varepsilon > 0$, there exists a linear isometric copy X of c_0 in c_0 such that the diameter of $\chi(S_X)$ is less than ε .*

Nguyen Van Thé and Sauer proved an analogous result for the Urysohn sphere:

Theorem 1.8.21 (Nguyen Van Thé–Sauer [NVTS09]). *Let K be a compact metric space, let \mathbf{U}_1 be the Urysohn sphere and let $\chi: \mathbf{U}_1 \rightarrow K$ be a Lipschitz map. For every $\varepsilon > 0$, there exists an isometric copy X of \mathbf{U}_1 in \mathbf{U}_1 such that the diameter of $\chi(X)$ is less than ε .*

The proof of this theorem goes by first proving indivisibility for $\{1, \dots, n\}$ -Urysohn spaces and then, based on ε , approximating \mathbf{U}_1 by such spaces. This method was later generalised by Sauer [Sau13b] to various S -Urysohn spaces.

Note that for metric structures (such as Banach spaces), Lipschitz maps are a good analogue of the discrete concept of a colouring, and a compact metric space is a good analogue of a finite set of colours. As nothing comes for free, we have to pay for this generalisation by settling for a copy which is only close to being monochromatic. Note that if we consider the discrete metric and Lipschitz maps to a finite metric space, we recover the standard category we have been working in so far. See also [BYU10] for a connection to continuous logic.

One can thus ask for analogues of small or big Ramsey degrees for metric structures with Lipschitz colourings into compact spaces. There are presently multiple approaches being developed, see for example [BLALM17] for a small Ramsey result by Bartošová, Lopez-Abad, Lupini, and Mbombo, and [BdRHK23] for an announcement of some big Ramsey results by Bice, de Racourt, Hubička, and the author.

2. Problems, questions and conjectures

In this section I tried to collect various questions, open problems and conjectures from the included papers, as well as from other papers which I have co-authored, there are also some which have not yet been published and a few which originally come from other people but which I endorse.

In all cases, whenever applicable, I link to the original source, and for those questions, problems, and conjectures which are collected from other chapters of this thesis, I use the number with which it appears in the given chapter in order for the reader to be able to easily look up the context.

Whenever it makes sense, I try to give as much context as possible, as well as my opinion on the difficulty, possible approaches or anything else relevant. For this reason, I temporarily switch back to “I” in this chapter in order to emphasize that these are my opinions as opposed to mathematical facts.

I hope to keep an occasionally updated version of this list on my personal website. In any case, please feel free to contact me to find out if I am aware of any progress not mentioned here. Obviously, I will also be happy to hear about new progress, and/or talk about any of these problems.

Conjecture 2.0.1. Let \mathbf{M} be the $\{1, \dots, n\}$ -Urysohn space and let \downarrow be a stationary independence relation of \mathbf{M} . Then \downarrow comes from a shortest path completion with respect to the following operation for some parameter $p \in \{\lceil \frac{n}{2} \rceil\}$:

$$a \oplus_p b = \begin{cases} a + b & \text{if } a + b < p \\ |a - b| & \text{if } |a - b| > p \\ p & \text{otherwise.} \end{cases}$$

The order for the shortest path completion is the natural partial order of \oplus_M , that is, $1 \preceq 2 \preceq \dots \preceq p$ and $n \preceq n - 1 \preceq \dots \preceq p$. All other pairs are incomparable.

This is a special case of the *magic completion algorithm* or *magic monoid* from [ABWH⁺17b] (see also [Kon19]). In general, one can conjecture that the only SIR’s on Cherlin’s primitive 3-constrained metrically homogeneous graphs are those given by the magic completion algorithm. This conjecture has not been explored much to my best knowledge. It is certainly possible that it is simple to prove it or that it is wrong for an easy reason.

The following conjecture and question are probably hard. The first one is a weak form of conjecturing that all finite-language binary symmetric homogeneous structures with a primitive automorphism group and no algebraicity are in fact generalised metric spaces (roughly Conjecture 1 from [Kon19]), the second one is asking whether a metric-like SIR can be always upgraded to give a generalised shortest path completion.

Conjecture 3.5.3 (Conjecture 5.4 from [EHK⁺21]). Every countable homogeneous complete L -edge-labelled graph with $2 \leq |L| < \infty$, primitive automorphism group and trivial algebraic closure admits a metric-like SIR.

Question 3.5.4 (Question 5.5 from [EHK⁺21]). Assume that \mathbb{F} is a transitive countable structure with a metric-like SIR \downarrow such that $\text{tp}(ab) = \text{tp}(ba)$ for every $a, b \in \mathbb{F}$. Can one define a partially ordered commutative semigroup \mathfrak{M} on the 2-types of \mathbb{F} such that \downarrow is the \mathfrak{M} -shortest path independence relation? If the answer is yes, is it true that for every $a \neq b \neq c \in \mathbb{F}$ it holds that $\text{tp}(ab) \preceq \text{tp}(ac) \oplus \text{tp}(bc)$?

Question 3.5.1 (Question 5.1 from [EHK⁺21]). Consider the structure \mathbb{M}_k from Example 3.1.3, that is, the Fraïssé limit of all finite $[n]^k$ -metric spaces (which are in fact semigroup-valued metric spaces in the sense of Section 3.4.1). Is the automorphism group of \mathbb{M}_k simple for $k \geq 2$ and n large enough? (If, for example, $n = 3$, it is in fact a free amalgamation class, as $(2, \dots, 2)$ is a free relation.)

Answering this question is a first necessary step for understanding what role 1-supportedness of a SIR plays. I believe that the assumption of 1-supportedness was necessary for our adaptation of the Tent–Ziegler method in Chapter 3. However, it might just be an artifact of the method. I believe that Question 3.5.1 for $k = 2$ and n large enough is a nice problem to look at and that it might not be difficult to obtain some results.

Question 3.5.2 (Question 5.2 from [EHK⁺21]). Let $\mathbb{U}^\#$ be the Fraïssé limit of the class of all finite complete \mathbb{Q}^+ -edge-labelled graphs (where \mathbb{Q}^+ is the set of all positive rational numbers) which contain no triangles a, b, c with $d(a, b) \geq d(a, c) + d(b, c)$ (that is, the triangle inequality is sharp). Is the automorphism group of $\mathbb{U}^\#$ simple modulo bounded automorphisms?

The *sharp Urysohn space* is a very peculiar structure, because while it does not admit a SIR, it behaves just as the standard Urysohn space from the point of view of finitary combinatorics on its age (Ramsey, EPPA, ...). The reason is that when solving a finite problem, one can always pick a small enough ε and pretend that one is in fact working with generalised metric spaces where the triangle inequality is ε -sharp which allows us to view this structure as a semigroup-valued metric space. While this question asks for simplicity of the automorphism group, the real vague question is how one should approach such structures. It seems that there is a lack of some kind of *saturation* which is not visible by finite eyes.

This structure is also a cautionary example if one wanted to generalise Conjecture 3.5.3 and its relatives to infinite languages, it seems that, when looking at triangle-constrained binary symmetric structures, one needs to distinguish cases based on the complexity of the constraints with respect to some topology on the language.

Question 2.0.2 (Aranda et al., Question 10.3 from [ABWH⁺17b]). *What are the normal subgroups of the automorphism groups of the non-tree-like countably infinite metrically homogeneous graphs from Cherlin’s catalogue?*

For the finite-diameter primitive cases this is handled by [EHK⁺21] resp. Chapter 3. Generalising this approach to the remaining cases could, possibly, be a nice student project.

Question 2.0.3 (Herwig–Lascar [HL00]). *Does the class of all finite tournaments have (coherent) EPPA?*

Originally, Herwig and Lascar only ask for EPPA, asking for coherent EPPA is a natural extension if the original question is answered in the affirmative. They proved that this question is equivalent to a problem in profinite group theory, see Section 1.6.4. The reason why the class of all finite tournaments has stood up to all attacks so far is the following: Most of the EPPA constructions construct witnesses where a lot of pairs of vertices are not in any relations together (and these witnesses are then, potentially, completed in an automorphism-preserving way). This in particular means that the EPPA-witnesses contain a lot of symmetric pairs. Completing to a tournament means that we need to choose an orientation for each such pair, thereby very likely killing many automorphisms. I expect that if the answer is positive, the solution will need to do a lot of group-theoretical arguments, because all automorphisms of tournaments have odd order.

Note that Huang, Pawliuk, Sabok, and Wise [HPSW19] disproved EPPA for a certain version of hypertournaments (different from the one defined in Section 1.4 and used in Chapter 7) defined specifically to obtain a variant of the profinite-topology equivalence for which the topological statement is false. This naturally motivates the following question.

Question 2.0.4. *Does the class of all finite n -hypertournaments (see Section 1.4 for definitions) have (coherent) EPPA? How about the H_4 -free 3-hypertournaments or the even hypertournaments?*

Note that the $\{H_4, O_4\}$ -free 3-hypertournaments do not have EPPA as these structures are just linear orders after naming a least vertex. I expect that the even 3-hypertournaments will behave as a combination of two-graphs and 3-hypertournaments with no restrictions.

The following two questions point out to other simple classes for which EPPA is open, the reason this time being the presence of non-unary algebraicity:

Question 4.13.2 (Question 13.2 from [HKN22]). Let L be the language consisting of a single partial binary function and let \mathcal{C} be the class of all finite L -structures. Does \mathcal{C} have EPPA?

Question 4.13.3 (Question 13.3 from [HKN22]). Does the class of all finite partial Steiner triple systems have EPPA, where one only wants to extend partial automorphism between closed substructures? (A sub-hypergraph H of a Steiner triple system S is *closed* if whenever $\{x, y, z\}$ is a triple of S and $x, y \in H$, then $z \in H$.)

Of course, both of them would be answered if the following question is answered in the affirmative:

Question 4.13.1 (Question 13.1 from [HKN22]). Do Theorems 4.1.1, 4.1.5 and 4.1.6 hold for languages with non-unary functions?

Question 2.0.5 ([EHKN20]). *Does the class of all finite two-graphs have coherent EPPA?*

See Remark 5.4.1. In Chapter 5 we prove that it has EPPA and as such it is the only example that I know of which has EPPA and for which coherent EPPA

is open. Note that the automorphism group of the generic two-graph does have a dense locally-finite subgroup (see Section 5.5). This is thus a candidate for the following question of Siniora:

Question 2.0.6 (Siniora, Question 3 in [Sin17]). *Is there a class of finite structures with EPPA but not coherent EPPA?*

Question 5.8.1 (Question 8.4 from [EHKN20]). For which pairs of classes \mathcal{C} , \mathcal{C}^- such that \mathcal{C}^- is a reduct of \mathcal{C} does it hold that for every $\mathbf{A} \in \mathcal{C}$ there is $\mathbf{B} \in \mathcal{C}$ such that \mathbf{B} is an EPPA-witness for \mathbf{A} (in \mathcal{C}) and furthermore if \mathbf{A}^- and \mathbf{B}^- are the corresponding reducts in \mathcal{C}^- then \mathbf{B}^- is an EPPA-witness for \mathbf{A}^- (in \mathcal{C}^-)?

In Chapter 5 we prove that this property does hold for \mathcal{C} being the class of all finite graphs and \mathcal{C}^- being the class of all finite two-graphs (see Remark 5.8.3). As witnessed by the class of all finite two-graphs, reducts can be interesting examples for studying EPPA even without the extra property from this question. In particular, one can ask the following:

Question 2.0.7 ([EHKN20]). *Which of the reducts of the random graph have (coherent) EPPA?*

There are five reducts. The random graph itself and the set with no structure have been solved. For two-graphs, coherence is open. The last two reducts are the *complementing graph* (adding an isomorphism between the random graph and its complement) and the “complementing two-graph” (the join of the automorphism group of the generic two-graph and the complementing graph) and for them both EPPA and coherent EPPA are open. Let me remark that in a Bachelor thesis in preparation, Beliayeu, a student of Hubička, identified an interesting kind of expansion of the complementing graph which does have EPPA.

Problem 2.0.8 ([HKN22]). *Obtain upper and lower bounds on the size of EPPA-witnesses for various classes of graphs, hypergraphs etc.*

See Remark 4.3.5. This is a broad problem. One possible way of interpreting it are Observations 4.7.8 and 4.9.5. Another possible interpretation is to try to mimic the area of Ramsey numbers and study EPPA numbers of graphs, hypergraphs etc. In Section 1.6.1 we have, for example, seen that every graph on n vertices has an EPPA-witness on $n2^n$ vertices (an easy variant of the proof gives $n2^{n-1}$), and that every graph on n vertices with maximum degree Δ has an EPPA-witness on $\mathcal{O}((n\Delta)^\Delta)$ vertices which gives, in particular, polynomial bounds for bounded degree graphs. Hrushovski [Hru92] asked these questions and also showed that the half-graph (or ladder graph) with each part of size m does not have EPPA-witnesses with fewer vertices than $\Omega(2^m)$.

Yet another possible interpretation of Problem 2.0.8 is the following question which has been asked by Sabok in private communication:

Question 2.0.9. *Is there a class with EPPA in which some members have superexponential lower bound on the size of EPPA-witnesses?*

Question 2.0.10 (Siniora, Question 1 of [Sin17]). *Is there, for every n , a countably infinite structure with n -generic automorphisms but not $(n + 1)$ -generic automorphisms?*

To my best knowledge, the only positive examples are for $n = 1$, hence already for $n = 2$ this question is very interesting.

Problem 2.0.11. *The observation that EPPA and JEP imply amalgamation can be refined to prove that if \mathcal{C} is a class of finite structures and $\mathcal{E} \subseteq \mathcal{C}$ such that \mathcal{C} has a refined variant of EPPA where we only consider partial automorphisms whose domains (and ranges) are isomorphic to a structure from \mathcal{E} , then if \mathcal{C} has JEP then it has the amalgamation property over structures from \mathcal{E} . Find some interesting examples of classes without EPPA but with this “ \mathcal{E} -EPPA”.*

Question 2.0.12. *Recall the definition of automorphism-preserving locally finite subclass (Definition 1.6.2) and of automorphism-preserving completion (Definition 1.6.1). In them, we require that the map $\alpha \mapsto \alpha'$ is a group homomorphism. This is necessary in order to preserve coherent EPPA, but it is not necessary for standard, non-coherent EPPA. Is there an example of a pair of amalgamation classes such that one is not an automorphism-preserving locally finite subclass of the other, but it would be if we did not require $\alpha \mapsto \alpha'$ to be a group homomorphism?*

Such a pair of classes would, likely, exhibit new kinds of behaviours and, as such, would be an important example to keep in mind when trying to advance the area. I believe that proving a negative answer to this question would be very hard.

Conjecture 4.1.7 (Conjecture 1.7 from [HKN19a]). Every strong amalgamation class with EPPA has a precompact Ramsey expansion.

Note that this conjecture implies every ω -categorical structure with EPPA having a precompact Ramsey expansion. Classes known to have EPPA where it is not known if they have a precompact Ramsey expansion include the class of all finite groups [Sin17, PS18] and the class of all finite skew-symmetric bilinear forms¹. This is also a special case of a question of Ivanov [Iva15] asking whether every structure with an amenable automorphism group has a precompact Ramsey expansion (see also Question 7.3 of [EHN19]) as EPPA implies amenability of the automorphism group [KR07].

Question 7.4.4 (Question 1 from [CHKN21]). What is the optimal Ramsey expansion for the class of all finite H_4 -free 3-hypertournaments? Does it have a Ramsey expansion in a finite language? What about H_{n+1} -free n -hypertournaments for $n \geq 4$?

In the area of structural Ramsey theory, there are very general questions (e.g. Question 1.7.7) as well as very general methods for answering these questions positively in concrete instances (Theorem 1.7.17). In essentially all testcases that Theorem 1.7.17 faced, it came out victorious. But it might be because of at least two reasons:

1. Theorem 1.7.17 is very strong, or
2. our testcases are very weak.

¹David M. Evans, personal communication. See also [CH03].

The issue is that it is hard to come up with novel homogeneous structures which do not mimic old homogeneous structures for which we already know that Theorem 1.7.17 works. Thus one needs to turn to the classification programme of homogeneous structures to produce examples, but as classification is very hard and work-intensive, it does not produce new examples in the same rate as we are able to process them. The other source of interesting examples are the ad hoc homogeneous structures discovered as a side-product of some other research (for example, the study of infinite CSP's seems to produce those from time to time). In any case, very few challenging examples have been appearing recently, and thus one needs to cherish those which do appear.

The class of all H_4 -free 3-hypertournaments is, at this point, at the very top of the list of classes for which I would like to understand their Ramsey behaviour. As opposed to other, longer-standing naughty classes, this one is relatively recent (and thus has seen much fewer person-hours spent trying to solve it), it is a strong amalgamation class in a finite relational language and has an extremely simple description. This means that, based on how this question resolves, there are various possible futures, but most of them are quite exciting:

- If the class does not have a precompact (which is equal to finite-language) Ramsey expansion then we have found a counterexample to Question 1.7.7. While I acknowledge that my opinion on this fluctuates quite a lot, I would be rather surprised if this turned out to be the case. To rectify my intuition if this happened, I would probably start claiming that, beyond binary languages (for which we have a lot of quite challenging examples, as opposed to higher-arity languages), the landscape is very wild and that there are probably many more homogeneous structures without nice Ramsey expansions which we have simply not yet found.
- If the class does have a precompact Ramsey expansion, then several things might happen:
 - The most boring possibility (but at the same time one needs to give it the highest probability) is that I was wrong in claiming that the H_4 -free 3-hypertournaments are difficult and that, in fact, one can use Theorem 1.7.17 to prove its Ramseyness without much complications.
 - Another possibility (the second most likely in my opinion) is that once the correct Ramsey expansion is guessed, it will not be hard to prove its Ramseyness using Theorem 1.7.17 and, as such, it will either earn its spot among the exceptions to the Ramsey cookbook (Section 1.7.4), or – which would be more exciting – it will direct us to a new recipe in the cookbook.
 - The last possibility is that H_4 -free 3-hypertournaments have a nice Ramsey expansion but there is no proof using Theorem 1.7.17. This is either due to us not finding the correct interpretation yet, or there provably being no proof using Theorem 1.7.17. The latter case would be extremely exciting because it would enable us to work on a strengthening of Theorem 1.7.17 which would then, hopefully, give us a better understanding of the general conditions for Ramseyness, perhaps even allowing us to take a step towards a general result on the existence of

nice Ramsey expansions. An argument why this might happen is that H_3 -free 2-hypertournaments are simply linear orders and their Ramsey property does **not** follow from Theorem 1.7.17, Ramsey’s theorem is always an ingredient in the proofs of more complicated Ramsey theorems. I do not give this argument too much weight because while linear orders are “unstructured” (in that there is only one isomorphism type on n vertices), there are many non-isomorphic H_4 -free 3-hypertournaments on n vertices.

In Chapter 7 we prove that the class of all finite linearly ordered H_4 -free 3-hypertournaments has arbitrarily large obstacles to completions which suggests that Theorem 1.7.17 should not be able to prove that the class is Ramsey. My expectation is that this is because one needs to consider a richer expansion (remember the class of all finite two-graphs and note that O_4 -free 3-hypertournaments are a close relative of two-graphs), but at this point I do not have any ideas what this expansion might be. I tried advertising this problem on various conferences and I was informed that, in discrete geometry, 3-hypertournaments are called *signotopes* and they represent arrangements of pseudohyperplanes, but it seemed to me that the H_4 -free 3-hypertournaments did not correspond to any subclass of signotopes which I have seen considered in the area. But my confidence in all of this is rather low and I invite the reader to try to do their own research and prove me wrong.

Question 2.0.13. *Does the class of all finite groups have a precompact Ramsey expansion?*

Here my weak conjecture would be that the answer is negative as I would expect there to be a complicated structure of imaginaries. Note that, in the definition of Ramsey property, one only colours subgroups, hence, in particular, there are infinitely many types of vertices according to their order. Note that there are some partition theorems for finite abelian groups but they do not seem to be relevant for this question. [Voi80]

Conjecture 2.0.14 (Aranda et al., Conjecture 10.4 from [ABWH⁺17b]). *Let $T_{m,n}$ be the tree-like metrically homogeneous graph with parameters m, n . If $(m, n) \neq (2, 2)$ then $\text{Aut}(T_{m,n})$ is not amenable.*

Here, given $2 \leq m, n \leq \infty$, the graph $T_{m,n}$ is defined to be the (regular) graph in which the blocks (two-connected components) are cliques of order n and every vertex is a cut vertex, lying in precisely m blocks. This is the last missing piece to fully classify which automorphism groups of Cherlin’s metrically homogeneous graphs are amenable. I believe that with the right tools, this is a relatively easy problem and that it might be a good problem for a student. For more discussion, see the original paper [ABWH⁺17b].

Problem 9.7.1 (Problem 7.1 from [BCdR⁺23]). Characterise the exact big Ramsey degrees of countable universal L -hypergraphs for finite L with no unary relations.

Problem 9.7.2 (Problem 7.2 from [BCdR⁺23]). Characterise the exact big Ramsey degrees of countable universal L -hypergraphs for L which has no unary relations and has finitely many relations of every arity.

This is a harder variant of Problem 9.7.1. The following two problems are further strengthenings:

Problem 9.7.3 (Problem 7.3 from [BCdR⁺23]). Characterise the exact big Ramsey degrees of countable universal L -hypergraphs for L which has at most countably many unary relations and has finitely many relations of every arity greater than one.

Problem 9.7.4 (Problem 7.4 from [BCdR⁺23]). Characterise the exact big Ramsey degrees for structures considered in Theorem 9.1.1.

If the problems are solved then there is the natural question whether they admit a big Ramsey structure.

There is a draft in preparation in which we intend to solve Problems 9.7.1 and 9.7.2. Adding unary relations is usually more of a technical burden than a conceptual one, although in the presence of infinitely many unaries some of the intuitions begin to break (for example, while any persistent and universal colouring does come from a tree, the big Ramsey structure is no longer tree per se, it needs to have at least one decreasing chain of levels).

Conjecture 9.7.8 (Conjecture 7.8 from [BCdR⁺23]). Let L be a relational language with finitely many unary relations and infinitely many relations of some arity $a \geq 2$. Let \mathbf{H} be an unrestricted L -structure realising all relations from L . Then there is a finite L -structure \mathbf{A} whose big Ramsey degree in \mathbf{H} is infinite. Moreover, the number of vertices of \mathbf{A} only depends on a .

Question 10.3.4 (Question 3.9 from [BCD⁺23a]). Fix a finite partial order \mathbf{A} and let CT be the coding tree of the generic poset. Let $r < \omega$ and let $\gamma: \text{AEmb}(\text{CT}^{\mathbf{A}}, \text{CT}) \rightarrow r$ be a coloring. Is there $h \in \text{AEmb}(\text{CT}, \text{CT})$ such that $h \circ \text{AEmb}(\text{CT}^{\mathbf{A}}, \text{CT})$ is monochromatic?

See Chapter 10 for the relevant definitions. In the bigger picture, there are currently two kinds of tree theorems used for proving big Ramsey degrees – one works with ordinary trees while the other works with *coding trees* where some vertices have a unary mark and these unary marks need to be respected by the morphisms. In a way, after the structure is defined on them, ordinary trees represent all possible enumerations of the structure glued to each other while coding trees pick an enumeration from the start and work within the enumeration. Coding tree theorems tend to be proved by very infinitary methods (forcing, the Erdős–Rado theorem etc.) while standard theorems can be proved by the more finitary method of combinatorial forcing. It is, at this point, not clear whether one of the techniques is stronger than the other. The generic poset is an example which can be handled by the standard tree theorems but coding tree theorems seem to fundamentally break on it. This question formalizes this and asks whether the coding tree theorems only seem to fundamentally break, or whether they really do break.

Problem 12.1.3 (Problem 1.4 from [ABC⁺23]). Let $L = \{E, H\}$ be a language with one binary relation E and one ternary relation H . Let \mathbf{F} be the L -structure where $F = \{0, 1, 2, 3\}$, $E_{\mathbf{F}} = \{(1, 0), (1, 2), (1, 3)\}$, $H_{\mathbf{F}} = \{(0, 2, 3)\}$. Denote by \mathcal{K} the class of all L -structures \mathbf{A} such that there is no monomorphism $\mathbf{F} \rightarrow \mathbf{A}$. Do \mathcal{K} -universal L -structures have finite big Ramsey degrees?

This is an example of a free amalgamation class which does not have the type-respecting amalgamation property (see Chapter 12) and hence our current methods cannot prove finiteness of its big Ramsey degrees. In a way, this is at this point a central problem in the area, as the current general methods (papers still in preparation) seem to be quite powerful on structures with the type-respecting amalgamation property. If the answer to this questions is positive then this means that we need to develop new general theorems which do not require this property. On the other hand, if the answer is negative, then maybe one can prove an analogue of Nešetřil’s observation that Ramsey classes have the amalgamation property:

Question 2.0.15. *Does every big Ramsey structure “have” the type-respecting amalgamation property?*

Note that this question is somewhat ambiguous in the “have” part. There are multiple possible ways of phrasing this (cf. Theorem 12.1.5) and at this point it is not clear to me which way one should choose.

Problem 2.0.16 (Balko et al., Problem 9.4.1 from [BCD⁺21b]). *Characterize the big Ramsey degrees of the S -Urysohn space for or $S = \{1, \dots, n\}$ with $n \geq 3$. Does it admit a big Ramsey structure?*

Problem 2.0.17 (Balko et al., Problem 9.4.2 from [BCD⁺21b]). *Let S be a finite set of non-negative real numbers for which the S -Urysohn space exists. Characterize the big Ramsey degrees of the S -Urysohn space. Does it admit a big Ramsey structure?*

Problem 2.0.18. *Characterize the metric big Ramsey degrees of the S -Urysohn spaces for infinite sets S .*

See [BdRHK23] for definitions. The first example to look at (and which we are actively looking at) is the Urysohn sphere. For unbounded sets S they will likely not have compact metric big Ramsey degrees at all. A solution of the previous problem will probably be an ingredient in a possible solution to this problem.

Problem 2.0.19 (Balko et al., Problem 9.4.3 from [BCD⁺21b]). *Characterize exact big Ramsey degrees for the 3-constrained finite diameter non-bipartite non-antipodal metrically homogeneous graphs from Cherlin’s catalog. Do they admit big Ramsey structures?*

Let us remark that for the S -Urysohn spaces with finite S as well as the metrically homogeneous graphs from the above problem, finiteness of big Ramsey degrees follows from the results in Chapter 11.

Problem 2.0.20 (Balko et al., Problem 9.5.1 from [BCD⁺21b]). *Which of the following classes have finite big Ramsey degrees? Can the exact big Ramsey degrees be characterized? Which of the corresponding Fraïssé limits admit big Ramsey structures?*

1. *The class of all finite structures in a language consisting of a single unary function or its variant where there is a bound on the size of the closure of a vertex,*
2. *The small Ramsey expansion of some variant of the Hrushovski construction, see [EHN19],*
3. *The class of all finite structures in a language consisting of two unary relations and one binary function such that each vertex is in exactly one unary relation and the functions go from vertices of one unary to vertices of the other unary (cf. Section 4.12.4),*
4. *The class of all finite structures in a language consisting of one binary function.*

Problem 2.0.21 (Kechris–Pestov–Todorčević [KPT05]). *Which homogeneous structures admit generalizations of the Galvin–Příkrý and/or Ellentuck theorems?*

Conjecture 2.0.22 (Balko et al., Conjecture 9.2.1 from [BCD⁺21b]). *Let L be a finite language, let \mathcal{F} be an infinite collection of finite irreducible L -structures such that no member of \mathcal{F} embeds to any other member of \mathcal{F} and let \mathbf{M} be the Fraïssé limit of $\text{Forb}_{\text{he}}(\mathcal{F})$. Then there is $\mathbf{A} \in \text{Forb}_{\text{he}}(\mathcal{F})$ such that the big Ramsey degree of \mathbf{A} in \mathbf{M} is infinite. Furthermore, the number of vertices of \mathbf{A} only depends on L .*

3. Simplicity of the automorphism groups of generalised metric spaces

DAVID M. EVANS, JAN HUBIČKA, MATĚJ KONEČNÝ, YIBEI LI, MARTIN ZIEGLER

Abstract Tent and Ziegler proved that the automorphism group of the Urysohn sphere is simple and that the automorphism group of the Urysohn space is simple modulo bounded automorphisms. A key component of their proof is the definition of a stationary independence relation (SIR). In this paper we prove that the existence of a SIR satisfying some extra axioms is enough to prove simplicity of the automorphism group of a countable structure. The extra axioms are chosen with applications in mind, namely homogeneous structures which admit a “metric-like amalgamation”, for example all primitive 3-constrained metrically homogeneous graphs of finite diameter from Cherlin’s list.

3.1 Introduction

In 2011, Macpherson and Tent [MT11] proved that the automorphism groups of Fraïssé limits of free amalgamation classes are simple. This was followed by two papers of Tent and Ziegler [TZ13b, TZ13a] where they prove that the isometry group of the Urysohn space (the unique complete separable homogeneous metric space universal for all finite metric spaces) modulo bounded isometries (i.e. isometries f with a finite bound on the distance between x and $f(x)$) is simple and that the isometry group of the Urysohn sphere is simple. Later, Evans, Ghadernezhad and Tent [EGT16] proved simplicity for automorphism groups of some Hrushovski constructions, and Li [Li18] proved simplicity for the structures from Cherlin’s list of 26 primitive triangle-constrained homogeneous structures with 4 binary symmetric relations (see appendix of [Che98]).

More recently, Tent and Ziegler’s method was generalised to asymmetric structures. Li [Li19] proved that the automorphism groups of some of Cherlin’s asymmetric structures in the appendix of [Che98] are simple. The same result for non-trivial linearly ordered free homogeneous structures has been proved independently by Calderoni, Kwiatkowska and Tent [CKT20] and Li [Li20a]. Also in [Li20a], simplicity was proved for the automorphism groups of the universal n -linear orders for $n \geq 2$. Another recent example where (non-stationary) independence relations have been used to prove strong results about automorphism groups of structures is a paper by Kaplan and Simon [KS19].

In this paper, we adapt the methods of Tent and Ziegler and prove the following theorem (definitions and examples will be given in the upcoming paragraphs).

Theorem 3.1.1. *Let \mathbb{F} be a transitive countable relational structure with a bounded 1-supported metric-like stationary independence relation \perp . Then $\text{Aut}(\mathbb{F})$ is simple.*

As direct corollaries of Theorem 3.1.1, we get the following two more concrete results, for which the definitions will be given in Section 3.4.

Theorem 3.1.2. *Let $\mathfrak{M} = (M, \oplus, \preceq)$ be a finite archimedean partially ordered commutative semigroup with at least two elements and let \mathbb{F} be a homogeneous \mathfrak{M} -metric space which realises all distances. Assume that \mathbb{F} admits an \mathfrak{M} -shortest path independence relation \downarrow and that \downarrow is a 1-supported SIR. Then $\text{Aut}(\mathbb{F})$ is simple.*

Theorem 3.1.3. *If G is a countably infinite metrically homogeneous graph which corresponds to one of the primitive 3-constrained finite-diameter classes from Cherlin’s catalogue [Che11], then $\text{Aut}(G)$ is simple.*

3.1.1 Stationary independence relations

The notion of stationary independence relations (Definition 3.1.1) was developed by Tent and Ziegler in their paper on the Urysohn space [TZ13b]. It has several generalisations (e.g. for structures with closures [EGT16]), but for our purposes the original variant suffices.

Let \mathbb{F} be a relational structure and let $A, B \subseteq \mathbb{F}$ be finite subsets. We will identify them with the substructures induced by \mathbb{F} on A and B respectively and by AB we will denote the union $A \cup B$ (and hence also the substructure induced by \mathbb{F} on AB). If the set $A = \{a\}$ is singleton, we may write a instead of $\{a\}$. Uppercase letters will denote sets while lowercase will denote the elements of the structure, which we call *vertices* owing to the combinatorial background of part of the authors. As is usual in this area, if $A \subseteq \mathbb{F}$, we sometimes assume that it has some implicit enumeration. This is clear from the context and should not cause any confusion.

Let $A, X \subseteq \mathbb{F}$. By the *type of A over X* (denoted by $\text{tp}(A/X)$) we mean the orbit of A under the action of the stabilizer subgroup of $\text{Aut}(\mathbb{F})$ with respect to X . If $p = \text{tp}(A/X)$, we say that $B \subseteq \mathbb{F}$ *realises p* (and denote it as $B \models p$) if B lies in p , in other words, if there is an automorphism of \mathbb{F} fixing X pointwise which maps A to B . To simplify the notation, we write $\text{tp}(A)$ for $\text{tp}(A/\emptyset)$. Our types correspond to realised types in a (strongly) homogeneous structure in the standard model-theoretic terminology. In fact, we may assume that the language is chosen so that \mathbb{F} is homogeneous, that is, partial automorphisms between finite substructures of \mathbb{F} extend to automorphisms.

Definition 3.1.1 (Stationary Independence Relation). Let \mathbb{F} be a relational structure. A ternary relation \downarrow on finite subsets of \mathbb{F} is called a *stationary independence relation* (SIR, with $A \downarrow_C B$ being pronounced “ A is independent from B over C ”) if the following conditions are satisfied:

SIR1 (*Invariance*). The independence of finite subsets of \mathbb{F} only depends on their type. In particular, for every automorphism f of \mathbb{F} , we have $A \downarrow_C B$ if and only if $f(A) \downarrow_{f(C)} f(B)$.

SIR2 (*Symmetry*). If $A \downarrow_C B$, then $B \downarrow_C A$.

SIR3 (*Monotonicity*). If $A \downarrow_C BD$, then $A \downarrow_C B$ and $A \downarrow_{BC} D$.

SIR4 (*Existence*). For every A, B and C in \mathbb{F} , there is some $A' \models \text{tp}(A/C)$ with $A' \downarrow_C B$.

SIR5 (*Transitivity*) If $A \downarrow_C B$ and $A \downarrow_{BC} B'$, then $A \downarrow_C B'$.

SIR6 (*Stationarity*) If A and A' have the same type over C and are both independent over C from some set B then they also have the same type over BC .

Note that by an observation of Baudisch [Bau16], these axioms are redundant as Monotonicity can be derived from the rest of them. Stationary independence relations correspond to “canonical amalgamations” by putting $A \downarrow_C B$ if and only if the canonical amalgamation of AC and BC over C is isomorphic to ABC . The notion of canonical amalgamations can be formalised, see [ABWH⁺17b].

To make our proofs shorter, we will sometimes use Symmetry, Monotonicity and Existence implicitly. The following observation which follows from Invariance will be useful later.

Observation 3.1.4. *If \mathbb{F} is a relational structure, \downarrow a SIR on \mathbb{F} and $A \downarrow_C B$, then $A \downarrow_C BC$.*

Definition 3.1.2 (k -supported SIR). Let k be a positive integer. We say that a SIR \downarrow is k -supported if for every a, b, C such that $a \downarrow_C b$ there is $C' \subseteq C$ such that $|C'| \leq k$ and $a \downarrow_{C'} b$.

Observation 3.1.5. *For $k = 1$, k -supportedness is equivalent to:*

(1-supportedness) *If $a \downarrow_C b$ and $C = C_1 \cup C_2$ then $a \downarrow_{C_1} b$ or $a \downarrow_{C_2} b$.*

We say that a structure \mathbb{F} is *transitive* if $\text{tp}(a) = \text{tp}(b)$ for every $a, b \in \mathbb{F}$.

Definition 3.1.3 (Metric-like SIR). Let \mathbb{F} be a relational structure with a SIR \downarrow . We say that \downarrow is *metric-like* if the following conditions are satisfied:

1. If $a \notin A$, then $a \not\downarrow_A a$.
2. For every $a \in \mathbb{F}$ there is $b \in \mathbb{F}$ such that $a \neq b$ and $a \not\downarrow_\emptyset b$.
3. (*Perfect triviality*) If $A \downarrow_C B$ and $C \subseteq C'$ then $A \downarrow_{C'} B$.

Lemma 3.1.6. *Let \mathbb{F} be a relational structure with a SIR \downarrow which satisfies Perfect triviality. Then \downarrow satisfies*

1. (Metricity) *If $A \downarrow_{C_1 C_2} B$ and $C_1 \downarrow_D B$ then $A \downarrow_{C_2 D} B$.*
2. (Triviality) *If $A \downarrow_B C$ and $A \downarrow_B D$ then $A \downarrow_B CD$.*

Proof. First assume that $A \downarrow_{C_1 C_2} B$ and $C_1 \downarrow_D B$. By Perfect triviality we have that $C_1 \downarrow_{C_2 D} B$ and $A \downarrow_{C_1 C_2 D} B$. Using Transitivity it follows that $A \downarrow_{C_2 D} B$, which proves Metricity.

Now assume that $A \downarrow_B C$ and $A \downarrow_B D$. By Perfect triviality we get $A \downarrow_{BC} D$ and by Observation 3.1.4 and Monotonicity it then follows that $A \downarrow_{BC} CD$. Using Transitivity together with $A \downarrow_B C$ then implies $A \downarrow_B CD$. \square

In fact, Metricity is equivalent to Perfect triviality if \downarrow is a SIR. The following is a simple corollary of Triviality which will be useful later.

Corollary 3.1.7. *If $a \downarrow_{\emptyset} x$ for every $x \in X$, then $a \downarrow_{\emptyset} X$.*

Definition 3.1.4 (Geodesic sequence). Let \mathbb{F} be a relational structure with a SIR \downarrow . We say that a sequence $a_1, \dots, a_n \in \mathbb{F}$ of pairwise distinct vertices of \mathbb{F} is *geodesic* if for every $1 \leq i < j < k \leq n$ it holds that $a_i \downarrow_{a_j} a_k$.

Definition 3.1.5. Let \mathbb{F} be a relational structure with a SIR \downarrow . We say that \downarrow is *bounded* if it satisfies

(*Boundedness*) There exists an integer k_0 such that if a_0, \dots, a_k is a geodesic sequence with $k \geq k_0$, then $a_0 \downarrow_{\emptyset} a_k$.

We denote the smallest such k_0 by $\|\downarrow\|$.

The reader is encouraged to have the following examples in mind when reading this paper.

Example 3.1.1. Let \mathbb{F} be the Fraïssé limit of all finite metric spaces using only distances $\{0, 1, \dots, n\}$ for some fixed $n \geq 2$ (clearly, one can view a metric space as a relational structure by introducing a binary relation for every distance). Define \downarrow on \mathbb{F} by putting $A \downarrow_C B$ if and only if for every $a \in A$ and every $b \in B$ it holds that $d(a, b) = \min(\{n\} \cup \{d(a, c) + d(b, c) : c \in C\})$. It is straightforward to check that \downarrow is a bounded 1-supported metric-like SIR with $\|\downarrow\| = n$.

For the Urysohn sphere, the only axiom which we do not have at hand is, paradoxically, Boundedness.

Example 3.1.2. Let \mathbb{U}_1 be the Urysohn sphere, that is, the unique homogeneous separable complete metric space with distances from $[0, 1]$ which is universal for all finite metric spaces with distances from $[0, 1]$. We will denote its metric by d . Define the relation \downarrow on finite subsets of \mathbb{U}_1 by putting $A \downarrow_C B$ if and only if for every $a \in A$ and every $b \in B$ it holds that $d(a, b) = \min(\{1\} \cup \{d(a, c) + d(b, c) : c \in C\})$. One can check that \downarrow is a 1-supported metric-like SIR, but does not satisfy Boundedness, as for every k one can find a geodesic sequence with $k + 1$ vertices such that the distance of every consecutive pair of them is smaller than $\frac{1}{k}$.

Example 3.1.3 (k -supported metric-like SIR). Let $k \geq 1$ and $n \geq 3$ be integers. Put $S = \{1, \dots, n\}^k \cup \{0\}^k$, let A be a set and let $d: A^2 \rightarrow S$ be a function. Let \preceq be the product order on S (i.e. $(a_1, \dots, a_k) \preceq (b_1, \dots, b_k)$ if and only if $a_i \leq b_i$ for every $1 \leq i \leq k$) and let \oplus be the component-wise addition on S capped at n (i.e. $(a_1, \dots, a_k) \oplus (b_1, \dots, b_k) = (c_1, \dots, c_k)$, where $c_i = \min(n, a_i + b_i)$ for every $1 \leq i \leq k$).

We say that (A, d) is an $[n]^k$ -metric space if the following holds for every $x, y, z \in A$:

1. $d(x, y) = d(y, x)$,
2. $d(x, y) = (0, \dots, 0)$ if and only if $x = y$,

$$3. d(x, z) \preceq d(x, y) \oplus d(y, z).$$

One can verify that the class of all finite $[n]^k$ -metric spaces is a Fraïssé class. Consider the structure $\mathbb{M}_k = (M_k, d)$, which is the Fraïssé limit of the class of all $[n]^k$ -metric spaces, and define \downarrow on \mathbb{M}_k by putting $A \downarrow_C B$ if and only if for every $a \in A$ and every $b \in B$ it holds that $d(a, b) = \inf_{\preceq} \{d(a, c) \oplus d(c, b) : c \in C\}$. As \preceq has a maximum, the infimum of the empty set is (n, \dots, n) .

It is easy to verify that \downarrow is a bounded metric-like SIR. Moreover, it is k -supported, but not k' -supported for any $k' < k$, which is witnessed by vertices $a, b, c_1, \dots, c_k \in \mathbb{M}_k$ such that $a \downarrow_{\{c_1, \dots, c_k\}} b$, $d(a, c_i) = (1, \dots, 1)$ for every i and $d(b, c_i)$ is equal to 1 on the i -th coordinate and equal to 2 everywhere else.

3.2 Geodesic sequences

In this section we prove some auxiliary results about geodesic sequences which will be used later. Fix a transitive relational structure \mathbb{F} with a metric-like SIR \downarrow .

Lemma 3.2.1. *Let a_1, \dots, a_n be a geodesic sequence of vertices of \mathbb{F} and let $b \in \mathbb{F} \setminus \{a_n\}$. Then there is $a_{n+1} \models \text{tp}(b/a_n)$ such that a_1, \dots, a_{n+1} is a geodesic sequence.*

Proof. Using Existence, pick $a_{n+1} \models \text{tp}(b/a_n)$ such that $a_1 \cdots a_{n-1} \downarrow_{a_n} a_{n+1}$. Consider any $1 \leq i < j \leq n - 1$. By Monotonicity, $a_i \downarrow_{a_n} a_{n+1}$ and hence, by Perfect triviality, $a_i \downarrow_{a_j a_n} a_{n+1}$. Since a_1, \dots, a_n is a geodesic sequence, we know that $a_i \downarrow_{a_j} a_n$. Transitivity now implies that $a_i \downarrow_{a_j} a_{n+1}$ and hence a_1, \dots, a_{n+1} is a geodesic sequence. \square

Lemma 3.2.2. *Let $a, b, c \in \mathbb{F}$ be distinct such that $a \downarrow_{\emptyset} b$. There is a geodesic sequence $a_0, a_1, \dots, a_n \in \mathbb{F}$ satisfying the following:*

1. $a = a_0$ and $b = a_n$, and
2. for every $0 \leq i \leq n - 1$ it holds that $\text{tp}(a_i a_{i+1}) = \text{tp}(ac)$,
3. $n = \|\downarrow\|$.

Proof. First observe that since all vertices have the same type, for every $v \in \mathbb{F}$ there is $v' \in \mathbb{F}$ such that $\text{tp}(vv') = \text{tp}(ac)$. Put $n = \|\downarrow\|$ and use Lemma 3.2.1 repeatedly to obtain a geodesic sequence a, x_1, \dots, x_n such that all consecutive pairs of vertices have the type $\text{tp}(ac)$. We know that $a \downarrow_{\emptyset} x_n$. By Stationarity, $\text{tp}(x_n/a) = \text{tp}(b/a)$, hence there exists an automorphism f of \mathbb{F} which fixes a and maps x_n to b . By Invariance, $f(a), f(x_1), \dots, f(x_n)$ has the desired properties. \square

Lemma 3.2.3. *Let v_1, \dots, v_k and w_1, \dots, w_k be geodesic sequences of vertices of \mathbb{F} such that for every $1 \leq i < k$ we have $\text{tp}(v_i v_{i+1}) = \text{tp}(w_i w_{i+1})$. Then $\text{tp}(v_1 \cdots v_k) = \text{tp}(w_1 \cdots w_k)$.*

Proof. We shall prove by induction on m that $\text{tp}(v_1 \cdots v_m) = \text{tp}(w_1 \cdots w_m)$. For $m = 2$ this is true by the assumption. Assume now that the statement is true for some m . Using the fact that v_1, \dots, v_k and w_1, \dots, w_k are geodesic sequences and Triviality we get that $v_1 \cdots v_{m-1} \downarrow_{v_m} v_{m+1}$ and $w_1 \cdots w_{m-1} \downarrow_{w_m} w_{m+1}$. By the assumption we have $\text{tp}(v_m v_{m+1}) = \text{tp}(w_m w_{m+1})$, hence Stationarity together with Invariance and the induction hypothesis give $\text{tp}(v_1 \cdots v_{m+1}) = \text{tp}(w_1 \cdots w_{m+1})$. \square

Proposition 3.2.4. *Let a, b, c be vertices of \mathbb{F} satisfying the following:*

1. $a \downarrow_b c$,
2. *there is a geodesic sequence $a = v_1, \dots, v_k = b$,*
3. *there is a geodesic sequence $b = w_1, \dots, w_\ell = c$.*

Then there is a geodesic sequence $a = x_1, \dots, x_{k+\ell-1} = c$ such that $\text{tp}(x_1 \cdots x_k) = \text{tp}(v_1 \cdots v_k)$ and $\text{tp}(x_k \cdots x_{k+\ell-1}) = \text{tp}(w_1 \cdots w_\ell)$.

Proof. Use Lemma 3.2.1 and the fact that all vertices have the same type $\ell - 1$ times to extend v_1, \dots, v_k by vertices w'_2, \dots, w'_ℓ such that $v_1, \dots, v_k, w'_2, \dots, w'_\ell$ is a geodesic sequence and for every $1 \leq i < \ell$ we have $\text{tp}(w'_i w'_{i+1}) = \text{tp}(w_i w_{i+1})$, where we put $w'_1 = v_k$ to simplify the notation.

In particular, w'_1, \dots, w'_ℓ is a geodesic sequence. Using Lemma 3.2.3 we get that $\text{tp}(w_1 \cdots w_\ell) = \text{tp}(w'_1 \cdots w'_\ell)$, so in particular $\text{tp}(w_1 w_\ell) = \text{tp}(w'_1 w'_\ell)$. Since $w_1 = w'_1 = v_k$, we have that $\text{tp}(w_\ell / v_k) = \text{tp}(w'_\ell / v_k)$. By the hypothesis and the construction, $w_\ell \downarrow_{v_k} v_1$ and $w'_\ell \downarrow_{v_k} v_1$. Stationarity implies that $w'_\ell \models \text{tp}(w_\ell / v_1 v_k)$, so in particular $w'_\ell \models \text{tp}(w_\ell / v_1)$.

In other words, there is an automorphism g of \mathbb{F} such that $g(v_1) = v_1$ and $g(w'_\ell) = w_\ell$. The image of $v_1, \dots, v_k, w'_2, \dots, w'_\ell$ under g then gives the desired geodesic sequence $x_1, \dots, x_{k+\ell-1}$. \square

Let $a, b \in \mathbb{F}$ be distinct. We say that b is almost free from a if $a \not\downarrow_\emptyset b$ and for every $c \in \mathbb{F}$ different from a, b such that $a \downarrow_b c$ it holds that $a \downarrow_\emptyset c$.

Observation 3.2.5. *Let $a, b \in \mathbb{F}$ be such that b is almost free from a . For every $a', b' \in \mathbb{F}$ such that $\text{tp}(a'b') = \text{tp}(ab)$ it holds that b' is almost free from a' .*

Lemma 3.2.6. *Suppose that \downarrow is bounded. For every $a \in \mathbb{F}$ and every finite $X \subseteq \mathbb{F}$ such that $a \notin X$ there is $b \in \mathbb{F}$ such that a is almost free from b , b is almost free from a , and $b \downarrow_a X$. In particular, $b \not\downarrow_\emptyset a$ and $b \downarrow_\emptyset X$.*

Proof. We claim that there exist $a', b' \in \mathbb{F}$ such that b' is almost free from a' and a' is almost free from b' . Suppose that this is true. Since \mathbb{F} is transitive, there is an automorphism f such that $f(a') = a$. Pick $b \models \text{tp}(f(b')/a)$ such that $b \downarrow_a X$. By Observation 3.2.5, b is almost free from a and a is almost free from b . The “in particular” part is immediate using Corollary 3.1.7.

Hence it suffices to prove the claim. Pick $a', b' \in \mathbb{F}$ such that $b' \not\downarrow_\emptyset a'$ and the length of the longest geodesic sequence starting at a' finishing at b' is as large as possible. (As \downarrow is bounded, such a', b' exist.) Pick $c \in \mathbb{F}$ such that $a' \downarrow_{b'} c$. By Proposition 3.2.4, we can extend the geodesic sequence from a' to b' by some

$c' \models \text{tp}(c/b')$. By the properties of a', b' we get that $a' \perp_{\emptyset} c'$. Invariance and Stationarity then imply that $a' \perp_{\emptyset} c$ and consequently b' is almost free from a' .

To prove that a' is almost free from b' , pick $c \in \mathbb{F}$ such that $b' \perp_{a'} c$. Since the reverse of a geodesic sequence is a geodesic sequence, we extend the geodesic sequence from b' to a' by some $c' \models \text{tp}(c/a')$ as above. Suppose that $b' \not\perp_{\emptyset} c'$. Since \mathbb{F} is transitive, there is an automorphism f such that $f(b') = a'$. The image of the geodesic sequence from b' to c' is then a geodesic sequence starting at a' which is longer than the geodesic sequence from a' to b' we started with. This is a contradiction, hence $b' \perp_{\emptyset} c'$. As before, we get that a' is almost free from b' which concludes the proof. \square

3.3 Proof of Theorem 3.1.1

We will closely follow the proof from the Tent–Ziegler paper on the Urysohn sphere [TZ13a] and use the following result by Tent and Ziegler [TZ13b].

Definition 3.3.1. Let \mathbb{F} be a countable structure with a stationary independence relation \perp , let $g \in \text{Aut}(\mathbb{F})$, let $A \subseteq \mathbb{F}$ be finite and let $p = \text{tp}(a/A)$ be a type. We say that g moves p almost maximally if there is a realisation $x \models p$ such that

$$x \perp_A g(x).$$

Theorem 3.3.1 (Corollary 5.4, [TZ13b]). *Let \mathbb{F} be a countable structure with a stationary independence relation and let g be an automorphism of \mathbb{F} which moves every type over every finite set almost maximally. Then every element of $\text{Aut}(\mathbb{F})$ is a product of sixteen conjugates of g .*

Throughout the section, we fix \mathbb{F} and \perp as in Theorem 3.1.1 (\mathbb{F} is a transitive countable relational structure with a bounded 1-supported metric-like stationary independence relation \perp) and put $G = \text{Aut}(\mathbb{F})$. As before, we may assume that \mathbb{F} is homogeneous (this will slightly simplify the proof of Lemma 3.3.5).

Lemma 3.3.2. *If $g \in G$ is not the identity then there is $a \in \mathbb{F}$ and $h \in G$ which is a product of $\|\perp\|$ conjugates of g such that $a \perp_{\emptyset} h(a)$.*

Proof. Let $a \in \mathbb{F}$ be such that $a \neq g(a)$ and pick $b \in \mathbb{F}$ such that $a \perp_{\emptyset} b$ (Existence). Use Lemma 3.2.2 to obtain a geodesic sequence $a = a_0, \dots, a_n = b$ such that $n = \|\perp\|$ and for every $0 \leq i \leq n-1$ we have $\text{tp}(a_i a_{i+1}) = \text{tp}(a g(a))$. This means that there are automorphisms h_0, \dots, h_{n-1} such that $h_i(a) = a_i$ and $h_i(g(a)) = a_{i+1}$. Then $h_i g h_i^{-1}$ moves a_i to a_{i+1} and the statement follows. \square

Lemma 3.3.3. *Let $g \in G$ be such that for some $a \in \mathbb{F}$ we have $a \perp_{\emptyset} g(a)$. Then for every finite set $A \subset \mathbb{F}$ there is $x \in \mathbb{F}$ with $x \perp_{\emptyset} A$ and $x \neq g(x)$.*

Proof. We may assume that $a \in A$. Put $Y = A \cup g^{-1}(A)$ and choose $b \in \mathbb{F}$ with $b \neq a$ and $b \not\perp_{\emptyset} a$ (\perp is metric-like) such that moreover $b \perp_a Y$ (Existence and Invariance). This means that $b \notin g^{-1}(A)$ (if $b \in g^{-1}(A)$, then $b \in Y$, so $b \perp_a b$, which is in contradiction with part (1) of Definition 3.1.3) and hence $g(b) \notin A$. We know that $a \perp_{\emptyset} g^{-1}(a)$ (by Invariance) and also $b \perp_a g^{-1}(a)$,

thus $b \perp_{\emptyset} g^{-1}(a)$ (Transitivity) and so $g(b) \perp_{\emptyset} a$ (Invariance). This means that $b \neq g(b)$ and therefore $g(b) \notin A \cup \{b\}$.

Use Lemma 3.2.6 to obtain $x \in \mathbb{F}$ such that $x \not\perp_{\emptyset} g(b)$ and $x \perp_{\emptyset} Ab$. By Monotonicity, $x \perp_{\emptyset} A$ and $x \perp_{\emptyset} b$, hence also $g(x) \perp_{\emptyset} g(b)$, thus $x \neq g(x)$. \square

Let $X \subset \mathbb{F}$ be a finite set and let $a \in \mathbb{F}$ be such that $a \perp_{\emptyset} X$. We call the type $\text{tp}(a/X)$ a *free type*. (It is the unique such type over X .)

Lemma 3.3.4. *Let $g \in G$ be such that for every free type p there is a realisation $a \models p$ with $g(a) \neq a$. Then for every finite $X \subset \mathbb{F}$ and every type $q = \text{tp}(x/X)$ with $x \notin X$, there is a realisation $c \models q$ such that $g(c) \neq c$.*

Proof. Let a be a vertex such that $a \perp_{\emptyset} X$ and $g(a) \neq a$ (a exists by the assumptions of this lemma) and let $b \models q$ be such that $b \perp_X g(a)$.

If $b \not\perp_{\emptyset} g(a)$ then pick $c \models q$ such that $c \perp_X ag(a)$. This means that $c \not\perp_{\emptyset} g(a)$ (by Stationarity and Invariance) and $c \perp_{\emptyset} a$ (by Transitivity), giving us $g(c) \neq c$.

So we have $b \perp_{\emptyset} g(a)$. Use Lemma 3.2.6 to obtain $a' \in \mathbb{F}$ such that $a' \not\perp_{\emptyset} b$, $a' \perp_{\emptyset} X$, and a' is almost free from b . By Stationarity, we have that $a \models \text{tp}(a'/X)$, hence there is $f \in G$ fixing X pointwise such that $f(a') = a$. Put $c' = f(b)$. In particular, $c' \models q$, $a \not\perp_{\emptyset} c'$, and a is almost free from c' (Observation 3.2.5).

Choose $c \models \text{tp}(c'/Xa)$ such that $c \perp_{Xa} g(a)$. In particular, $c \not\perp_{\emptyset} a$ (Invariance). By Observation 3.2.5, a is almost free from c . Using 1-supportedness, $c \perp_{Xa} g(a)$ implies that either $c \perp_a g(a)$ (in which case $c \perp_{\emptyset} g(a)$ and hence $g(c) \neq c$), or $c \perp_X g(a)$. In this case we know that $\text{tp}(c/X) = \text{tp}(b/X)$ and $b \perp_X g(a)$ (using Perfect triviality on $b \perp_{\emptyset} g(a)$), hence by Stationarity and Invariance, $c \perp_{\emptyset} g(a)$, thus again $g(c) \neq c$. \square

We say that $g \in G$ *moves type p by distance k* if there is $a \models p$ and a geodesic sequence $a = a_0, \dots, a_k = g(a)$. If $p = \text{tp}(x/X)$ is a type and h is an automorphism or a partial automorphism defined on a finite set such that $X \subseteq \text{Dom}(h)$, we denote $h(p) = \text{tp}(h'(x)/h'(X))$, where h' is some automorphism of \mathbb{F} extending h (remember that we assumed that \mathbb{F} is homogeneous).

Lemma 3.3.5. *Let $g \in G$ be such that g moves all types almost maximally or by distance n . Then there exists $h \in G$ such that $[g, h] = g^{-1}h^{-1}gh$ moves all types almost maximally or by distance $2n$.*

Proof. As in [TZ13a], we construct h by a “back-and-forth” construction as the union of a chain of finite partial automorphisms. We show the following: Let h' be already defined on a finite set U and let $p = \text{tp}(x/X)$ be a type. Then h' has an extension h such that $[g, h]$ moves p almost maximally or by distance $2n$.

We can assume that $X \cup g^{-1}(X) \subseteq U$. Put $V = h'(U)$. Let a' be a realisation of p such that $a' \perp_X Ug^{-1}(U)$ and let b' be a realisation of $h'(\text{tp}(a'/U))$ (which is a type over V). By the hypothesis on g there are realisations $a \models \text{tp}(a'/Ug^{-1}(U))$ and $b \models \text{tp}(b'/V)$ such that either $a \perp_{Ug^{-1}(U)} g(a)$, or there is a geodesic sequence $a = a_0, \dots, a_n = g(a)$ and similarly for b . We also have

$$a \underset{X}{\perp} Ug^{-1}(U) \text{ and } b \underset{h'(X)}{\perp} V.$$

Let h_0 be the isomorphism $Ua \simeq Vb$, and c be a realisation of $h_0^{-1}(\text{tp}(g(b)/Vb))$ (which is a type over Ua) such that $c \downarrow_{Ua} g(a)$. Put h to be the isomorphism $Uac \simeq Vbg(b)$. Observe that $[g, h](a) = g^{-1}(c)$. It remains to prove that a witnesses that $[g, h]$ moves p almost maximally or by distance $2n$.

Since $a \downarrow_X g^{-1}(U)$, we know that $g(a) \downarrow_{g(X)} U$. Using Metricity, we get

$$c \downarrow_{g(X)a} g(a),$$

thus from 1-supportedness we know that either $c \downarrow_a g(a)$ or $c \downarrow_{g(X)} g(a)$. In the second case we get $g^{-1}(c) \downarrow_X a$, which implies that $[g, h]$ moves p almost maximally. Hence we can assume that

$$c \downarrow_a g(a).$$

By the choice of a and b we know that one of the following cases occurs:

1. First suppose that there are geodesic sequences $b = b_0, \dots, b_n = g(b)$ and $g(a) = a_0, \dots, a_n = a$ (the reverse of a geodesic sequence is a geodesic sequence by Symmetry). From the construction we know that $\text{tp}(ac) = \text{tp}(bg(b))$. This implies that there is a geodesic sequence $a = c_0, \dots, c_n = c$. Since $g(a) \downarrow_a c$, Proposition 3.2.4 gives a geodesic sequence starting at $g(a)$ and finishing at c using $2n+1$ vertices (including c and $g(a)$). Finally, taking the image of this sequence under g^{-1} gives a geodesic sequence starting at a and finishing at $g^{-1}(c) = [g, h](a)$ using $2n+1$ vertices. This means that a witnesses that $[g, h]$ moves p by distance $2n$.
2. Now assume that $a \downarrow_{Ug^{-1}(U)} g(a)$. Then in fact we have $a \downarrow_X g(a)$, because $a \downarrow_X Ug^{-1}(U)$ (Metricity). As $U \supseteq Xg^{-1}(X)$, $a \downarrow_X U$ also implies $g(a) \downarrow_{g(X)} X$ (by Invariance and Monotonicity), which together with $a \downarrow_X g(a)$ implies $a \downarrow_{g(X)} g(a)$ (Metricity). Thus from $c \downarrow_a g(a)$ we get $c \downarrow_{g(X)} g(a)$ (yet again Metricity) and thus $g^{-1}(c) \downarrow_X a$, i.e. a witnesses that $[g, h]$ moves p almost maximally.
3. Otherwise we have $b \downarrow_V g(b)$. Using that h is an isomorphism of Uac and $Vbg(b)$ and Invariance we obtain $a \downarrow_U c$. Then we get $a \downarrow_X c$, because $a \downarrow_X U$ (Metricity), and then, combining with $c \downarrow_a g(a)$ using Metricity again, we obtain $c \downarrow_X g(a)$. As in the previous case, $a \downarrow_X U$ implies $g(a) \downarrow_{g(X)} X$ and hence $c \downarrow_{g(X)} g(a)$, or $g^{-1}(c) \downarrow_X a$, i.e. a witnesses that $[g, h]$ moves p almost maximally.

□

Now we prove the following proposition, Theorem 3.1.1 is then its direct consequence.

Proposition 3.3.6. *Let \mathbb{F} be a countable relational structure with a bounded 1-supported metric-like stationary independence relation \downarrow and let g be a non-identity automorphism of \mathbb{F} . Then there is an automorphism of \mathbb{F} which is a product of at most $2\|\downarrow\|^2$ conjugates of g and g^{-1} and moves every type over every finite set almost maximally.*

Proof. From Lemma 3.3.2 we get an automorphism g_0 which is a product of at most $\|\downarrow\|$ conjugates of g such that there is $a \in \mathbb{F}$ with $a \downarrow_{\emptyset} g_0(a)$. Using Lemma 3.3.3 we get that in fact for every free type there is a realisation which is not fixed by g_0 .

Let $p = \text{tp}(x/X)$ be a type. Either $x \in X$ (then $x \downarrow_X g(x)$, hence g_0 moves q almost maximally), or $x \notin X$ and thus by Lemma 3.3.4 there is a realisation of p which is not fixed by g_0 . This means that g_0 moves all types almost maximally or by distance 1.

Put $n = \lceil \log_2(\|\downarrow\|) \rceil$ and construct a sequence g_0, g_1, \dots, g_n of automorphisms of \mathbb{F} using Lemma 3.3.5 such that every g_i moves all types almost maximally or by distance 2^i , and if $i \geq 1$ then g_i is a product of two conjugates of g_{i-1} and g_{i-1}^{-1} . For g_n we get that it moves every type almost maximally or by distance at least $\|\downarrow\|$. In the latter case, we have for every type p a realisation $a \models p$ and a geodesic sequence $a = a_0, \dots, a_k = g(a)$, where $k \geq \|\downarrow\|$. Boundedness (Definition 3.1.5) implies that $a \downarrow_{\emptyset} g(a)$, i.e. g_n moves p almost maximally, and hence g_n moves all types almost maximally.

By the construction, g_n is a product of at most $2^{\lceil \log_2(\|\downarrow\|) \rceil}$ conjugates of g_0 and g_0^{-1} , hence a product of at most $2^{\lceil \log_2(\|\downarrow\|) \rceil} \|\downarrow\| \leq 2\|\downarrow\|^2$ conjugates of g and g^{-1} . \square

Proof of Theorem 3.1.1. Let g be a non-identity automorphism of \mathbb{F} . We need to prove that if N is a normal subgroup of G such that $g \in N$, then $N = G$. If $g \in N$, then clearly $g^{-1} \in N$. Let $h \in G$. By Proposition 3.3.6 and Theorem 3.3.1, we know that h can be written as a product of conjugates of g and g^{-1} , hence $h \in N$. This is true for every $h \in G$, hence $N = G$ and G is simple. \square

3.4 Corollaries

In this section we prove Theorems 3.1.2 and 3.1.3.

3.4.1 Semigroup-valued metric spaces

We say that a tuple $\mathfrak{M} = (M, \oplus, \preceq)$ is a *partially ordered commutative semigroup* if the following hold:

1. (M, \oplus) is a commutative semigroup,
2. (M, \preceq) is a partial order which is reflexive ($a \preceq a$ for every $a \in M$),
3. for every $a, b \in M$ it holds that $a \preceq a \oplus b$, and
4. for every $a, b, c \in M$ it holds that if $b \preceq c$ then $a \oplus b \preceq a \oplus c$ (\oplus is monotone with respect to \preceq).

\mathfrak{M} is *archimedean* if for every $a, b \in \mathfrak{M}$ there is an integer n such that $n \times a \succeq b$, where by $n \times a$ we mean

$$\underbrace{a \oplus a \oplus \dots \oplus a}_{n \text{ times}}.$$

Note that if \mathfrak{M} is archimedean and non-trivial, it follows that \mathfrak{M} does not have an identity.

Let L be a set. An L -edge-labelled graph is a tuple $\mathbf{A} = (A, E, d)$, where $E \subseteq \binom{A}{2}$ and d is a function $E \rightarrow L$. Clearly, the set E can be inferred from the function d and thus we omit it. For simplicity, we write $d(x, y)$ instead of $d(\{x, y\})$ and we put $d(x, x) = 0$, where 0 is a symbol which is not an element of \mathfrak{M} . When convenient, we naturally understand 0 as the neutral element with respect to \oplus and as the minimum element of \preceq .

We say that \mathbf{A} is *complete* if the graph (A, E) is a complete graph. Note that an L -edge-labelled graph can equivalently be viewed as a relational structure with an irreflexive binary symmetric relation R^m for every $m \in L$ such that every pair of vertices is in at most one relation.

For a partially ordered commutative semigroup $\mathfrak{M} = (M, \oplus, \preceq)$, a complete \mathfrak{M} -edge-labelled graph $\mathbf{A} = (A, d)$ is an \mathfrak{M} -metric space if for every triple $a, b, c \in A$ of distinct vertices it holds that $d(a, b) \preceq d(a, c) \oplus d(b, c)$ (the *triangle inequality*).

Let \mathbb{F} be an \mathfrak{M} -metric space. We say that \mathbb{F} *admits an \mathfrak{M} -shortest path independence relation* if for every $a, b \in \mathbb{F}$ and $C \subseteq \mathbb{F}$ finite we have that $\{d(a, c) \oplus d(c, b) : c \in C\}$ has an infimum with respect to \preceq (note that C can be empty which implies that \mathfrak{M} has maximum $\inf_{\preceq}(\emptyset)$). If \mathbb{F} admits an \mathfrak{M} -shortest path independence relation, then its *\mathfrak{M} -shortest path independence relation* is a ternary relation \downarrow defined on finite subsets of \mathbb{F} by putting $A \downarrow_C B$ if and only if for every $a \in A$ and every $b \in B$ it holds that $d(a, b) = \inf_{\preceq} \{d(a, c) \oplus d(c, b) : c \in C\}$.

Generalising concepts of Sauer [Sau12], Conant [Con19] (see also [HKN21]) and Braunfeld [Bra17] (see also [KPR18]), Hubička, Konečný and Nešetřil [Kon19, HKN18] introduced the framework of semigroup-valued metric spaces, which served as a motivation for this paper. Given a partially ordered commutative semigroup $\mathfrak{M} = (M, \oplus, \preceq)$ and a “nice” family \mathcal{F} of \mathfrak{M} -edge-labelled cycles, the structures of interest are \mathfrak{M} -metric spaces which moreover contain no homomorphic images of members of \mathcal{F} . We will denote the class of all such finite structures $\mathcal{M}_{\mathfrak{M}}^{\mathcal{F}}$.

The conditions of \mathcal{F} are strong enough that one can then prove that $\mathcal{M}_{\mathfrak{M}}^{\mathcal{F}}$ is a strong amalgamation class, its Fraïssé limit admits an \mathfrak{M} -shortest path independence relation which is a SIR (provided that \mathfrak{M} has a maximum, otherwise one can still get a *local* SIR), it has EPPA (for background, see [HKN22, Sin17]) and a precompact Ramsey expansion (for background, see [HN19, NVT15]), but they are general enough that most known binary symmetric homogeneous structures can be viewed as such a semigroup-valued metric space. In fact, it is conjectured that every primitive transitive homogeneous structure in a finite binary symmetric language with trivial algebraic closures admits such an interpretation (Conjecture 1 in [Kon19]).

Now we are ready to prove Theorem 3.1.2.

Proof of Theorem 3.1.2. We need to prove that \downarrow is metric-like and bounded. (In fact, we do not need 1-supportedness for this, we only need it later in order to apply Theorem 3.1.1.)

Since \mathbb{F} is homogeneous, all vertices have the same type. As $d(x, y) = 0$ if and only if $x = y$ and \oplus is monotone with respect to \preceq , it follows that if $a \notin A$, then $a \not\downarrow_A a$. The fact that there are $a \neq b \in \mathbb{F}$ such that $a \not\downarrow_{\emptyset} b$ follows from

Stationarity, the fact that \mathfrak{M} has at least two elements (remember that $0 \notin \mathfrak{M}$) and the fact that \mathbb{F} realises all distances.

Suppose now that $a \downarrow_C b$. If there was $c' \in \mathbb{F} \setminus C$ such that $a \not\downarrow_{C \cup \{c'\}} b$, this would mean that $\inf_{\preceq} \{d(a, c) \oplus d(c, b) : c \in C \cup \{c'\}\} \prec \{d(a, c) \oplus d(c, b) : c \in C\} = d(a, b)$, hence $d(a, c') \oplus d(c', b) \not\preceq d(a, b)$, in other words, abc' violates the triangle inequality which is a contradiction. Consequently, \downarrow satisfies Perfect triviality and hence \downarrow is metric-like.

Next we prove that \downarrow is bounded. Denote by 1 the maximum element of \mathfrak{M} (\mathfrak{M} is finite and hence there is such an element). Assume that there are $a, b \in \mathfrak{M}$ such that $a \oplus b = a$. This means (by associativity) that $a \oplus (n \times b) = a$ for every n . Let $c \in \mathfrak{M}$ be arbitrary. By archimedeanity there is n such that $n \times b \succeq c$. But then $a = a \oplus (n \times b) \succeq c$. Hence $a \succeq c$ for every $c \in \mathfrak{M}$, that is, $a = 1$. In other words, for every $a, b \in \mathfrak{M} \setminus \{1\}$ it holds that $a \oplus b \succ a$, which implies that whenever $a_1, \dots, a_{|\mathfrak{M}|} \in \mathfrak{M}$, then

$$\bigoplus_{i=1}^{|\mathfrak{M}|} a_i = 1.$$

We can use this observation to prove that $\|\downarrow\| \leq |\mathfrak{M}|$. Indeed, if $a_0, \dots, a_{|\mathfrak{M}|}$ is a geodesic sequence, we know that $d(a_0, a_{i+1}) = d(a_0, a_i) \oplus d(a_i, a_{i+1})$. Using induction we get that

$$d(a_0, a_{|\mathfrak{M}|}) = d(a_1, a_2) \oplus d(a_2, a_3) \oplus \dots \oplus d(a_{|\mathfrak{M}|-1}, a_{|\mathfrak{M}|}),$$

that is, $d(a_0, a_{|\mathfrak{M}|})$ is a sum of $|\mathfrak{M}|$ elements of \mathfrak{M} and hence $d(a_0, a_{|\mathfrak{M}|}) = 1$, which means that indeed $a_0 \downarrow_{\emptyset} a_{|\mathfrak{M}|}$.

We have proved that \downarrow is bounded and metric-like, hence we can apply Theorem 3.1.1 to show that $\text{Aut}(\mathbb{F})$ is simple. \square

Note that whenever \preceq is a linear order, the corresponding \mathfrak{M} -shortest path independence relation is necessarily 1-supported. The following theorem is a direct consequence of this fact, Theorem 3.1.2 and existing results on semigroup-valued metric spaces [Kon19, HKN18].

Let $S \subseteq \mathbb{R}^+$ be a finite subset of positive reals such that the following operation $\oplus_S: S^2 \rightarrow S$ is associative:

$$a \oplus_S b = \max\{x \in S : x \leq a + b\}.$$

Delhommé, Laflamme, Pouzet, and Sauer [DLPS07] studied and Sauer later classified [Sau13a, Sau13b] such subsets. Ramsey expansions for all such classes of (S, \oplus_S, \leq) -metric spaces were obtained by Hubička and Nešetřil [HN19], and Hubička, Konečný, Nešetřil and Sauer [HKNS20] (Nguyen Van Thé [NVT09] earlier proved some partial results). We contribute to the study of such classes by the following result:

Theorem 3.4.1. *Let $S \subseteq \mathbb{R}^+$ be a finite subset of positive reals such that $\mathfrak{M}_S = (S, \oplus_S, \leq)$ is an archimedean partially ordered commutative semigroup. Then the automorphism group of the Fraïssé limit of the class of all finite \mathfrak{M}_S -metric spaces is simple.*

3.4.2 Metrically homogeneous graphs

A *metrically homogeneous graph* is a graph whose path-metric is a homogeneous metric space. Cherlin [Che11, Che22] gave a list of such graphs by describing the corresponding amalgamation classes of metric spaces. The vast majority of the list is occupied by the 5-parameter classes $\mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$, where δ denotes the diameter of such spaces (i.e. they only use distances $\{1, \dots, \delta\}$) and the other four parameters describe four different families of forbidden triangles (for example, all triangles of odd perimeter smaller than $2K_1$ are forbidden).

Aranda, Bradley-Williams, Hubička, Karamanlis, Kompatscher, Konečný and Pawliuk [ABWH⁺17b, ABWH⁺17a, ABWH⁺21] studied EPPA, Ramsey expansions and (local) SIR's for these classes (see also [Kon18, EHKN20, Kon20]). In particular, if $\mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$ is primitive (i.e. it is neither antipodal nor bipartite) and δ is finite, it can be shown using another result of Hubička, Kompatscher and Konečný [HKK18] that these (local) stationary independence relations are 1-supported and can be viewed as \mathfrak{M} -shortest path independence relations [Kon19] with a finite archimedean \mathfrak{M} , which means that Theorem 3.1.3 is a direct consequence of Theorem 3.1.2.

3.5 Conclusion

We conclude with two questions and a conjecture. The first question is a particular instance of the general question whether 1-supportedness is necessary.

Question 3.5.1. *Consider the structure \mathbb{M}_k from Example 3.1.3, that is, the Fraïssé limit of all finite $[n]^k$ -metric spaces (which are in fact semigroup-valued metric spaces in the sense of Section 3.4.1). Is the automorphism group of \mathbb{M}_k simple? (For $k \geq 2$ and n large enough – if, for example, $n = 3$, it is in fact a free amalgamation class, as $(2, \dots, 2)$ is a free relation.)*

The obvious next step is to generalise our results to countable archimedean semigroups which do not have to contain a maximum element, thereby obtaining an analogue of Tent and Ziegler's result on the Urysohn space [TZ13b]. We believe that such a generalisation is quite straightforward. However, there are structures in infinite language which do not even admit a SIR, although they are also very much metric-like. One example is the *sharp Urysohn space*:

Question 3.5.2. *Let $\mathbb{U}^\#$ be the Fraïssé limit of the class of all finite complete \mathbb{Q}^+ -edge-labelled graphs (here \mathbb{Q}^+ is the set of all positive rational numbers) which contain no triangles a, b, c with $d(a, b) \geq d(a, c) + d(b, c)$ (that is, the triangle inequality is sharp). Is the automorphism group of $\mathbb{U}^\#$ simple modulo bounded automorphisms?*

Note that if we consider \mathbb{N} instead of \mathbb{Q}^+ , the resulting structure can be understood as an \mathfrak{M} -metric space (putting $a \oplus b = a + b - 1$ and $a \preceq b$ if $a \leq b$).

Remark 3.5.1. The sharp Urysohn space is a very peculiar structure, because although it does not admit a SIR, it has EPPA, APA and it is Ramsey when equipped with a (free) linear order.

The following conjecture and question are closely related to a conjecture from [Kon19].

Conjecture 3.5.3. *Every countable homogeneous complete L -edge-labelled graph with $2 \leq |L| < \infty$, primitive automorphism group and trivial algebraic closure admits a metric-like SIR.*

Question 3.5.4. *Assume that \mathbb{F} is a transitive countable structure with a metric-like SIR \downarrow such that $\text{tp}(ab) = \text{tp}(ba)$ for every $a, b \in \mathbb{F}$. Can one define a partially ordered commutative semigroup \mathfrak{M} on the 2-types of \mathbb{F} such that \downarrow is the \mathfrak{M} -shortest path independence relation? If the answer is yes, is it true that for every $a \neq b \neq c \in \mathbb{F}$ it holds that $\text{tp}(ab) \preceq \text{tp}(ac) \oplus \text{tp}(bc)$?*

The obvious special cases of Question 3.5.4 are for finitely many 2-types, 1-supported \downarrow , bounded \downarrow , and their combinations. It is not true that the conditions of Question 3.5.4 imply that the structure at hand is an \mathfrak{M} -metric space in the sense of [Kon19, HKN18]. For example, suppose that \mathbb{F} is the Fraïssé limit of the class of all $[n]^1$ -metric spaces which also contain a ternary relation R such that if $(a, b, c) \in R$, then $d(a, b) = d(b, c) = d(c, a) = 1$. The standard $([n], +, \leq)$ -shortest path independence relation is the desired SIR on \mathbb{F} .

4. All those EPPA classes (Strengthenings of the Herwig–Lascar theorem)

JAN HUBIČKA, MATĚJ KONEČNÝ, JAROSLAV NEŠETŘIL

Abstract Let \mathbf{A} be a finite structure. We say that a finite structure \mathbf{B} is an EPPA-witness for \mathbf{A} if it contains \mathbf{A} as a substructure and every isomorphism of substructures of \mathbf{A} extends to an automorphism of \mathbf{B} . Class \mathcal{C} of finite structures has the extension property for partial automorphisms (EPPA, also called the Hrushovski property) if it contains an EPPA-witness for every structure in \mathcal{C} .

We develop a systematic framework for combinatorial constructions of EPPA-witnesses satisfying additional local properties and thus for proving EPPA for a given class \mathcal{C} . Our constructions are elementary, self-contained and lead to a common strengthening of the Herwig–Lascar theorem on EPPA for relational classes defined by forbidden homomorphisms, the Hodkinson–Otto theorem on EPPA for relational free amalgamation classes, its strengthening for unary functions by Evans, Hubička and Nešetřil and their coherent variants by Siniora and Solecki. We also prove an EPPA analogue of the main results of J. Hubička and J. Nešetřil: *All those Ramsey classes (Ramsey classes with closures and forbidden homomorphisms)*, thereby establishing a common framework for proving EPPA and the Ramsey property.

There are numerous applications of our results, we include a solution of a problem related to a class constructed by the Hrushovski predimension construction. We also characterize free amalgamation classes of finite Γ_L -structures with relations and unary functions which have EPPA.

4.1 Introduction

Let \mathbf{A} and \mathbf{B} be finite structures (e.g. graphs, hypergraphs or metric spaces) such that \mathbf{A} is a substructure of \mathbf{B} . We say that \mathbf{B} is an *EPPA-witness* for \mathbf{A} if every isomorphism of substructures of \mathbf{A} (a *partial automorphism of \mathbf{A}*) extends to an automorphism of \mathbf{B} . We say that a class \mathcal{C} of finite structures has the *extension property for partial automorphisms (EPPA, also called the Hrushovski property)* if for every $\mathbf{A} \in \mathcal{C}$ there is $\mathbf{B} \in \mathcal{C}$ which is an EPPA-witness for \mathbf{A} .

In 1992, Hrushovski [Hru92] established that the class of all finite graphs has EPPA. This result was used by Hodges, Hodkinson, Lascar, and Shelah to show the small index property for the random graph [HHLS93]. After this, the quest of identifying new classes of structures with EPPA continued with a series of papers including [Her95, Her98, HL00, HO03, Sol05, Ver08, Con19, Ott20, ABWH⁺17b, HKN19a, Kon19, HKN18, EHKN20].

In particular, Herwig and Lascar [HL00] proved EPPA for certain relational classes with forbidden homomorphisms. Solecki [Sol05] used this result to prove EPPA for the class of all finite metric spaces. This was independently obtained

by Vershik [Ver08], see also [Pes08, Ros11b, Ros11a, Sab17, HKN19a] for other proofs. Some of these proofs are combinatorial [HKN19a], others are using the profinite topology on free groups and the Ribes–Zaleskii [RZ93] and Marshall Hall [Hal49] theorems. Solecki’s argument was refined by Conant [Con19] for certain classes of generalised metric spaces and metric spaces with (some) forbidden subspaces. In [ABWH⁺17b], these techniques were carried further and a layer was added on top of the Herwig–Lascar theorem to show EPPA for many classes of metrically homogeneous graphs from Cherlin’s catalogue [Che22] (see also exposition in [Kon18]).

There are known EPPA classes for which the Herwig–Lascar theorem is not well suited. In particular, EPPA for free amalgamation classes of relational structures was shown by Sinióra and Solecki [SS19] using results of Hodkinson and Otto [HO03]. It was noticed by Ivanov [Iva15] that a lemma on permomorphisms from Herwig’s paper [Her98, Lemma 1] can be used to show EPPA for structures with definable equivalences on n -tuples with infinitely many equivalence classes. Evans, Hubička, and Nešetřil [EHN21] strengthened the aforementioned construction of Hodkinson and Otto and established EPPA for free amalgamation classes in languages with relations and unary functions (e.g. the class of k -orientations arising from a Hrushovski construction [EHN19] or the class of all finite bowtie-free graphs [EHN19]).

We give a combinatorial, elementary, and fully self-contained proof of a strengthening of all the aforementioned results [Her98, HL00, HO03, EHN21] and their coherent variants by Sinióra and Solecki [SS19, Sin17]. This has a number of applications which we list in Section 4.1.3. In particular, in Section 4.12.5 we present a solution of a problem related to a class constructed by the Hrushovski predimension construction.

4.1.1 Γ_L -structures

Before presenting the summary of our results, let us introduce the structures we are dealing with. With applications in mind, we generalise the standard notion of model-theoretic L -structures in two directions. We consider functions which go to subsets of the vertex set and we also equip the languages with a permutation group Γ_L . Our morphisms will consist of a map between vertices together with a permutation of the language: The standard notions of homomorphism, embedding etc. are generalised naturally, see Section 4.2 for formal definitions. If Γ_L consists of the identity only and the ranges of all functions consist of singletons, one gets back the standard model-theoretic L -structures together with the standard mappings, the standard definition of EPPA etc.

4.1.2 The main results

We now state the principal results of this paper together with a short discussion.

A structure is *irreducible* if it is not a free amalgamation of its proper substructures. If \mathbf{B} is an EPPA-witness for \mathbf{A} , we say that it is *irreducible structure faithful* if whenever \mathbf{C} is an irreducible substructure of \mathbf{B} , then there is an automorphism $g \in \text{Aut}(\mathbf{B})$ such that $g(C) \subseteq A$. A class has *irreducible structure faithful EPPA* if it has EPPA and all EPPA-witnesses can be chosen to be irre-

ducible structure faithful. (This is a natural generalization of the clique faithful EPPA introduced by Hodkinson and Otto [HO03] to structures with functions.) Coherent EPPA is a “functorial” strengthening of EPPA introduced by Siniora and Solecki [SS19, Sin17] and is defined in Section 4.2.6.

In this paper we prove two main theorems. The “base” unrestricted theorem, formulated as Theorem 4.1.1, gives irreducible structure faithful coherent EPPA for the class of all finite Γ_L -structures, strengthening results of Herwig [Her98, Lemma 1], Hodkinson and Otto [HO03], its coherent variant of Siniora and Solecki [SS19], and Evans, Hubička, and Nešetřil [EHN21].

Theorem 4.1.1 (Construction of an unrestricted EPPA-witness). *Let L be a language consisting of relations and unary functions equipped with a permutation group Γ_L and let \mathbf{A} be a finite Γ_L -structure. If \mathbf{A} lies in a finite orbit of the action of Γ_L by relabelling, then there is a finite Γ_L -structure \mathbf{B} , which is an irreducible structure faithful coherent EPPA-witness for \mathbf{A} .*

Consequently, the class of all finite Γ_L -structures has irreducible structure faithful coherent EPPA if every finite Γ_L -structure lies in a finite orbit of the action of Γ_L by relabelling.

Here, the action of Γ_L by relabelling is defined such that $g \in \Gamma_L$ sends a Γ_L -structure \mathbf{A} to a Γ_L -structure \mathbf{A}' on the same vertex set where the relations and functions are relabelled according to g (see Definition 4.2.1). Note that in particular if Γ_L is finite (e.g. $\Gamma_L = \{\text{id}_L\}$ or L is finite), then every Γ_L -structure lies in a finite orbit of the action of Γ_L by relabelling.

After that, we provide a theorem which, given a finite irreducible structure faithful (coherent) EPPA-witness \mathbf{B}_0 for \mathbf{A} , produces a finite irreducible structure faithful (coherent) EPPA-witness \mathbf{B} for \mathbf{A} while providing extra control over the local structure of \mathbf{B} :

Theorem 4.1.2 (Construction of a restricted EPPA-witness). *Let L be a language consisting of relations and unary functions equipped with a permutation group Γ_L , let \mathbf{A} be a finite irreducible Γ_L -structure, let \mathbf{B}_0 be a finite EPPA-witness for \mathbf{A} and let $n \geq 1$ be an integer. There is a finite Γ_L -structure \mathbf{B} satisfying the following.*

1. \mathbf{B} is an irreducible structure faithful EPPA-witness for \mathbf{A} .
2. There is a homomorphism-embedding $\mathbf{B} \rightarrow \mathbf{B}_0$.
3. For every substructure \mathbf{C} of \mathbf{B} on at most n vertices there is a tree amalgamation \mathbf{D} of copies of \mathbf{A} and a homomorphism-embedding $f: \mathbf{C} \rightarrow \mathbf{D}$.
4. If \mathbf{B}_0 is coherent then so is \mathbf{B} .

Here, a *tree amalgamation of copies of \mathbf{A}* is any structure which can be created by a series of free amalgamations of copies of \mathbf{A} over its substructures (see Definition 4.9.1).

Theorem 4.1.1 contains the condition that \mathbf{A} needs to lie in a finite orbit of the action of Γ_L by relabelling. We prove that this is in fact necessary:

Theorem 4.1.3. *Let L be a language equipped with a permutation group Γ_L and let \mathbf{A} be a finite Γ_L -structure. If \mathbf{A} lies in an infinite orbit of the action of Γ_L by relabelling then there is no finite Γ_L -structure \mathbf{B} which is an EPPA-witness for \mathbf{A} .*

We can combine Theorems 4.1.1 and 4.1.3 with an easy observation used earlier [HO03, EHN21, Sin17] and characterize free amalgamation classes of finite Γ_L -structures which have (irreducible structure faithful coherent) EPPA, provided that all functions in the language are unary. The following theorem strengthens results of Hodkinson and Otto [HO03], Evans, Hubička, and Nešetřil [EHN21], and Siniora [Sin17], and in particular implies irreducible structure faithful coherent EPPA for the class of all graphs, K_n -free graphs or k -regular hypergraphs.

Corollary 4.1.4. *Let L be a language consisting of relations and unary functions equipped with a permutation group Γ_L and let \mathcal{K} be a free amalgamation class of finite Γ_L -structures. Then \mathcal{K} has EPPA if and only if every $\mathbf{A} \in \mathcal{K}$ lies in a finite orbit of the action of Γ_L on Γ_L -structures by relabelling. Moreover, if \mathcal{K} has EPPA, then it has irreducible structure faithful coherent EPPA.*

We also provide two corollaries of Theorem 4.1.2, which might be easier to apply in some cases. The first corollary is a direct strengthening of the Herwig–Lascar theorem [HL00, Theorem 3.2] and its coherent variant of Solecki and Siniora [SS19, Theorem 1.10]. For a set \mathcal{F} of Γ_L -structures, we denote by $\text{Forb}_{\text{he}}(\mathcal{F})$ the set of all finite and countable Γ_L structures \mathbf{A} such that there is no $\mathbf{F} \in \mathcal{F}$ with a homomorphism-embedding $\mathbf{F} \rightarrow \mathbf{A}$.

Theorem 4.1.5. *Let L be a language consisting of relations and unary functions equipped with a permutation group Γ_L . Let \mathcal{F} be a finite family of finite Γ_L -structures and let $\mathbf{A} \in \text{Forb}_{\text{he}}(\mathcal{F})$ be a finite Γ_L -structure which lies in a finite orbit of the action of Γ_L by relabelling. If there exists a (not necessarily finite) structure $\mathbf{M} \in \text{Forb}_{\text{he}}(\mathcal{F})$ containing \mathbf{A} as a substructure such that each partial automorphism of \mathbf{A} extends to an automorphism of \mathbf{M} , then there exists a finite structure $\mathbf{B} \in \text{Forb}_{\text{he}}(\mathcal{F})$ which is an irreducible structure faithful coherent EPPA-witness for \mathbf{A} .*

Hubička and Nešetřil [HN19] gave a structural condition for a class to be Ramsey. It turns out that, in papers studying Ramsey expansions of various classes using their theorem, EPPA is sometimes an easy corollary of one of the intermediate steps, see e.g. [ABWH⁺17b, ABWH⁺17a, ABWH⁺21, Kon19]. In this paper, we make this link explicit by proving a theorem on EPPA whose statement is very similar to [HN19, Theorem 2.18]. For the definition of a locally finite automorphism-preserving subclass, see Section 4.11.

Theorem 4.1.6. *Let L be a language consisting of relations and unary functions equipped with a permutation group Γ_L and let \mathcal{E} be a class of finite Γ_L -structures which has EPPA. Let \mathcal{K} be a hereditary locally finite automorphism-preserving subclass of \mathcal{E} with the strong amalgamation property which consists of irreducible structures. Then \mathcal{K} has EPPA.*

Moreover, if EPPA-witnesses in \mathcal{E} can be chosen to be coherent then EPPA-witnesses in \mathcal{K} can be chosen to be coherent, too.

4.1.3 Applications of our results

When proving EPPA for (some of) the antipodal classes of metrically homogeneous graphs in [ABWH⁺17b], an additional ad hoc layer was added on top of an application of the Herwig–Lascar theorem to ensure that edges of length δ form a matching [ABWH⁺17b, Theorem 7.7]. Another ad hoc layer was needed in the same paper for the bipartite classes [ABWH⁺17b, Theorem 6.13]. These ad hoc constructions can be avoided using the main theorem of this paper, adding unary functions to represent edges of length δ for the antipodal classes, or adding two unary predicates which can be swapped by Γ_L to describe the partition for the bipartite classes.

When proving EPPA for the antipodal classes of odd diameter and the bipartite antipodal classes of even diameter of metrically homogeneous graphs in [Kon20], the full strength of Theorem 4.1.6 (from an early draft of this paper) was used: One needed unary functions to represent that edges of length δ form a matching, control over substructures to ensure that no non-metric cycles are created, and language permutations to generalise a construction from [EHKN20].

Section 4.12 of this paper is devoted to applications. In particular, we outline how to further strengthen our results to languages with constants or certain non-unary functions (see Theorems 4.12.4 and 4.12.5), and we prove that a class connected to Hrushovski’s predimension construction has EPPA (see Theorem 4.12.7). The two latter proofs use a general method for dealing with higher-arity functions by chaining several applications of Theorem 4.1.2 on top of each other using language permutations.

We are confident that there are many more applications of the main theorems of this paper to be discovered.

4.1.4 EPPA and Ramsey

The results and techniques of this paper are motivated by recent developments of the structural Ramsey theory, particularly the efforts to characterise Ramsey classes of finite structures. As this paper demonstrates, many techniques and proof strategies from structural Ramsey theory may serve as a motivation for results about EPPA classes. We were inspired by the scheme of proofs of corresponding Ramsey results in [HN19], by the construction of clique faithful EPPA-witnesses for relational structures given by Hodkinson and Otto [HO03], by the treatment of unary function in [EHN21], and by the recent proofs of EPPA for metric spaces [HKN19a] and for two-graphs [EHKN20].

In each section, we fix a Γ_L -structure \mathbf{A} and give an explicit construction of a Γ_L -structure \mathbf{B} and an embedding $\mathbf{A} \rightarrow \mathbf{B}$. Then, given a partial automorphism of \mathbf{A} , we show how to construct an automorphism of \mathbf{B} extending it, that is, we prove that \mathbf{B} is an EPPA-witness for \mathbf{A} . Finally, we prove that \mathbf{B} has the given special properties (e.g. irreducible structure faithfulness, or control over small substructures) and that the extension is coherent. Usually, the constructions are the interesting part and the proofs are just verification that a function is an automorphism and that it composes correctly.

While all this may be surprising on the first glance and it is one of the novelties of this paper, we want to stress at this point some of the main differences between EPPA and Ramsey. (Further open problems will be in Section 4.13.)

Both EPPA and the Ramsey property imply the amalgamation property (see Observation 4.2.3 and [Neš05]) and have strong consequences for the Fraïssé limits. Nonetheless, not every amalgamation class has EPPA or the Ramsey property. While there is a meaningful conjecture motivating the classification program of Ramsey classes (see [HN19, BPT11]), for the classification of EPPA classes this is not yet the case.

The classification programme for EPPA classes was initiated in [EHN19, EHN21] by giving examples of classes with a non-trivial EPPA expansion. (See also the survey by the first author [Hub20b].) There exist many classes which have a non-trivial Ramsey expansion but fail to have a non-trivial EPPA expansion. Examples include the class of all finite linear orders or the class of all finite finite partially ordered sets. On the other hand, to the author’s best knowledge, whenever a Ramsey expansion of an EPPA class is known, the expansion only adds a “small amount of information” (compared to what is promised for ω -categorical structures by [Kec12, Theorem 4.5]).


The correspondence between the structural conditions for EPPA and Ramsey classes then motivates the following conjecture.

Conjecture 4.1.7. *Every strong amalgamation class with EPPA has a precompact Ramsey expansion.*

(See Section 4.2.3 for a definition of strong amalgamation and for example [NVT15] for a definition of a precompact expansion.) Note that this conjecture implies every ω -categorical structure with EPPA having a precompact Ramsey expansion. Classes known to have EPPA where it is not known if they have a precompact Ramsey expansion include the class of all finite groups [Sin17, PS18] and the class of all finite skew-symmetric bilinear forms¹. More open problems are listed in Section 4.13.

It is worth to mention a result of Jahel and Tsankov [JT22] who prove that for a large number of classes, EPPA implies the ordering property (which is closely related to the Ramsey property, see [KPT05]). In particular, this implies that while for Ramsey classes, there exists an ordering of Fraïssé limit which is compatible with the group of automorphisms, for EPPA classes satisfying the conditions of [JT22] such a global ordering cannot be definable. This in fact may be one of the main dividing lines.

Based on all this information and an analogous scheme in the Ramsey context [HN19], this may be schematically depicted as follows.



This paper is organised as follows: In Section 4.2, we give all the necessary notions and definitions. In Section 4.3, which is supposed to serve as a warm-up, we

¹David M. Evans, personal communication. See also [CH03].

give a new proof of (a coherent strengthening of) Hrushovski’s theorem [Hru92]. Then, in Sections 4.4 and 4.5, we show that this new construction generalises naturally to relational Γ_L -structures. In Section 4.6, we add a new layer which allows the language to also contain unary functions. In Section 4.7, we combine this with techniques introduced earlier [HO03, EHN21] to obtain irreducible structure faithfulness, and in Section 4.8, we once again use a similar construction to deal with forbidden homomorphic images, which allows us to prove the main theorems of this paper in Sections 4.9, 4.10 and 4.11. Finally, in Section 4.12, we apply our results and prove EPPA for the class of k -orientations with d -closures, thereby confirming the first part of [EHN19, Conjecture 7.5]. We also prove Corollary 4.1.4, illustrate the usage of Theorems 4.1.1 and 4.1.6 on the example of integer-valued metric spaces with no large subspaces, where all vertices are in distance 1 from each other and prove EPPA for languages with constants or certain classes with non-unary functions.

4.2 Background and notation

We find it convenient to work with model-theoretic structures generalised in two ways: We equip the language with a permutation group (giving a more systematic treatment to the concept of *permomorphisms* introduced by Herwig [Her98]) and consider functions to the powerset (a further generalisation of [EHN21]). This is motivated by applications, see Section 4.12.5.

Let $L = L_{\mathcal{R}} \cup L_{\mathcal{F}}$ be a language consisting of relation symbols $R \in L_{\mathcal{R}}$ and function symbols $F \in L_{\mathcal{F}}$ each having its *arity* denoted by $a(R) \geq 1$ for relations and $a(F) \geq 0$ for functions.

Let Γ_L be a permutation group on L which preserves types and arities of all symbols. We say that L is a *language equipped with a permutation group* Γ_L . Observe that when Γ_L is trivial and the ranges of all functions consist of singletons, one obtains the usual notion of model-theoretic language (and structures). All results and constructions in this paper presented on Γ_L -structures thus hold also for standard L -structures. By this we mean that given a class of standard L -structures, one can treat them as Γ_L -structures with Γ_L trivial, use the results of this paper and then, perhaps after some straightforward adjustments, obtain the same results for the original class (see Observations 4.2.1 and 4.2.4).

Denote by 2^A the set of all subsets of A . A Γ_L -*structure* \mathbf{A} is a structure with *vertex set* A , functions $F_{\mathbf{A}}: A^{a(F)} \rightarrow 2^A$ for every $F \in L_{\mathcal{F}}$ and relations $R_{\mathbf{A}} \subseteq A^{a(R)}$ for every $R \in L_{\mathcal{R}}$. Notice that the domain of a function is a tuple while the range is a set, the reason for this is that it allows to explicitly represent algebraic closures by functions. If the set A is finite, we call \mathbf{A} a *finite structure*. We consider only structures with finitely or countably infinitely many vertices. If $L_{\mathcal{F}} = \emptyset$, we call L a *relational language* and say that a Γ_L -structure is a *relational Γ_L -structure*. A function F such that $a(F) = 1$ is a *unary function*.

In this paper, the language and its permutation group are often fixed and understood from the context (and they are in most cases denoted by L and Γ_L respectively), we also only consider unary functions.

4.2.1 Maps between Γ_L -structures

A *homomorphism* $f: \mathbf{A} \rightarrow \mathbf{B}$ is a pair $f = (f_L, f_A)$ where $f_L \in \Gamma_L$ and f_A is a mapping $A \rightarrow B$ such that for every $R \in L_{\mathcal{R}}$ and $F \in L_{\mathcal{F}}$ we have:

- (a) $(x_1, \dots, x_{a(R)}) \in R_{\mathbf{A}} \implies (f_A(x_1), \dots, f_A(x_{a(R)})) \in f_L(R)_{\mathbf{B}}$, and
- (b) $f_A(F_{\mathbf{A}}(x_1, \dots, x_{a(F)})) \subseteq f_L(F)_{\mathbf{B}}(f_A(x_1), \dots, f_A(x_{a(F)}))$.

If $f = (f_L, f_A): \mathbf{A} \rightarrow \mathbf{B}$ and $g = (g_L, g_B): \mathbf{B} \rightarrow \mathbf{C}$ are homomorphisms, we denote by $gf = g \circ f = (g_L \circ f_L, g_B \circ f_A)$ the homomorphism $\mathbf{A} \rightarrow \mathbf{C}$ obtained by their composition. (It is straightforward to check that the composition is indeed a homomorphism $\mathbf{A} \rightarrow \mathbf{C}$.)

If f_A is injective then f is called a *monomorphism*. A monomorphism $f = (f_L, f_A)$ is an *embedding* if for every $R \in L_{\mathcal{R}}$ and $F \in L_{\mathcal{F}}$:

- (a) $(x_1, \dots, x_{a(R)}) \in R_{\mathbf{A}} \iff (f_A(x_1), \dots, f_A(x_{a(R)})) \in f_L(R)_{\mathbf{B}}$, and
- (b) $f_A(F_{\mathbf{A}}(x_1, \dots, x_{a(F)})) = f_L(F)_{\mathbf{B}}(f_A(x_1), \dots, f_A(x_{a(F)}))$.

If f is an embedding where f_A is one-to-one then f is an *isomorphism*. An isomorphism from a structure to itself is called an *automorphism*. If f_A is an inclusion and f_L is the identity then \mathbf{A} is a *substructure* of \mathbf{B} and we may write $\mathbf{A} \subseteq \mathbf{B}$ to denote this fact.

Given a Γ_L -structure \mathbf{B} and $A \subseteq B$, the *closure of A in \mathbf{B}* , denoted by $\text{Cl}_{\mathbf{B}}(A)$, is the smallest substructure of \mathbf{B} containing A . For $x \in B$, we will also write $\text{Cl}_{\mathbf{B}}(x)$ for $\text{Cl}_{\mathbf{B}}(\{x\})$ and for a tuple $\bar{x} = (x_1, \dots, x_n) \in B^n$ we will write $\text{Cl}_{\mathbf{B}}(\bar{x})$ for $\text{Cl}_{\mathbf{B}}(\{x_1, \dots, x_n\})$.

Let \mathbf{A} , \mathbf{B} , \mathbf{C} and \mathbf{C}' be Γ_L -structures such that $\mathbf{C} \subseteq \mathbf{A}$ and $\mathbf{C}' \subseteq \mathbf{B}$. If $f: \mathbf{C} \rightarrow \mathbf{C}'$ is an isomorphism, we may also call it a *partial isomorphism* between \mathbf{A} and \mathbf{B} (note that f also includes a permutation $f_L \in \Gamma_L$).

Let $f = (f_L, f_A): \mathbf{A} \rightarrow \mathbf{B}$ be a homomorphism. For brevity, we may write $f(x)$ for $f_A(x)$ in the context where $x \in A$, and $f(S)$ for $f_L(S)$ where $S \in L$. By $\text{Dom}(f)$ and $\text{Range}(f)$ we will always mean $\text{Dom}(f_A)$ and $\text{Range}(f_A)$ respectively. If $\bar{x} = (x_1, \dots, x_n) \in A^n$ then by $f(\bar{x}) = f_A(\bar{x})$ we mean the tuple $(f_A(x_1), \dots, f_A(x_n))$, and if $X \subseteq A^n$ then we put $f(X) = f_A(X) = \{f(\bar{x}) : \bar{x} \in X\}$.

Note that $f(R_{\mathbf{A}}) = f_A(R_{\mathbf{A}})$ is the image of a set of tuples, while $f(R)_{\mathbf{B}} = f_L(R)_{\mathbf{B}}$ is the realisation of the relation $f_L(R)$ in \mathbf{B} . These sets need not be equal in general (they will, however, be equal whenever f is an embedding).

If $f_L \in \Gamma_L$ and f_A is a function from A to some set X , we denote by $f(\mathbf{A})$ the *homomorphic image* of structure \mathbf{A} , that is, the Γ_L -structure with vertex set $f_A(A)$ such that for every $R \in L_{\mathcal{R}}$ and $F \in L_{\mathcal{F}}$ we have:

- (a) $R_{f(\mathbf{A})} = f_A(f_L^{-1}(R)_{\mathbf{A}})$, and
- (b) for every $\bar{x} \in f_A(A)^{a(F)}$ it holds that

$$F_{f(\mathbf{A})}(\bar{x}) = \bigcup_{\bar{y} \in A^{a(F)}, \bar{x} = f(\bar{y})} f_A(f_L^{-1}(F)_{\mathbf{A}}(\bar{y})).$$

Note that f is a homomorphism $\mathbf{A} \rightarrow f(\mathbf{A})$ and moreover all relations and functions of $f(\mathbf{A})$ are minimal possible for f to be a homomorphism. Also observe that if f_A is injective, then f is an isomorphism $\mathbf{A} \rightarrow f(\mathbf{A})$.

Definition 4.2.1. Let L be a language equipped with a permutation group Γ_L . We define the *action of Γ_L on Γ_L -structures by relabelling*, such that for a Γ_L -structure \mathbf{A} and $g \in \Gamma_L$, we define $g\mathbf{A}$ as $(g, \text{id}_A)(\mathbf{A})$.

Observation 4.2.1. *Let L be a language equipped with a permutation group Γ_L . If L is finite or $\Gamma_L = \{\text{id}_L\}$ then every finite Γ_L -structure lies in a finite orbit of the action of Γ_L by relabelling.*

Proof. If L is finite, then there are only finitely many Γ_L -structures on any given finite set A and the action of Γ_L by relabelling preserves the vertex set. If $\Gamma_L = \{\text{id}_L\}$, then the action is trivial and every orbit is a singleton. \square

4.2.2 Γ_L -structures as standard model-theoretic structures

Whenever there exists a structure \mathbf{L} such that $\text{Aut}(\mathbf{L}) = \Gamma_L$ (e.g. when L is finite or more generally when Γ_L is a closed subgroup of $\text{Sym}(\mathbb{N})$), there is a functorial correspondence between Γ_L -structures and structures in a bigger language with no permutation group. This makes it possible to extend many theorems about classical structures to (certain) Γ_L -structures without having to re-prove them.

Definition 4.2.2. Let L be a language equipped with a permutation group Γ_L . Let X be the set of all symbols of L which are not fixed by Γ_L and assume that there is a language L_0 disjoint from L and an L_0 -structure \mathbf{X} such that $\text{Aut}(\mathbf{X})$ is precisely the action of Γ_L on X . Let L° be the language defined as follows:

1. For every symbol from $L \setminus X$, we put the same symbol with the same arity into L° .
2. For every symbol from L_0 , we put the same symbol with the same arity into L° .
3. For every n such that X contains a relation symbol of arity n , we put an $(n + 1)$ -ary relation symbol R^n into L° (without loss of generality $R^n \notin L \cup L_0$).
4. For every n such that X contains a function symbol of arity n , we put an $(n + 1)$ -ary function symbol F^n into L° (without loss of generality $F^n \notin L \cup L_0$).
5. There is a constant symbol $c \in L^\circ$ (without loss of generality $c \notin L \cup L_0$).

Given a Γ_L -structure \mathbf{A} , we define an L° -structure \mathbf{A}° as follows:

1. The vertex set of \mathbf{A}° is the disjoint union $A \cup X$ (without loss of generality we can assume that $A \cap X = \emptyset$).
2. $c_{\mathbf{A}^\circ} = X$.




Figure 4.1: An amalgamation of \mathbf{B}_1 and \mathbf{B}_2 over \mathbf{A} .

3. The substructure of \mathbf{A}° induced on X is isomorphic to \mathbf{X} (in particular, there are no relations or functions from L).
4. For every symbol $S \in L \setminus X$ we have that $S_{\mathbf{A}^\circ} = S_{\mathbf{A}}$.
5. For every n -ary relation symbol $S \in X$ and every tuple $(x_1, \dots, x_n) \in A^n$ it holds that $(x_1, \dots, x_n, S) \in R_{\mathbf{A}^\circ}^n$ if and only if $(x_1, \dots, x_n) \in S_{\mathbf{A}}$.
6. For every n -ary function symbol $S \in X$ and every tuple $(x_1, \dots, x_n) \in A^n$ it holds that $(x_1, \dots, x_n, S) \in \text{Dom}(F_{\mathbf{A}^\circ}^n)$ if and only if $(x_1, \dots, x_n) \in \text{Dom}(S_{\mathbf{A}})$, and in that case $F_{\mathbf{A}^\circ}^n(x_1, \dots, x_n, S) = S_{\mathbf{A}}(x_1, \dots, x_n)$.

Fact 4.2.2. *In the setting of Definition 4.2.2, $f = (f_L, f_A)$ is an embedding of Γ_L -structures $\mathbf{A} \rightarrow \mathbf{B}$ if and only if $f_L \upharpoonright_X \cup f_A$ is an embedding of L° structures $\mathbf{A}^\circ \rightarrow \mathbf{B}^\circ$.*

This implies that the map $\mathbf{A} \mapsto \mathbf{A}^\circ$ is an isomorphism of categories. Note that whenever Γ_L is finite, we have that \mathbf{A} is finite if and only if \mathbf{A}° is. This construction still gives structures where the images of functions need not consist of singletons. In order to deal with this, one can replace functions by relations and consider only algebraically closed substructures as is standard in the area, see for example [Eva].

4.2.3 Amalgamation classes

Let \mathbf{A} , \mathbf{B}_1 , and \mathbf{B}_2 be Γ_L -structures, and let $\alpha_1: \mathbf{A} \rightarrow \mathbf{B}_1$, $\alpha_2: \mathbf{A} \rightarrow \mathbf{B}_2$ be embeddings. A structure \mathbf{C} with embeddings $\beta_1: \mathbf{B}_1 \rightarrow \mathbf{C}$ and $\beta_2: \mathbf{B}_2 \rightarrow \mathbf{C}$ such that $\beta_1 \circ \alpha_1 = \beta_2 \circ \alpha_2$ (remember that this must also hold for the language part of α_i 's and β_i 's) is called an *amalgamation* of \mathbf{B}_1 and \mathbf{B}_2 over \mathbf{A} with respect to α_1 and α_2 , see Figure 4.1. We will often call \mathbf{C} simply an *amalgamation* of \mathbf{B}_1 and \mathbf{B}_2 over \mathbf{A} (in most cases α_1 and α_2 can be chosen to be inclusion embeddings).

We say that the amalgamation is *strong* if it holds that $\beta_1(x_1) = \beta_2(x_2)$ if and only if $x_1 \in \alpha_1(A)$ and $x_2 \in \alpha_2(A)$. Strong amalgamation is *free* if $C = \beta_1(B_1) \cup \beta_2(B_2)$, and whenever a tuple \bar{x} of vertices of \mathbf{C} contains vertices of both $\beta_1(B_1 \setminus \alpha_1(A))$ and $\beta_2(B_2 \setminus \alpha_2(A))$, then \bar{x} is in no relation of \mathbf{C} and also for every function $F \in L$ with $a(F) = |\bar{x}|$ it holds that $F_{\mathbf{C}}(\bar{x}) = \emptyset$.

Definition 4.2.3. An *amalgamation class* is a class \mathcal{K} of finite Γ_L -structures which is closed for isomorphisms and satisfies the following three conditions:

1. *Hereditary property*: For every $\mathbf{B} \in \mathcal{K}$ and every structure \mathbf{A} with an embedding $f: \mathbf{A} \rightarrow \mathbf{B}$ we have $\mathbf{A} \in \mathcal{K}$;
2. *Joint embedding property*: For every $\mathbf{A}, \mathbf{B} \in \mathcal{K}$ there exists $\mathbf{C} \in \mathcal{K}$ with an embeddings $f: \mathbf{A} \rightarrow \mathbf{C}$ and $g: \mathbf{B} \rightarrow \mathbf{C}$;
3. *Amalgamation property*: For $\mathbf{A}, \mathbf{B}_1, \mathbf{B}_2 \in \mathcal{K}$ and embeddings $\alpha_1: \mathbf{A} \rightarrow \mathbf{B}_1$, $\alpha_2: \mathbf{A} \rightarrow \mathbf{B}_2$, there is $\mathbf{C} \in \mathcal{K}$ which is an amalgamation of \mathbf{B}_1 and \mathbf{B}_2 over \mathbf{A} with respect to α_1 and α_2 .

If the \mathbf{C} in the amalgamation property can always be chosen to be a strong amalgamation then \mathcal{K} is a *strong amalgamation class*, if it can always be chosen to be the free amalgamation then \mathcal{K} is a *free amalgamation class*.

By the Fraïssé theorem [Fra53], relational amalgamation classes in a countable language with trivial Γ_L containing only countably many members up to isomorphism correspond to countable homogeneous structures. By Section 4.2.2, one can extend this to various languages equipped with a permutation group using a variant of the Fraïssé theorem for languages with functions or for strong substructures, see for example [Eva].

Generalising the notion of a graph clique, we say that a structure is *irreducible* if it is not a free amalgamation of its proper substructures. A homomorphism $f: \mathbf{A} \rightarrow \mathbf{B}$ is a *homomorphism-embedding* if the restriction $f \upharpoonright_{\mathbf{C}}$ is an embedding whenever \mathbf{C} is an irreducible substructure of \mathbf{A} . Given a family \mathcal{F} of Γ_L -structures, we denote by $\text{Forb}_{\text{he}}(\mathcal{F})$ the class of all finite or countably infinite Γ_L -structures \mathbf{A} such that there is no $\mathbf{F} \in \mathcal{F}$ with a homomorphism-embedding $\mathbf{F} \rightarrow \mathbf{A}$.

4.2.4 EPPA for Γ_L -structures

A *partial automorphism* of a Γ_L -structure \mathbf{A} is a partial isomorphism between \mathbf{A} and \mathbf{A} . Let \mathbf{A} and \mathbf{B} be finite Γ_L -structures. We say that \mathbf{B} is an *EPPA-witness for \mathbf{A}* if there is an embedding $\psi: \mathbf{A} \rightarrow \mathbf{B}$ and every partial automorphism of $\psi(\mathbf{A})$ extends to an automorphism of \mathbf{B} , that is, for every partial automorphism φ of $\psi(\mathbf{A})$ there is an automorphism $\tilde{\varphi}: \mathbf{B} \rightarrow \mathbf{B}$ such that $\varphi \subseteq \tilde{\varphi}$.

We say that a class of finite Γ_L -structures \mathcal{K} has the *extension property for partial automorphisms* (shortly *EPPA*, sometimes called the *Hrushovski property*) if for every $\mathbf{A} \in \mathcal{K}$ there is $\mathbf{B} \in \mathcal{K}$ which is an EPPA-witness for \mathbf{A} . Such a structure \mathbf{B} is *irreducible structure faithful* (with respect to $\psi(\mathbf{A})$) if it has the property that for every irreducible substructure \mathbf{C} of \mathbf{B} there exists an automorphism g of \mathbf{B} such that $g(C) \subseteq \psi(A)$.

Note that the classes which we are interested in are closed under taking isomorphisms, and hence if there is an EPPA-witness \mathbf{B} for \mathbf{A} in \mathcal{K} , then there is also an EPPA-witness $\mathbf{B}' \in \mathcal{K}$ such that ψ is just the inclusion $\mathbf{A} \subseteq \mathbf{B}'$. To simplify the arguments, we will often ignore this subtle technicality.

Homomorphism-embeddings were introduced in [HN19] and irreducible structure faithfulness was introduced in [EHN21] as a generalisation of clique faithfulness of Hodkinson and Otto [HO03]. The following observation provides a link to the study of homogeneous structures.

Observation 4.2.3 ([Her98]). *Every hereditary isomorphism-closed class of finite Γ_L -structures which has EPPA and the joint embedding property (see Definition 4.2.3) is an amalgamation class.*

Proof. Let \mathcal{K} be such a class and let $\mathbf{A}, \mathbf{B}_1, \mathbf{B}_2 \in \mathcal{K}$, $\alpha_1: \mathbf{A} \rightarrow \mathbf{B}_1$, $\alpha_2: \mathbf{A} \rightarrow \mathbf{B}_2$ be as in Definition 4.2.3. Let \mathbf{B} be the joint embedding of \mathbf{B}_1 and \mathbf{B}_2 (that is, we have embeddings $\beta'_1: \mathbf{B}_1 \rightarrow \mathbf{B}$ and $\beta'_2: \mathbf{B}_2 \rightarrow \mathbf{B}$) and let \mathbf{C} be an EPPA-witness for \mathbf{B} . Without loss of generality, we can assume that $\mathbf{B} \subseteq \mathbf{C}$.

Let φ be a partial automorphism of \mathbf{B} sending $\beta'_1(\alpha_1(\mathbf{A}))$ to $\beta'_2(\alpha_2(\mathbf{A}))$ and let θ be its extension to an automorphism of \mathbf{C} . Finally, put $\beta_1 = \theta \circ \beta'_1$ and $\beta_2 = \beta'_2$. It is easy to check that β_1 and β_2 certify that \mathbf{C} is an amalgamation of \mathbf{B}_1 and \mathbf{B}_2 over \mathbf{A} with respect to α_1 and α_2 . \square

We believe that the results of this paper may often be used in a more specialised setting such as for standard model-theoretic L -structures etc. In order to facilitate that, we state the following simple observation which allows translating the main theorems into this more specialised setting.

Observation 4.2.4. *Let \mathbf{A} and \mathbf{B} be finite Γ_L -structures such that \mathbf{B} is an irreducible structure faithful EPPA-witness for \mathbf{A} . If, in \mathbf{A} , the range of every function consists of singletons then this also holds in \mathbf{B} .*

Proof. Suppose for a contradiction that there is a function $F \in L$ and a tuple $\bar{x} \in B^{a(F)}$ such that $F_{\mathbf{B}}(\bar{x}) = X$ and $|X| > 1$. Since $\text{Cl}_{\mathbf{B}}(\bar{x})$ is irreducible, by irreducible structure faithfulness there is an automorphism $f: \mathbf{B} \rightarrow \mathbf{B}$ sending $\text{Cl}_{\mathbf{B}}(\bar{x})$ to A . In particular, $f(\bar{x}) \subseteq A$ and $f(X) \subseteq A$. As f is an automorphism, we know that $f_L(F)_{\mathbf{A}}(f_B(\bar{x})) = f_B(X)$ implying that $|f_B(X)| \leq 1$ which is a contradiction. \square

4.2.5 EPPA and automorphism groups

As was mentioned in the introduction, the significance of EPPA comes from the fact that, while being a property of a class of finite structures, it is closely connected with topological properties of the automorphism group of an infinite structure, namely the Fraïssé limit of the class. We do not aim for this section to be self-contained nor complete (and refer the reader for example to [Sin17]), we only outline some of these connections and discuss how they extend to Γ_L -structures. In contrast to the rest of this paper, in this section we are mostly going to be interested in (countably) infinite structures. For simplicity, we will assume that Γ_L is finite (so that Section 4.2.2 can be fully applied), the case of infinite Γ_L can be more complicated and deserves a study on its own.

Let \mathbf{M} be a countably infinite Γ_L -structure. Assume without loss of generality that the vertex set of \mathbf{M} are the natural numbers and put $G = \text{Aut}(\mathbf{M})$. Remember that members of G are pairs $f = (f_L, f_M)$ with $f_L \in \Gamma_L$. This means that G can be understood as a subgroup of $\Gamma_L \times \text{Sym}(\mathbb{N})$ and as such inherits the topology of this product, where $\text{Sym}(\mathbb{N})$ is equipped with the pointwise-convergence topology and Γ_L , being finite, is equipped with the discrete topology. Note that this precisely corresponds to what one gets by using Section 4.2.2 and recalling the standard definitions.

Let \mathbf{M} be a homogeneous Γ_L -structure. We say that \mathbf{M} is *locally finite* if for every finite $X \subseteq M$ it holds that $\text{Cl}_{\mathbf{M}}(X)$ is also finite. By $\text{Age}(\mathbf{M})$ we denote the class of all finite Γ_L -structures which embed into \mathbf{M} . The following theorem has been proved by Kechris and Rosendal [KR07] for classical structures and by Section 4.2.2 extends naturally to Γ_L -structures:

Theorem 4.2.5 ([KR07]). *Let L be a language equipped with a finite permutation group Γ_L and let \mathbf{M} be a countable locally finite homogeneous Γ_L -structure. Then $\text{Age}(\mathbf{M})$ has EPPA if and only if $\text{Aut}(\mathbf{M})$ can be written as the closure of a chain of compact subgroups. Moreover, if $\text{Age}(\mathbf{M})$ has EPPA, then $\text{Aut}(\mathbf{M})$ is amenable.*

Definition 4.2.4 ([HHLS93, KR07]). Let L be a language equipped with a finite permutation group Γ_L , let \mathbf{M} be a countable locally finite homogeneous Γ_L -structure and let $n \geq 1$ be an integer. We say that \mathbf{M} has *n -generic automorphisms* if G has a comeagre orbit on G^n in its action by diagonal conjugation. We say that \mathbf{M} has *ample generics* if it has n -generic automorphisms for every $n \geq 1$.

Here, the action by diagonal conjugation is defined by

$$g \cdot (h_1, \dots, h_n) = (gh_1g^{-1}, \dots, gh_ng^{-1}).$$

The existence of ample generics has many consequences for the automorphism group such as the small index property. From the point of view of this paper, ample generics are relevant, because EPPA is very often a key ingredient in proving them. We will outline this connection in the rest of this section.

Definition 4.2.5. Let L be a language equipped with a permutation group Γ_L , let \mathcal{C} be a class of finite Γ_L -structures and let $n \geq 1$ be an integer. An *n -system over \mathcal{C}* is a tuple $(\mathbf{A}, p_1, \dots, p_n)$, where $\mathbf{A} \in \mathcal{C}$ and p_1, \dots, p_n are partial automorphisms of \mathbf{A} . We denote by \mathcal{C}^n the class of all n -systems over \mathcal{C} .

If $P = (\mathbf{A}, p_1, \dots, p_n)$ and $Q = (\mathbf{B}, q_1, \dots, q_n)$ are both n -systems over \mathcal{C} and $f: \mathbf{A} \rightarrow \mathbf{B}$ is an embedding of Γ_L -structures, we say that f is an *embedding of n -systems $P \rightarrow Q$* if for every $1 \leq i \leq n$ it holds that $f \circ p_i \subseteq q_i \circ f$ (in particular, $f(\text{Dom}(p_i)) \subseteq \text{Dom}(q_i)$ and $f(\text{Range}(p_i)) \subseteq \text{Range}(q_i)$).

Definition 4.2.6. Let L be a language equipped with a permutation group Γ_L , let \mathcal{C} be a class of finite Γ_L -structures and let $n \geq 1$ be an integer. We say that \mathcal{C}^n has the *joint embedding property* if for every $P, Q \in \mathcal{C}^n$ there exists $S \in \mathcal{C}^n$ with embeddings of n -systems $f: P \rightarrow S$ and $g: Q \rightarrow S$. We say that \mathcal{C}^n has the *weak amalgamation property* if for every $T \in \mathcal{C}^n$ there exists $\hat{T} \in \mathcal{C}^n$ and an embedding of n -systems $\iota: T \rightarrow \hat{T}$ such that for every pair of n -systems $P, Q \in \mathcal{C}^n$ and embeddings of n -systems $\alpha_1: \hat{T} \rightarrow P$ and $\alpha_2: \hat{T} \rightarrow Q$ there exists $S \in \mathcal{C}^n$ with embeddings on n -systems $\beta_1: P \rightarrow S$ and $\beta_2: Q \rightarrow S$ such that $\beta_1\alpha_1\iota = \beta_2\alpha_2\iota$.

We only state the following theorem for $\Gamma_L = \{\text{id}\}$, as \mathcal{C}^n does not have the joint embedding property for any n if \mathcal{C} is a class of finite Γ_L -structures with $2 \leq |\Gamma_L| < \infty$ (see Example 4.2.2).

Theorem 4.2.6 ([KR07]). *Let L be a language, let \mathbf{M} be a countable locally finite homogeneous L -structure, put $\mathcal{C} = \text{Age}(\mathbf{M})$ and fix $n \geq 1$. Then \mathbf{M} has n -generic automorphisms if and only if \mathcal{C}^n has the joint embedding property and the weak amalgamation property.*

In order to explain the connection between EPPA and ample generics, we need one more definition

Definition 4.2.7. Let L be a language equipped with a permutation group Γ_L and let \mathcal{C} be a class of finite Γ_L -structures. We say that \mathcal{C} has the *amalgamation property with automorphisms* (abbreviated as *APA*) if for every $\mathbf{A}, \mathbf{B}_1, \mathbf{B}_2 \in \mathcal{C}$ and embeddings $\alpha_1: \mathbf{A} \rightarrow \mathbf{B}_1, \alpha_2: \mathbf{A} \rightarrow \mathbf{B}_2$ there exists $\mathbf{C} \in \mathcal{C}$ with embeddings $\beta_1: \mathbf{B}_1 \rightarrow \mathbf{C}$ and $\beta_2: \mathbf{B}_2 \rightarrow \mathbf{C}$ such that $\beta_1 \circ \alpha_1 = \beta_2 \circ \alpha_2$ (i.e. \mathbf{C} is an amalgamation of \mathbf{B}_1 and \mathbf{B}_2 over \mathbf{A} with respect to α_1 and α_2) and moreover whenever we have $f \in \text{Aut}(\mathbf{B}_1)$ and $g \in \text{Aut}(\mathbf{B}_2)$ such that $f(\alpha_1(A)) = \alpha_1(A), g(\alpha_2(A)) = \alpha_2(A)$ and for every $a \in A$ it holds that $\alpha_1^{-1}(f(\alpha_1(a))) = \alpha_2^{-1}(g(\alpha_2(a)))$ (that is, f and g agree on the copy of \mathbf{A} we are amalgamating over), then there is $h \in \text{Aut}(\mathbf{C})$ which extends $\beta_1 f \beta_1^{-1} \cup \beta_2 g \beta_2^{-1}$. We call such \mathbf{C} with embeddings β_1 and β_2 an *APA-witness* for \mathbf{B}_1 and \mathbf{B}_2 over \mathbf{A} with respect to α_1 and α_2

Proposition 4.2.7 ([KR07]). *Let L be a language equipped with a permutation group Γ_L and let \mathcal{C} be a class of finite Γ_L -structures. If \mathcal{C} has EPPA and APA then \mathcal{C}^n has the weak amalgamation property for every $n \geq 1$.*

Proof. Fix $n \geq 1$. If $S = (\mathbf{S}, s_1, \dots, s_n) \in \mathcal{C}^n$ is an n -system, we denote by $\hat{S} = (\hat{\mathbf{S}}, \hat{s}_1, \dots, \hat{s}_n) \in \mathcal{C}^n$ the n -system where $\hat{\mathbf{S}}$ is an EPPA-witness for \mathbf{S} (with respect to the inclusion embedding) and for every $1 \leq i \leq n$ it holds that \hat{s}_i is an automorphism of $\hat{\mathbf{S}}$ extending s_i .

We now prove that \mathcal{C}^n has the weak amalgamation property. Towards that, fix some $T = (\mathbf{T}, t_1, \dots, t_n) \in \mathcal{C}^n$. Let $P = (\mathbf{P}, p_1, \dots, p_n), Q = (\mathbf{Q}, q_1, \dots, q_n) \in \mathcal{C}^n$ be arbitrary n -systems with embeddings $\alpha_1: \hat{T} \rightarrow P$ and $\alpha_2: \hat{T} \rightarrow Q$.

Use APA for \mathcal{C} to get $\mathbf{S} \in \mathcal{C}$ and embeddings $\beta_1: \hat{\mathbf{P}} \rightarrow \mathbf{S}$ and $\beta_2: \hat{\mathbf{Q}} \rightarrow \mathbf{S}$ such that \mathbf{S} with β_1 and β_2 form an APA-witness for $\hat{\mathbf{P}}$ and $\hat{\mathbf{Q}}$ over $\hat{\mathbf{T}}$ with respect to α_1 and α_2 . Let $S = (\mathbf{S}, s_1, \dots, s_n) \in \mathcal{C}^n$ be some n -system such that for every $1 \leq i \leq n$ we have that s_i extends $\beta_1 \hat{p}_i \beta_1^{-1} \cup \beta_2 \hat{q}_i \beta_2^{-1}$. It is straightforward to verify that S is the desired n -system witnessing the weak amalgamation property for P, Q and T . \square

Example 4.2.1. Consider the class \mathcal{C} of all finite graphs. By a theorem of Hrushovski [Hru92] (or by Section 4.3) we know that \mathcal{C} has EPPA. APA for \mathcal{C} is an easy exercise (in general, APA for free amalgamation classes is always true). Hence, by Proposition 4.2.7, \mathcal{C}^n has the weak amalgamation property for every $n \geq 1$. To prove ample generics for the countable random graph it thus remains to prove the joint embedding property for \mathcal{C}^n . However, it is again an easy exercise (simply take the disjoint union of the graphs and the partial automorphisms).

Example 4.2.2. Let L be a language consisting of two unary relations U and V , put $\Gamma_L = \text{Sym}(L)$ and let \mathcal{C} be the class of all finite Γ_L -structures where every vertex is in precisely one of the two unary relations. Clearly, \mathcal{C} can equivalently be seen as the class of all finite structures with one equivalence relation with two equivalence classes. Since \mathcal{C} is a free amalgamation class, Corollary 4.1.4 gives us

EPPA for \mathcal{C} , APA for \mathcal{C} is straightforward. Hence, by Proposition 4.2.7, \mathcal{C}^n has the weak amalgamation property for every $n \geq 1$.

However, \mathcal{C}^n fails to have the joint embedding property for any $n \geq 1$ and this is already visible on n -systems with the empty structure. Let \mathbf{E} be the Γ_L -structure with no vertices and assume that Γ_L is enumerated as $\{\text{id}, t\}$ where id is the identity and t is the transposition $U \leftrightarrow V$. Put $P = (\mathbf{E}, (\text{id}, \emptyset))$ and $Q = (\mathbf{E}, (t, \emptyset))$. Clearly, there is no 1-system which embeds both P and Q .

However, this kind of obstacle is the only reason why \mathcal{C}^n does not have the joint embedding property (in general for free amalgamation classes): One can define an equivalence relation \sim_n on \mathcal{C}^n for every n and for every pair of n -systems $P = (\mathbf{P}, (p_L^1, p_P^1), \dots, (p_L^n, p_P^n)) \in \mathcal{C}^n$ and $Q = (\mathbf{Q}, (q_L^1, q_Q^1), \dots, (q_L^n, q_Q^n)) \in \mathcal{C}^n$ by putting $P \sim_n Q$ if and only if there is $f \in \Gamma_L$ such that for every $1 \leq i \leq n$ we have that $f \circ p_L^i = q_L^i \circ f$. Then $P, Q \in \mathcal{C}^n$ have a joint embedding if and only if $P \sim_n Q$. In fact, it is then possible to prove a relativised version of Theorem 4.2.6 and obtain generic automorphisms for every equivalence class of \sim_n . For example, if \mathcal{C} is the class from this example, $n = 1$ and the language part is the identity, then the generic automorphism is just a pair of permutations of vertices of each unary such that both of them have no infinite cycles and infinitely many k -cycles for every finite $k \geq 1$.

4.2.6 Coherence of EPPA-witnesses

Siniora and Solecki [Sol09, SS19] strengthened the notion of EPPA in order to get a dense locally finite subgroup of the automorphism group of the corresponding Fraïssé limit. In order to state their definitions, we need to define how partial maps compose. Let L be a language equipped with a permutation group Γ_L , let \mathbf{A} , \mathbf{B} and \mathbf{C} be Γ_L -structures, let f be a partial isomorphism between \mathbf{A} and \mathbf{B} and let g be a partial isomorphism between \mathbf{B} and \mathbf{C} such that $\text{Dom}(g_B) = \text{Range}(f_A)$. We define their composition gf (also denoted by $g \circ f$) to be the partial isomorphism between \mathbf{A} and \mathbf{C} such that $(gf)_L = g_L f_L$, $(gf)_A(x)$ is defined if and only if $x \in \text{Dom}(f_A)$ and $f_A(x) \in \text{Dom}(g_B)$, and in this case we put $(gf)_A(x) = g_A(f_B(x))$.

Definition 4.2.8 (Coherent maps). Let L be a language equipped with a permutation group Γ_L , let \mathbf{A} be a Γ_L -structure and let \mathcal{P} be a family of partial automorphisms of \mathbf{A} . A triple $f, g, h \in \mathcal{P}$ is called a *coherent triple* if $\text{Range}(f_A) = \text{Dom}(g_A)$ and $h = gf$. A pair $f, g \in \mathcal{P}$ is called a *coherent pair* if there is $h \in \mathcal{P}$ such that f, g, h is a coherent triple.

Let \mathbf{A} and \mathbf{B} be Γ_L -structures, and let \mathcal{P} and \mathcal{Q} be families of partial automorphisms of \mathbf{A} and \mathbf{B} , respectively. A function $\varphi: \mathcal{P} \rightarrow \mathcal{Q}$ is said to be a *coherent map* if for each coherent triple (f, g, h) from \mathcal{P} , its image $(\varphi(f), \varphi(g), \varphi(h))$ in \mathcal{Q} is also coherent.

Definition 4.2.9 (Coherent EPPA). A class \mathcal{K} of finite Γ_L -structures is said to have *coherent EPPA* if \mathcal{K} has EPPA and moreover the extension of partial automorphisms is coherent. That is, for every $\mathbf{A} \in \mathcal{K}$, there exists $\mathbf{B} \in \mathcal{K}$ and an embedding $\psi: \mathbf{A} \rightarrow \mathbf{B}$ such that every partial automorphism f of $\psi(\mathbf{A})$ extends to some $\hat{f} \in \text{Aut}(\mathbf{B})$ with the property that the map $f \mapsto \hat{f}$ from partial automorphisms of $\psi(\mathbf{A})$ to $\text{Aut}(\mathbf{B})$ is coherent. We say that \mathbf{B} is a *coherent EPPA-witness* for \mathbf{A} .

The following easy proposition will be used several times. We include its proof to make this paper self-contained.

Proposition 4.2.8 (Lemma 2.1 in [SS19]). *Every finite set is a coherent EPPA-witness for itself. Consequently, the class of all finite sets has coherent EPPA.*

Explicitly, for every finite set A there is a map assigning to every partial injective function $\varphi: A \rightarrow A$ a permutation $\widehat{\varphi}$ of A such that $\varphi \subseteq \widehat{\varphi}$ and moreover for every coherent pair $\varphi_1, \varphi_2: A \rightarrow A$ it holds that $\widehat{\varphi_2}\widehat{\varphi_1} = \widehat{\varphi_2\varphi_1}$.

Proof. Fix a set A . Without loss of generality we can assume that $A = \{1, \dots, n\}$. Let φ be a partial automorphism of A , in other words, a partial injective function $A \rightarrow A$. We construct a permutation $\widehat{\varphi}: A \rightarrow A$ extending φ in the following way:

Put $X = A \setminus \text{Dom}(\varphi)$ and $Y = A \setminus \text{Range}(\varphi)$ and enumerate $X = \{x_1, \dots, x_k\}$ and $Y = \{y_1, \dots, y_k\}$ such that $x_1 < \dots < x_k$ and $y_1 < \dots < y_k$. Define $\widehat{\varphi}$ by

$$\widehat{\varphi}(x) = \begin{cases} \varphi(x) & \text{if } x \in \text{Dom}(\varphi) \\ y_i & \text{if } x = x_i. \end{cases}$$

It is obvious that $\widehat{\varphi}$ is a permutation of A which extends φ . Thus it only remains to prove coherence.

Consider $x \in A$. If $x \in \text{Dom}(\varphi_1)$, then we have $\varphi_1(x) \in \text{Dom}(\varphi_2)$ and hence $\widehat{\varphi_2}\widehat{\varphi_1}(x) = \widehat{\varphi_2}(\widehat{\varphi_1}(x))$. Put $X = A \setminus \text{Dom}(\varphi_1)$, $Y = A \setminus \text{Range}(\varphi_1) (= A \setminus \text{Dom}(\varphi_2))$ and $Z = A \setminus \text{Range}(\varphi_2)$ and again enumerate them in an ascending order. If $x = x_i$, we have $\widehat{\varphi_1}(x_i) = y_i$, $\widehat{\varphi_2}(y_i) = z_i$ and $\widehat{\varphi_2\varphi_1}(x_i) = z_i$, therefore indeed $\widehat{\varphi_2}\widehat{\varphi_1}(x_i) = \widehat{\varphi_2\varphi_1}(x_i)$. \square

When using this result, we will often simply say that we extend a partial permutation in an *order-preserving way* or *coherently*.

4.3 Warm-up: new proof of EPPA for graphs

We start with a simple proof of the theorem of Hrushovski [Hru92]. (Our proof is different from another simple proof given by Herwig and Lascar [HL00, Section 4.1].) This is the simplest case where the construction of coherent EPPA-witnesses is non-trivial and we encourage the reader to spend enough time on this section, as it can provide a very useful intuition for the subsequent sections. We consider graphs to be (relational) structures in a language with a single binary relation E which is symmetric and irreflexive.

Fix a graph \mathbf{A} with vertex set $A = \{1, \dots, n\}$.

Witness construction

We give a construction of a coherent EPPA-witness \mathbf{B} . It will be constructed as follows:

1. The vertices of \mathbf{B} are all pairs (x, χ) where $x \in A$ and χ is a function from $A \setminus \{x\}$ to $\{0, 1\}$ (called a *valuation function for x*).




Figure 4.2: Scheme of the construction of $\tilde{\varphi}$.

2. Vertices (x, χ) and (x', χ') form an edge of \mathbf{B} if and only if $x \neq x'$ and $\chi(x') \neq \chi'(x)$.

We now introduce a *generic copy* \mathbf{A}' of \mathbf{A} in \mathbf{B} using an embedding $\psi: \mathbf{A} \rightarrow \mathbf{B}$ defined by $\psi(x) = (x, \chi_x)$, where $\chi_x(y) = 1$ if $x > y$ and $\{x, y\} \in E_{\mathbf{A}}$ and $\chi_x(y) = 0$ otherwise (remember that we enumerated $A = \{1, \dots, n\}$). We put \mathbf{A}' to be the graph induced by \mathbf{B} on $\psi(A)$. It follows directly that ψ is indeed an embedding of \mathbf{A} into \mathbf{B} .

Remark 4.3.1. Note that the functions χ_x from the definition of ψ are in fact the rows of an asymmetric variant of the adjacency matrix of \mathbf{A} .

Let $\pi: B \rightarrow A$ be the *projection* mapping $(x, \chi) \mapsto x$. Note that $\pi(\psi(x)) = x$ for every $x \in \mathbf{A}$. This means that \mathbf{A}' is *transversal*, that is, π is injective on \mathbf{A}' .

Constructing the extension

The construction from the following paragraphs is schematically depicted in Figure 4.2.

Let φ be a partial automorphism of \mathbf{A}' . Using π we get a partial permutation of (the set) A and we denote by $\hat{\varphi}$ its order-preserving extension to a permutation of A (cf. Proposition 4.2.8).

We now construct a set $F \subseteq \binom{A}{2}$ of *flipped pairs* by putting $\{x, y\} \in F$ if $x \neq y$, $(x, \chi_x) \in \text{Dom}(\varphi)$ and $\chi_x(y) \neq \chi'(\hat{\varphi}(y))$, where $\varphi((x, \chi_x)) = (\hat{\varphi}(x), \chi')$. Note that if we also have that $(y, \chi_y) \in \text{Dom}(\varphi)$, then $\{x, y\} \in F$ if and only if $\chi_y(x) \neq \chi''(\hat{\varphi}(x))$, where $\varphi((y, \chi_y)) = (\hat{\varphi}(y), \chi'')$. This follows from the fact that φ is a partial automorphism.

For every $x \in A$ we define a function f_x on valuation functions for x putting

$$f_x(\chi)(\hat{\varphi}(y)) = \begin{cases} \chi(y) & \text{if } \{x, y\} \notin F \\ 1 - \chi(y) & \text{if } \{x, y\} \in F. \end{cases}$$

Finally, we define a function $\tilde{\varphi}: B \rightarrow B$ by putting $\tilde{\varphi}((x, \chi)) = (\hat{\varphi}(x), f_x(\chi))$. This function will be the coherent extension of φ .

Proofs

Both proofs in this section are only an explicit verification that our constructions work as expected.

Lemma 4.3.1. *$\tilde{\varphi}$ is an automorphism of \mathbf{B} extending φ . In other words, \mathbf{B} is an EPPA-witness for \mathbf{A} .*

Proof. Clearly, $\hat{\varphi}$ has an inverse $\hat{\varphi}^{-1}$. Observe also that the function f_x^{-1} defined as

$$f_x^{-1}(\chi)(\hat{\varphi}^{-1}(y)) = \begin{cases} \chi(y) & \text{if } \{\hat{\varphi}^{-1}(x), \hat{\varphi}^{-1}(y)\} \notin F \\ 1 - \chi(y) & \text{if } \{\hat{\varphi}^{-1}(x), \hat{\varphi}^{-1}(y)\} \in F \end{cases}$$

is an inverse of f_x . It follows that $\tilde{\varphi}$ is a bijection $B \rightarrow B$.

Let (x, χ) and (y, ξ) be vertices of \mathbf{B} . If $x = y$, then by the definition of \mathbf{B} neither of $(x, \chi), (y, \xi)$ and $\tilde{\varphi}((x, \chi)), \tilde{\varphi}((y, \xi))$ form an edge. If $x \neq y$, then we have $f_x(\chi)(\hat{\varphi}(y)) \neq f_y(\xi)(\hat{\varphi}(x))$ if and only if $\chi(y) \neq \xi(x)$ (by the definition of f_x and f_y), hence $\tilde{\varphi}$ preserves both edges and non-edges, that is, it is an automorphism of \mathbf{B} .

Let $(x, \chi_x) \in \text{Dom}(\varphi)$ with $\varphi((x, \chi_x)) = (z, \chi_z)$. We have for every $y \in A$ that $\{x, y\} \in F$ if and only if $\chi_x(y) \neq \chi_z(\hat{\varphi}(y))$. By the definition of f_x , it follows that $\chi_x(y) \neq f_x(\chi)(\hat{\varphi}(y))$ if and only if $\{x, y\} \in F$, therefore $\chi_z = f_x(\chi)$. This means that $\tilde{\varphi}$ indeed extends φ . \square

Lemma 4.3.2. *Let φ_1, φ_2 and φ be partial automorphisms of \mathbf{A} such that $\varphi = \varphi_2 \circ \varphi_1$ and $\tilde{\varphi}_1, \tilde{\varphi}_2$ and $\tilde{\varphi}$ their corresponding extensions as above. Then $\tilde{\varphi} = \tilde{\varphi}_2 \circ \tilde{\varphi}_1$.*

Proof. Denote by $\hat{\varphi}_1, \hat{\varphi}_2$ and $\hat{\varphi}$ the corresponding permutations of A constructed above, and by F_1, F_2 and F the corresponding sets of flipped pairs.

By Proposition 4.2.8 we get that $\hat{\varphi} = \hat{\varphi}_2 \circ \hat{\varphi}_1$. To see that $\tilde{\varphi}$ is a composition of $\tilde{\varphi}_1$ and $\tilde{\varphi}_2$ it remains to verify that pairs flipped by $\tilde{\varphi}$ are precisely those pairs that are flipped by the composition of $\tilde{\varphi}_1$ and $\tilde{\varphi}_2$.

This follows from the construction of F . Only pairs with at least one vertex in the domain of φ_1 are put into sets F and F_1 and again only pairs with at least one vertex in the domain of φ_2 (which is the same as the value range of φ_1) are put into F_2 .

Consider $\{x, y\} \in F$. This means that at least one of them (without loss of generality x) is in $\pi(\text{Dom}(\varphi)) = \pi(\text{Dom}(\varphi_1))$. Furthermore we know that $f_x(\chi_x)(\hat{\varphi}(y)) \neq \chi_x(y)$. Because $\varphi = \varphi_2 \circ \varphi_1$, we get that either $\{x, y\} \in F_1$, or $\{\hat{\varphi}_1(x), \hat{\varphi}_1(y)\} \in F_2$ (and precisely one of these happens). And this means that both $\tilde{\varphi}$ and $\tilde{\varphi}_2 \circ \tilde{\varphi}_1$ flip $\{x, y\}$.

On the other hand, if $\{x, y\} \notin F$, then either $\{x, y\}$ is in both F_1 and F_2 or in neither of them and then, again, neither $\tilde{\varphi}$ nor $\tilde{\varphi}_2 \circ \tilde{\varphi}_1$ flip $\{x, y\}$. This implies that indeed $\tilde{\varphi} = \tilde{\varphi}_2 \circ \tilde{\varphi}_1$. \square

The previous lemmas immediately imply the following proposition.

Proposition 4.3.3. *The graph \mathbf{B} is a coherent EPPA-witness for \mathbf{A}' .*

Remark 4.3.2. Note that \mathbf{B} only depends on the number of vertices of \mathbf{A} and, as such, is a coherent EPPA-witness for all graphs with at most $|A|$ vertices.

Remark 4.3.3. There is a simple generalisation of the ideas of this section which gives coherent EPPA for (not only) k -uniform hypergraphs directly, producing EPPA-witnesses on fewer vertices than Corollary 4.1.4 (see the next paragraph). For example, when $k = 3$, vertices of \mathbf{B} are pairs (x, χ) , where χ is a function from $\binom{A \setminus \{x\}}{2}$ (the set of all (unordered) pairs of vertices of \mathbf{A} different from x) to $\{0, 1\}$ and we say that $(x, \chi), (x', \chi'), (x'', \chi'')$ form a hyperedge of \mathbf{B} if and only if x, x' and x'' are distinct and $\chi(\{x', x''\}) + \chi'(\{x, x''\}) + \chi''(\{x, x'\})$ is odd. The rest of the construction is generalised in the same way, see also Section 4.4

Corollary 4.1.4 also implies coherent EPPA for the class of all k -uniform hypergraphs and other such classes. However, its proof takes a detour by first constructing EPPA-witnesses where the k -ary relation is not symmetric and contains tuples with repeated occurrences of the same vertices (a generalisation of loops), and then relying on a construction of irreducible structure faithful EPPA-witnesses to get a k -uniform hypergraph.

Remark 4.3.4. A minor change to the construction makes it possible to prove the extension property for partial switching automorphisms (which is a strengthening of standard EPPA), and hence also EPPA for two-graphs and antipodal metric spaces. This was done by Evans and the authors in [EHKN20].

Remark 4.3.5. Hrushovski's construction gives EPPA-witnesses on at most $(2n2^n)!$ vertices (where $|A| = n$) and he asks if this can be improved [Hru92, Section 3].² The combinatorial construction of Herwig and Lascar [HL00] provides EPPA-witnesses on roughly $(kn)^k$ vertices, where $|A| = n$ and k is the maximum degree of a vertex of \mathbf{A} . Our construction gives EPPA-witnesses on $n2^{n-1}$ vertices, thereby providing the best uniform bound over all graphs on n vertices. The same construction was found independently by Andr eka and N emeti [AN19].

Hrushovski also proves a lower bound $|B| \geq 2^m + m$ for $|A| = 2m$. It remains open to improve either of the bounds. Some partial progress on obtaining EPPA-witnesses of small size for some special classes of graphs has been made by Bradley-Williams and Cameron [BWC20]. We believe that studying bounds on the number of vertices of EPPA-witnesses is an interesting and meaningful project which can deepen our understanding of symmetries of graphs.

What now follows is a series of strengthenings of the main ideas from this section. Each of the constructions will proceed in several steps:

1. Define a structure \mathbf{B} using a suitable variant of valuations.
2. Give a construction of a generic copy \mathbf{A}' of \mathbf{A} in \mathbf{B} .
3. For a partial automorphism φ of \mathbf{A}' , give a construction of its coherent extension $\tilde{\varphi}: \mathbf{B} \rightarrow \mathbf{B}$.
4. Prove that $\tilde{\varphi}$ is indeed a coherent extension of φ and that \mathbf{B} and $\tilde{\varphi}$ have the extra properties required in the respective section.

We believe that the constructions are what is interesting. However, the proofs often contain some steps which are conceptually straightforward but slightly technical due to the nature of the constructions. We decided to state these technicalities as Claims and prove them at the very end of each section. We believe that this helps separate the key parts of the arguments from technical verifications.

²We would like to thank H. Andr eka and I. N emeti for bringing this question to our attention.

4.4 Coherent EPPA for relational structures

In this section we generalise the construction from the previous section to prove the following proposition:

Proposition 4.4.1. *Let L be a finite relational language equipped with a permutation group Γ_L and let \mathbf{A} be a finite Γ_L -structure. There exists a finite Γ_L -structure \mathbf{B} which is a coherent EPPA-witness for \mathbf{A} .*

Fix a finite relational language L equipped with a permutation group Γ_L and a finite Γ_L -structure \mathbf{A} with $A = \{1, \dots, k\}$. We will construct a Γ_L -structure \mathbf{B} and give an embedding $\psi: \mathbf{A} \rightarrow \mathbf{B}$ such that \mathbf{B} is a coherent EPPA-witness for \mathbf{A} (with respect to ψ). Proposition 4.4.1 then immediately follows.

Witness construction

Given a vertex $x \in A$ and an integer n , we denote by $U_n^A(x)$ the set of all n -tuples (i.e. n -element sequences) of elements of A containing x . Note that $U_n^A(x)$ also includes n -tuples with repeated occurrences of vertices.

Given a relation $R \in L$ of arity n and a vertex $x \in A$, we say that a function $\xi: U_n^A(x) \rightarrow \{0, 1\}$ is an R -valuation function for x . An L -valuation function for a vertex $x \in A$ is a function χ assigning to every $R \in L$ an R -valuation function $\chi(R)$ for x .

Now we are ready to give the definition of \mathbf{B} :

1. The vertices of \mathbf{B} are all pairs (x, χ) , where $x \in A$ and χ is an L -valuation function for x .
2. For every relation symbol R of arity n we put

$$((x_1, \chi_1), \dots, (x_n, \chi_n)) \in R_{\mathbf{B}}$$

if and only if for every $1 \leq i < j \leq n$ such that $x_i = x_j$ it also holds that $\chi_i = \chi_j$ and furthermore

$$\sum_{\chi \in \{\chi_i: 1 \leq i \leq n\}} \chi(R)((x_1, \dots, x_n)) \text{ is odd}$$

(summing over $\chi \in \{\chi_i: 1 \leq i \leq n\}$ ensures that possible multiple occurrences of (x_i, χ_i) are only counted once).

Next we give an embedding $\psi: \mathbf{A} \rightarrow \mathbf{B}$ by putting ψ_L to be the identity, and

$$\psi_A(x) = (x, \chi_x),$$

where χ_x is an L -valuation function for x such that for every $R \in L$ we have

$$\chi_x(R)((y_1, \dots, y_n)) = \begin{cases} 1 & \text{if } (y_1, \dots, y_n) \in R_{\mathbf{A}} \text{ and } x = y_1, \\ 0 & \text{otherwise.} \end{cases}$$

The following claim follows from the construction:

Claim 4.4.2. ψ is an embedding $\mathbf{A} \rightarrow \mathbf{B}$.

Proof. Fix an n -ary relation $R \in L$. Recall that for $x \in A$, we have $\psi(x) = (x, \chi_x)$ with $\chi_x(R)(\bar{y}) = 1$ if and only if $\bar{y}_1 = x$ and $\bar{y} \in R_{\mathbf{A}}$. In particular, if $\bar{y} \notin R_{\mathbf{A}}$, then $\psi(\bar{y}) \notin R_{\mathbf{B}}$, as for every i we have that $\chi_{\bar{y}_i}(R)(\bar{y}) = 0$.

Suppose now that $\bar{y} \in R_{\mathbf{A}}$. For every i we have $\chi_{\bar{y}_i}(R)(\bar{y}) = 1$ if and only if $\bar{y}_i = \bar{y}_1$. Hence

$$\sum_{\chi \in \{\chi_{\bar{y}_i} : 1 \leq i \leq n\}} \chi(R)(\bar{y}) = 1,$$

so it is odd and thus $\psi(\bar{y}) \in R_{\mathbf{B}}$. \square

Put $\mathbf{A}' = \psi(\mathbf{A})$. This is the copy whose partial automorphisms we will later extend. Let $\pi: B \rightarrow A$ defined as $\pi((x, \chi)) = x$ be the *projection*.

Constructing the extension

As in Section 4.3, we fix a partial automorphism $\varphi: \mathbf{A}' \rightarrow \mathbf{A}'$ and extend the projection of φ to a permutation $\hat{\varphi}$ of A in an order-preserving way. Note that φ already contains a permutation of the language, therefore we will focus on extending the structural part.

For every relation symbol $R \in L$ of arity n , we construct a function $F_R: A^n \rightarrow \{0, 1\}^n$. These functions will play a similar role as the set F in Section 4.3 (i.e., they will control the *flips*) and are constructed as follows:

For an n -tuple $\bar{x} = (x_1, \dots, x_n)$ and $1 \leq i \leq n$, we put $F_R(\bar{x})_i = 1$ if and only if one of the following two cases is true:

1. $x_i \in \pi(\text{Dom}(\varphi))$ and $\chi_{x_i}(R)(\bar{x}) \neq \chi_{x_j}(\varphi(R))(\hat{\varphi}(\bar{x}))$, where $\varphi((x_i, \chi_{x_i})) = (x_j, \chi_{x_j})$.
2. $|\{x_j : x_j \in \pi(\text{Dom}(\varphi)) \text{ and } F_R(\bar{x})_j = 1\}|$ is odd and $x_i = x_m$, where $1 \leq m \leq n$ is the smallest index such that $x_m \notin \pi(\text{Dom}(\varphi))$ (note that m might not exist, but then all entries are covered by case 1).

All the other entries of $F_R(\bar{x})$ are equal to 0. Note that case 2 ensures that there is an even number of distinct vertices of \bar{x} whose corresponding entry in $F_R(\bar{x})$ is equal to 1.

For every $x \in A$ we define a function f_x on L -valuation functions for x , putting

$$f_x(\chi)(\varphi(R))(\hat{\varphi}(\bar{y})) = \begin{cases} \chi(R)(\bar{y}) & \text{if } F_R(\bar{y}) \text{ has 0 on entry corresponding to } x \\ 1 - \chi(R)(\bar{y}) & \text{if } F_R(\bar{y}) \text{ has 1 on entry corresponding to } x. \end{cases}$$

Finally, we define $\tilde{\varphi}: \mathbf{B} \rightarrow \mathbf{B}$ by putting $\tilde{\varphi}_L = \varphi_L$ and

$$\tilde{\varphi}_B((x, \chi)) = (\hat{\varphi}(x), f_x(\chi)).$$

Proofs

In the rest we proceed analogously to Section 4.3.

Lemma 4.4.3. $\tilde{\varphi}$ is an automorphism of \mathbf{B} extending φ .

Proof. In the same way as in Lemma 4.3.1 one can see that $\tilde{\varphi}$ is a bijection. Observe that by the construction we get that for every $R \in L$ of arity n and every n -tuple $\bar{x} = (x_1, \dots, x_n) \in A^n$ we have that $F_R(\bar{x})_i = F_R(\bar{x})_j$ whenever $x_i = x_j$ and that

$$|\{x_i : F_R(\bar{x})_i = 1\}| \text{ is even}$$

(where taking the size of the set means that each distinct vertex is counted only once even if it has repeated occurrences in \bar{x}): Indeed, if \bar{x} contains vertices from $A \setminus \pi(\text{Dom}(\varphi))$, this follows directly. Otherwise all vertices of \bar{x} are from $\pi(\text{Dom}(\varphi))$, but then (as φ is a partial automorphism), we get that

$$((x_1, \chi_{x_1}), \dots, (x_n, \chi_{x_n})) \in R_{\mathbf{B}} \iff (\varphi((x_1, \chi_{x_1})), \dots, \varphi((x_n, \chi_{x_n}))) \in \varphi(R)_{\mathbf{B}},$$

and using the definition of relations in \mathbf{B} we see that an even number of distinct vertices from $\psi(\bar{x})$ changed how they valuate \bar{x} with respect to R and $\hat{\varphi}(\bar{x})$ with respect to $\varphi(R)$ respectively.

Pick $\vec{x} = ((x_1, \chi_1), \dots, (x_n, \chi_n)) \in B^n$ and put $\bar{x} = (x_1, \dots, x_n)$. Recall that $\vec{x} \in R_{\mathbf{B}}$ if and only if for every $1 \leq i < j \leq n$ such that $x_i = x_j$ it also holds that $\chi_i = \chi_j$ and furthermore

$$\sum_{\chi \in \{\chi_i : 1 \leq i \leq n\}} \chi(R)(\bar{x}) \text{ is odd.}$$

Note that if $x_i = x_j$ then $f_{x_i}(\chi_i) = f_{x_j}(\chi_j)$ if and only if $\chi_i = \chi_j$. To get that $\tilde{\varphi}$ is an automorphism, it remains to show that

$$\sum_{\chi \in \{f_{x_i}(\chi_i) : 1 \leq i \leq n\}} \chi(\varphi(R))(\hat{\varphi}(\bar{x})) \text{ is odd}$$

if and only if

$$\sum_{\chi \in \{\chi_i : 1 \leq i \leq n\}} \chi(R)(\bar{x}) \text{ is odd.}$$

By the construction of f_x , we have for every i that

$$f_{x_i}(\chi_i)(\varphi(R))(\hat{\varphi}(\bar{x})) = \chi_i(R)(\bar{x})$$

if and only if $F_R(\bar{x})_i = 0$. This means that

$$|\{x_i : f_{x_i}(\chi_i)(\varphi(R))(\hat{\varphi}(\bar{x})) \neq \chi_i(R)(\bar{x})\}|$$

is equal to $|\{x_i : F_R(\bar{x})_i = 1\}|$, which is an even number. This concludes the proof that $\tilde{\varphi}$ is an automorphism of \mathbf{B} .

To see that $\tilde{\varphi}$ extends φ , pick a tuple $\bar{x} = (x_1, \dots, x_n) \in A^n$, an index $1 \leq i \leq n$ such that $(x_i, \chi_{x_i}) \in \text{Dom}(\varphi)$, and an arbitrary $R \in L$. Recall that $F_R(\bar{x})_i = 1$ if and only if $\chi_{x_i}(R)(\bar{x}) \neq \chi_{x_j}(\varphi(R))(\hat{\varphi}(\bar{x}))$, where $\varphi((x_i, \chi_{x_i})) = (x_j, \chi_{x_j})$. Since $f_{x_i}(\chi_{x_i})(\varphi(R))(\hat{\varphi}(\bar{x})) \neq \chi_{x_i}(R)(\bar{x})$ if and only if $F_R(\bar{x})_i = 1$, we get that $\chi_{x_j} = f_{x_i}(\chi_{x_i})$ and hence $\varphi \subseteq \tilde{\varphi}$. \square

Lemma 4.4.4. *Let φ_1, φ_2 and φ be partial automorphisms of \mathbf{A} such that $\varphi = \varphi_2 \circ \varphi_1$ and let $\tilde{\varphi}_1, \tilde{\varphi}_2$ and $\tilde{\varphi}$ be their corresponding extensions as above. Then $\tilde{\varphi} = \tilde{\varphi}_2 \circ \tilde{\varphi}_1$.*

Proof. Let $\hat{\varphi}_1$, $\hat{\varphi}_2$ and $\hat{\varphi}$ be the permutations of A constructed in the previous section for φ_1 , φ_2 and φ respectively, similarly define F_R^1 , F_R^2 and F_R for every $R \in L$. Since $\hat{\varphi}_1$, $\hat{\varphi}_2$ and $\hat{\varphi}$ were chosen in an order-preserving way, by Proposition 4.2.8 we get that $\hat{\varphi} = \hat{\varphi}_2 \circ \hat{\varphi}_1$.

Hence, by an argument analogous to the proof of Lemma 4.3.2, we can see that it suffices to show that for every $R \in L$, for every $\bar{x} \in A^{a(R)}$ and for every $1 \leq i \leq a(R)$, it holds that $F_R(\bar{x})_i = F_R^1(\bar{x})_i + F_{\varphi_1(R)}^2(\hat{\varphi}_1(\bar{x}))_i \pmod{2}$.

Fix such R , \bar{x} and i . Put $y = \bar{x}_i$ and $(y, \chi) = \psi(y)$. First suppose that $(y, \chi) \in \text{Dom}(\varphi) = \text{Dom}(\varphi_1)$ and denote $(y', \chi') = \varphi_1((y, \chi))$ and $(y'', \chi'') = \varphi((y, \chi)) = \varphi_2((y', \chi'))$. By the construction, we have the following:

$$\begin{aligned} F_R^1(\bar{x})_i = 1 &\iff \chi(\bar{x}) \neq \chi'(\hat{\varphi}_1(\bar{x})), \\ F_{\varphi_1(R)}^2(\hat{\varphi}_1(\bar{x}))_i = 1 &\iff \chi'(\hat{\varphi}_1(\bar{x})) \neq \chi''(\hat{\varphi}(\bar{x})), \\ F_R(\bar{x})_i = 1 &\iff \chi(\bar{x}) \neq \chi''(\hat{\varphi}(\bar{x})). \end{aligned}$$

It immediately follows that $F_R(\bar{x})_i = 1$ if and only if exactly one of $F_R^1(\bar{x})_i$ and $F_{\varphi_1(R)}^2(\hat{\varphi}_1(\bar{x}))_i$ is equal to one and we are done.

Otherwise $(y, \chi) \notin \text{Dom}(\varphi)$. Let m , m_1 and m_2 be the indices from case 2 of the definition of F_R for \bar{x} , F_R^1 for \bar{x} , and $F_{\varphi_1(R)}^2$ for $\hat{\varphi}_1(\bar{x})$, respectively. Since $(\varphi_1, \varphi_2, \varphi)$ is a coherent triple, we get that $m = m_1 = m_2$.

If $y = \bar{x}_i \neq \bar{x}_m$ is not in $\pi(\text{Dom}(\varphi))$, it follows that $F_R^1(\bar{x})_i = F_{\varphi_1(R)}^2(\hat{\varphi}_1(\bar{x}))_i = F_R(\bar{x})_i = 0$. Now we will assume that $\bar{x}_i = \bar{x}_m$. Define

$$\begin{aligned} I &= \{1 \leq j \leq n : \bar{x}_j \in \pi(\text{Dom}(\varphi)) \text{ and } F_R(\bar{x})_j = 1\}, \\ I_1 &= \{1 \leq j \leq n : \bar{x}_j \in \pi(\text{Dom}(\varphi_1)) \text{ and } F_R^1(\bar{x})_j = 1\}, \text{ and} \\ I_2 &= \{1 \leq j \leq n : \hat{\varphi}_1(\bar{x})_j \in \pi(\text{Dom}(\varphi_2)) \text{ and } F_{\varphi_1(R)}^2(\hat{\varphi}_1(\bar{x}))_j = 1\}. \end{aligned}$$

Observe that by the previous paragraphs we have that I is the symmetric difference of I_1 and I_2 , so in particular

$$|I| = |I_1| + |I_2| - 2|I_1 \cap I_2|.$$

Also note that $\bar{x}_j = \bar{x}_k$ if and only if $\hat{\varphi}_1(\bar{x})_j = \hat{\varphi}_1(\bar{x})_k$. It follows that

$$|\{x_j : j \in I\}| = |\{x_j : j \in I_1\}| + |\{x_j : j \in I_2\}| - 2|\{x_j : j \in I_1 \cap I_2\}|,$$

where $|\{x_j : j \in I\}|$ is the number of distinct vertices of \mathbf{A} such that their corresponding entry in $F_R(\bar{x})$ is equal to one. Looking at this equation modulo 2, we get that $|\{x_j : j \in I\}|$ is odd if and only if precisely one of $|\{x_j : j \in I_1\}|$ and $|\{x_j : j \in I_2\}|$ is odd.

This implies (comparing with case 2 of the definitions of F_R , F_R^1 and $F_{\varphi_1(R)}^2$) that even for $\bar{x}_i = \bar{x}_m$, we have that $F_R(\bar{x})_i = F_R^1(\bar{x})_i + F_{\varphi_1(R)}^2(\hat{\varphi}_1(\bar{x}))_i \pmod{2}$, which finishes the proof. \square

This finishes the proof of Proposition 4.4.1.

Remark 4.4.1. The EPPA-witness \mathbf{B} constructed in this section has at most

$$\mathcal{O}\left(|A|2^{|L||A|^m}\right)$$

vertices, where m is the largest arity of a relation in L . Consequently, the size of a coherent EPPA-witness for \mathbf{A} only depends on the language and on the number of vertices of \mathbf{A} .

4.5 Infinite languages

When a non-trivial permutation group is present it is not true that for every finite structure there is a finite EPPA-witness. Consider, for example, the language L consisting of infinitely many unary relations, where Γ_L is the symmetric group. Let \mathbf{A} be a structure with a single vertex which is in exactly one relation. Then every EPPA-witness for \mathbf{A} needs to, in particular, extend all partial automorphisms of \mathbf{A} of type (g, \emptyset) , where $g \in \Gamma_L$ and \emptyset is the empty map. This implies that every EPPA-witness for \mathbf{A} must contain a vertex in precisely one unary relation U for every $U \in L$, hence infinitely many vertices.

First, we generalise this argument and prove Theorem 4.1.3:

Proof of Theorem 4.1.3. \mathbf{A} being in an infinite orbit means that there is an infinite sequence $g_1, g_2, \dots \in \Gamma_L$ such that the sequence $(g_1, \text{id}_A)(\mathbf{A}), (g_2, \text{id}_A)(\mathbf{A}), \dots$ consists of pairwise distinct structures. For a contradiction, assume that there is a (finite) EPPA-witness \mathbf{B} for \mathbf{A} .

In particular, \mathbf{B} needs to extend all partial automorphisms (g_i, \emptyset) , $i \geq 1$, which means that for every $i \geq 1$, there is an embedding of $(g_i, \text{id}_A)(\mathbf{A})$ into \mathbf{B} . In other words, for every $i \geq 1$ we get a tuple $\bar{x}_i \in B^{|A|}$, and by the assumption, all these tuples are pairwise distinct. This implies that the set $B^{|A|}$ is infinite, and since $|A|$ is finite, it follows that B is infinite, a contradiction. \square

On the positive side, we prove the following proposition, thereby characterising relational languages L equipped with a permutation group Γ_L for which the class of all finite Γ_L -structures has EPPA.

Proposition 4.5.1. *Let L be a relational language equipped with a permutation group Γ_L and let \mathbf{A} be a finite Γ_L -structure such that \mathbf{A} lies in a finite orbit of the action of Γ_L by relabelling. There is a finite Γ_L -structure \mathbf{B} which is a coherent EPPA-witness for \mathbf{A} .*

In order to prove Proposition 4.5.1, we will need the following lemma.

Lemma 4.5.2. *Let M be a relational language equipped with a permutation group Γ_M and let \mathcal{C} be a class of finite Γ_M -structures. Suppose that there is a set $N \subseteq M$ such that for every $\mathbf{A} \in \mathcal{C}$ and every $R \in M \setminus N$, it holds that $R_{\mathbf{A}} = \emptyset$ and for every $\pi \in \Gamma_M$ we have $\pi(N) = N$ (i.e. π fixes N setwise). Put*

$$\Gamma_N = \{\pi \upharpoonright_N : \pi \in \Gamma_M\},$$

where for a permutation $\pi: M \rightarrow M$, $\pi \upharpoonright_N$ is its restriction to N .

Then Γ_N is a permutation group on N and \mathcal{C} has coherent EPPA if and only if \mathcal{D} does, where \mathcal{D} is the class consisting of the same structures as \mathcal{C} , but understood as Γ_N -structures.

Proof. Since every $\pi \in \Gamma_M$ fixes N setwise, we immediately get that Γ_N is a permutation group on N . If \mathcal{C} has coherent EPPA, then clearly \mathcal{D} does, too, because for each $\pi' \in \Gamma_N$, we can simply pick an arbitrary $\pi \in \Gamma_M$ such that $\pi' = \pi \upharpoonright_N$ and use coherent EPPA for \mathcal{C} . It thus remains to prove the other direction.

In the following, for $\mathbf{A} \in \mathcal{C}$, we denote by \mathbf{A}^N its corresponding Γ_N -structure from \mathcal{D} .

Fix $\mathbf{A} \in \mathcal{C}$. By the assumption that \mathcal{D} has coherent EPPA, we get $\mathbf{B} \in \mathcal{C}$ such that \mathbf{B}^N is a coherent EPPA-witness for \mathbf{A}^N . Let $f = (f_L, f_A)$ be a partial automorphism of \mathbf{A} . Then $(f_L \upharpoonright_N, f_A)$ is a partial automorphism of \mathbf{A}^N and it extends to an automorphism $(f_L \upharpoonright_N, \theta)$ of \mathbf{B}^N . It is straightforward to check that (f_L, θ) is an automorphism of \mathbf{B} extending f (it clearly extends f , and it is an automorphism of \mathbf{B} , because \mathbf{B} contains no relations from $M \setminus N$ and $f_L(N) = N$). Coherence follows by coherence in \mathcal{D} . \square

4.5.1 Proof of Proposition 4.5.1

First we will define some auxiliary notions. Given an n -tuple $\bar{x} = (x_1, \dots, x_n)$ and a function $\omega: \{1, \dots, m\} \rightarrow \{1, \dots, n\}$, we define an m -tuple

$$\bar{x} \circ \omega = (x_{\omega(1)}, \dots, x_{\omega(m)}).$$

For a Γ_L -structure \mathbf{B} and an n -tuple $\bar{x} \in B$ containing no repeated vertices (i.e. if $\bar{x}_i = \bar{x}_j$, then $i = j$), we define $\sigma(\bar{x}, \mathbf{B})$ to be the set of all pairs (R, ω) , where $R \in L$ is an m -ary relation and $\omega: \{1, \dots, m\} \rightarrow \{1, \dots, n\}$ is a surjective function, such that $\bar{x} \circ \omega \in R_{\mathbf{B}}$.

Next, we define sets M_1, M_2, \dots , such that M_n consists of all pairs (R, ω) , where $R \in L$ is an m -ary relation and ω is a surjection $\{1, \dots, m\} \rightarrow \{1, \dots, n\}$. For every n and for every $X \subseteq M_n$ we assume that L does not contain the symbol R^X and we let R^X be an n -ary relation. We define a language M' to consist of all these relations R^X .

We put $M = M' \cup L$ (remember that $M' \cap L = \emptyset$ by our assumption). For every $g \in \Gamma_L$, we define a permutation π_g on M such that

$$\pi_g(R) = \begin{cases} g(R) & \text{if } R \in L, \\ R^Y & \text{if } R = R^X \in M', \end{cases}$$

where $Y = \{(g(S), \omega) : (S, \omega) \in X\}$.

Finally, we put $\Gamma_M = \{\pi_g : g \in \Gamma_L\}$. It is easy to verify that the map $g \mapsto \pi_g$ is a group isomorphism $\Gamma_L \rightarrow \Gamma_M$ (this is the only place in the proof where we use that $L \subseteq M$).

Given a Γ_L -structure \mathbf{B} , we define a Γ_M -structure $T(\mathbf{B})$ such that the vertex set of $T(\mathbf{B})$ is B and for every $\bar{x} \in B^n$ containing no repeated vertices, we put $\bar{x} \in R_{T(\mathbf{B})}^{\sigma(\bar{x}, \mathbf{B})}$. There are no other tuples in any relations of $T(\mathbf{B})$.

In the other direction, given a Γ_M -structure \mathbf{B} such that $R_{\mathbf{B}} = \emptyset$ for every $R \in L$, we define a Γ_L -structure $U(\mathbf{B})$ such that the vertex set of $U(\mathbf{B})$ is B , and whenever $\bar{x} \in R_{\mathbf{B}}^X$, we put $\bar{x} \circ \omega \in R_{U(\mathbf{B})}^S$ for every $(S, \omega) \in X$. There are no other tuples in any relations of $U(\mathbf{B})$. It is easy to verify that T and U are mutually inverse, that is $UT(\mathbf{B}) = \mathbf{B}$ for every Γ_L -structure \mathbf{B} , and $TU(\mathbf{B}) = \mathbf{B}$ for every Γ_M -structure \mathbf{B} such that $R_{\mathbf{B}} = \emptyset$ for every $R \in L$.

In fact, these maps are functorial in the sense of the following lemma.

Lemma 4.5.3. *Let \mathbf{B}, \mathbf{C} be Γ_L -structures. Let (g, f) be an embedding $\mathbf{B} \rightarrow \mathbf{C}$ ($g \in \Gamma_L$, $f: B \rightarrow C$). Then (π_g, f) is an embedding $T(\mathbf{B}) \rightarrow T(\mathbf{C})$.*

Let \mathbf{B}, \mathbf{C} be Γ_M -structures such that $R_{\mathbf{B}} = R_{\mathbf{C}} = \emptyset$ for every $R \in L$. Let (π_g, f) be an embedding $\mathbf{B} \rightarrow \mathbf{C}$ ($\pi_g \in \Gamma_M$, $f: B \rightarrow C$). Then (g, f) is an embedding $U(\mathbf{B}) \rightarrow U(\mathbf{C})$.

Proof. We only need to verify the definition of an embedding. For the first part, we know that for every $S \in L$, every n -tuple $\bar{x} \in B$ containing no repeated vertices and for every surjection $\omega: \{1, \dots, m\} \rightarrow \{1, \dots, n\}$, we have $\bar{x} \circ \omega \in S_{\mathbf{B}}$ if and only if $f(\bar{x}) \circ \omega \in g(S)_{\mathbf{B}}$, which implies that

$$\sigma(f(\bar{x}), (g, f)(\mathbf{B})) = \{(g(S), \omega) : (S, \omega) \in \sigma(\bar{x}, \mathbf{B})\},$$

from which the claim follows. The second part can be proved in a complete analogy. \square

Continuing with the proof of Proposition 4.5.1 we define N to be the subset of M' consisting of all $\pi_g(R^{\sigma(\bar{x}, \mathbf{A})})$, where $\pi_g \in \Gamma_M$ and \bar{x} is a tuple of vertices of \mathbf{A} containing no repeated ones (remember that \mathbf{A} is the Γ_L -structure fixed in the statement of Proposition 4.5.1).

We claim that N is finite: Whenever $R^X \in N$, then there is $\bar{x} \in A$ and $\pi_g \in \Gamma_M$ such that $R^X = \pi_g(R^{\sigma(\bar{x}, \mathbf{A})})$, however, this is equivalent to saying that $X = \sigma(\bar{x}, (g, \text{id}_A)(\mathbf{A}))$. In other words, every n -ary relation $R^X \in N$ corresponds to at least one pair $(\bar{x}, (g, \text{id}_A)(\mathbf{A}))$, where \bar{x} is an n -tuple of vertices of \mathbf{A} with no repeated occurrences and $g \in \Gamma_L$. Since \mathbf{A} lies in a finite orbit of the action of Γ_L by relabelling, it follows that there are only finitely many different choices for $(g, \text{id}_A)(\mathbf{A})$. By definition, all relations in N have arity at most $|A|$, hence there are finitely many choices for \bar{x} and thus N is indeed finite. Observe also that for every $\pi_g \in \Gamma_M$ we have $\pi_g(N) = N$.

Let \mathcal{C} be the class consisting of all Γ_M -structures \mathbf{B} such that whenever $R \in M \setminus N$, then $R_{\mathbf{B}} = \emptyset$. Observe that $T(\mathbf{A}) \in \mathcal{C}$ and $U(\mathbf{B})$ is defined for every $\mathbf{B} \in \mathcal{C}$.

We have verified that \mathcal{C} satisfies the conditions of Lemma 4.5.2. Hence, we get a permutation group Γ_N on N and a class \mathcal{D} , which in this case is simply the class of all finite Γ_N -structures and hence has coherent EPPA by Proposition 4.4.1. By Lemma 4.5.2 we then get that \mathcal{C} also has coherent EPPA.

In particular, we get $\mathbf{C} \in \mathcal{C}$ which is a coherent EPPA-witness for $T(\mathbf{A})$. Putting $\mathbf{B} = U(\mathbf{C})$, we have a Γ_L -structure \mathbf{B} such that $T(\mathbf{B})$ is a coherent EPPA-witness for $T(\mathbf{A})$. In the last paragraph, we shall prove that \mathbf{B} is a coherent EPPA-witness for \mathbf{A} .

Let (g, f) be a partial automorphism of \mathbf{A} . From the construction it follows that (π_g, f) is a partial automorphism of $T(\mathbf{A})$ (equivalently, it follows from Lemma 4.5.3 and the observation that a partial automorphism of \mathbf{A} can be understood as a pair of embeddings of the same structure into \mathbf{A}), which extends to an automorphism (π_g, \tilde{f}) of $T(\mathbf{B})$ (that is, $f \subseteq \tilde{f}$). By Lemma 4.5.3 again, we get that (g, \tilde{f}) is an automorphism of \mathbf{B} , and since $f \subseteq \tilde{f}$, it extends (g, f) . Coherence follows from coherence of $T(\mathbf{B})$, because $\pi_{gf} = \pi_g \pi_f$. This finishes the proof of Proposition 4.5.1.

4.6 EPPA for structures with unary functions

We are now ready to introduce unary functions into the language. In order to do so, we will use valuation structures instead of valuation functions, which was first done in [EHN21]. Otherwise we follow the general scheme as above and prove the following proposition.

Proposition 4.6.1. *Let L be a language consisting of relation and unary function symbols equipped with a permutation group Γ_L and let \mathbf{A} be a finite Γ_L -structure which lies in a finite orbit of the action of Γ_L by relabelling. Then there is a coherent EPPA-witness for \mathbf{A} .*

Fix a language L consisting of relation and unary function symbols equipped with a permutation group Γ_L , and a finite Γ_L -structure \mathbf{A} which lies in a finite orbit of the action of Γ_L by relabelling.

Denote by $L_{\mathcal{R}} \subseteq L$ the language consisting of all relation symbols of L and let $\Gamma_{L_{\mathcal{R}}}$ be the group obtained by restricting permutations from Γ_L to $L_{\mathcal{R}}$. For a Γ_L -structure \mathbf{D} , we will denote by \mathbf{D}^- the $\Gamma_{L_{\mathcal{R}}}$ -reduct of \mathbf{D} (that is, the $\Gamma_{L_{\mathcal{R}}}$ -structure on the same vertex set as \mathbf{D} with $R_{\mathbf{D}^-} = R_{\mathbf{D}}$ for every $R \in L_{\mathcal{R}}$)

Witness construction

Let \mathbf{B}_0 be a finite $\Gamma_{L_{\mathcal{R}}}$ -structure which is a coherent EPPA-witness for \mathbf{A}^- (\mathbf{B}_0 exists by Proposition 4.5.1). We furthermore, for convenience, assume that $\mathbf{A}^- \subseteq \mathbf{B}_0$. Let $x \in B_0$ be a vertex of \mathbf{B}_0 and let \mathbf{V} be a Γ_L -structure. We say that \mathbf{V} is a *valuation structure for x* if the following hold:

1. $x \in V$,
2. there exists $y \in A$ and an isomorphism $\iota: \mathbf{V} \rightarrow \text{Cl}_{\mathbf{A}}(y)$ satisfying $\iota(x) = y$ (note that ι can permute the language),
3. \mathbf{V}^- is a substructure of \mathbf{B}_0 .

Note that if L contains no functions, then there is exactly one valuation structure for every $x \in B_0$, namely the substructure of \mathbf{B}_0 induced on $\{x\}$. In this case, the rest of this construction simply describes the identity.

We construct \mathbf{B} as follows:

1. The vertices of \mathbf{B} are all pairs (x, \mathbf{V}) where $x \in B_0$ and \mathbf{V} is a valuation structure for x ,
2. for every relation symbol $R \in L$ of arity n , we put $((x_1, \mathbf{V}_1), \dots, (x_n, \mathbf{V}_n)) \in R_{\mathbf{B}}$ if and only if $(x_1, \dots, x_n) \in R_{\mathbf{B}_0^-}$,
3. for every (unary) function symbol $F \in L$ we put

$$F_{\mathbf{B}}((x, \mathbf{V})) = \{(y, \text{Cl}_{\mathbf{V}}(y)) : y \in F_{\mathbf{V}}(x)\}.$$

Since \mathbf{V} is a valuation structure for x , it follows that $\text{Cl}_{\mathbf{V}}(y)$ is a valuation structure for y .

Next we define an embedding $\psi: \mathbf{A} \rightarrow \mathbf{B}$, putting $\psi_A(x) = (x, \text{Cl}_{\mathbf{A}}(x))$ and $\psi_L = \text{id}_L$. Note that $\text{Cl}_{\mathbf{A}}(x)$ (a substructure of \mathbf{A}) is indeed a valuation structure for ι being the identity, because we assumed that $\mathbf{A}^- \subseteq \mathbf{B}_0$. We put $\mathbf{A}' = \psi(\mathbf{A})$ to be the copy of \mathbf{A} in \mathbf{B} whose partial automorphisms we will extend.

Claim 4.6.2. *ψ is an embedding $\mathbf{A} \rightarrow \mathbf{B}$.*

Proof. It follows directly from the construction that ψ is injective and that for every relation $R \in L$ we have $\bar{x} \in R_{\mathbf{A}}$ if and only if $\psi(\bar{x}) \in R_{\mathbf{B}}$. It remains to verify that for every $x \in \mathbf{A}$ and for every function $F \in L$ we have $\psi(F_{\mathbf{A}}(x)) = F_{\mathbf{B}}(\psi(x))$.

By the construction of \mathbf{B} we have

$$F_{\mathbf{B}}(\psi(x)) = F_{\mathbf{B}}((x, \text{Cl}_{\mathbf{A}}(x))) = \{(y, \text{Cl}_{\text{Cl}_{\mathbf{A}}(x)}(y)) : y \in F_{\text{Cl}_{\mathbf{A}}(x)}(x)\}.$$

Since $\text{Cl}_{\text{Cl}_{\mathbf{A}}(x)}(y) = \text{Cl}_{\mathbf{A}}(y)$ and $F_{\text{Cl}_{\mathbf{A}}(x)}(x) = F_{\mathbf{A}}(x)$, we have

$$F_{\mathbf{B}}((x, \text{Cl}_{\mathbf{A}}(x))) = \{(y, \text{Cl}_{\mathbf{A}}(y)) : y \in F_{\mathbf{A}}(x)\},$$

which is exactly $\psi(F_{\mathbf{A}}(x))$. □

Observe that the vertex set of \mathbf{B} is finite: Assume for a contradiction that it is infinite. Since there are finitely many vertices in B_0 , this implies that there is $x \in B_0$ for which there are infinitely many valuation structures. Moreover, by the definition of a valuation structure, this implies that in fact there is a vertex $y \in A$, a structure $\mathbf{W} \subseteq \mathbf{B}_0$, an injection $\iota_W: W \rightarrow \text{Cl}_{\mathbf{A}}(y)$, a sequence of permutations g_1, g_2, \dots and a sequence of structures $\mathbf{V}_1, \mathbf{V}_2, \dots$, such that the following hold:

1. $\mathbf{V}_i^- = \mathbf{W}$ for every $i \geq 1$ (so, in particular, they have the same vertex set),
2. the structures $\mathbf{V}_1, \mathbf{V}_2, \dots$ are pairwise distinct, and
3. (g_i, ι_W) is an embedding $\mathbf{V}_i \rightarrow \text{Cl}_{\mathbf{A}}(y)$ for every $i \geq 1$.

Taking the inverse, we get that there is a substructure $(g_1, \iota_W)(\mathbf{V}_1) = \mathbf{X} \subseteq \text{Cl}_{\mathbf{A}}(y)$ such that the structures $(g_i^{-1}, \iota_W^{-1})(\mathbf{X}) = \mathbf{V}_i$, $i \geq 1$ are pairwise distinct, which gives a contradiction with \mathbf{A} lying in a finite orbit of the action of Γ_L by relabelling.

Constructing the extension

Let $\pi: B \rightarrow B_0$, defined by $\pi((x, \mathbf{V})) = x$, be the projection. Note that $\pi(\mathbf{A}') = \mathbf{A}^-$. Fix a partial automorphism φ of \mathbf{A}' . It induces (by π and restriction to $\Gamma_{L_{\mathcal{R}}}$) a partial automorphism φ_0 of \mathbf{A}^- . Denote by $\hat{\varphi}$ the extension of φ_0 to an automorphism of \mathbf{B}_0 . Put $\tilde{\varphi}_L = \varphi_L$ and

$$\tilde{\varphi}_B((x, \mathbf{V})) = (\hat{\varphi}(x), (\varphi_L, \hat{\varphi})(\mathbf{V})).$$

Proofs

We again proceed analogously to Section 4.3.

Lemma 4.6.3. *$\tilde{\varphi}$ is an automorphism of \mathbf{B} extending φ .*

Proof. Since $\hat{\varphi}$ is a bijection $B_0 \rightarrow B_0$, it follows that for every $x \in B_0$ the function $(\varphi_L, \hat{\varphi})$ is a bijection of valuation structures for x . Hence $\tilde{\varphi}_B$ is a bijection $B \rightarrow B$. The relations on \mathbf{B} only depend on the projection, and since $\hat{\varphi}$ is an automorphism, we get that $\tilde{\varphi}$ respects the relations. It remains to prove that for every $F \in \Gamma_L$ and every $(x, \mathbf{V}) \in B$ we have that $\tilde{\varphi}(F_{\mathbf{B}}((x, \mathbf{V}))) = \tilde{\varphi}(F)_{\mathbf{B}}(\tilde{\varphi}((x, \mathbf{V})))$. To make the notation more readable, define function h mapping every valuation structure \mathbf{V} to $(\varphi_L, \hat{\varphi})(\mathbf{V})$.

By the definition of \mathbf{B} we know that

$$F_{\mathbf{B}}((x, \mathbf{V})) = \{(y, \text{Cl}_{\mathbf{V}}(y)) : y \in F_{\mathbf{V}}(x)\},$$

hence

$$\tilde{\varphi}(F_{\mathbf{B}}((x, \mathbf{V}))) = \{(\hat{\varphi}(y), h(\text{Cl}_{\mathbf{V}}(y))) : y \in F_{\mathbf{V}}(x)\}.$$

Denote $X = \tilde{\varphi}(F)_{\mathbf{B}}(\tilde{\varphi}((x, \mathbf{V})))$. By the definition of \mathbf{B} , we have

$$X = \{(y, \text{Cl}_{h(\mathbf{V})}(y)) : y \in \tilde{\varphi}(F)_{h(\mathbf{V})}(\hat{\varphi}(x))\}.$$

Since $\hat{\varphi}$ is a bijection $B \rightarrow B$, we can write

$$X = \{(\hat{\varphi}(y), \text{Cl}_{h(\mathbf{V})}(\hat{\varphi}(y))) : \hat{\varphi}(y) \in \tilde{\varphi}(F)_{h(\mathbf{V})}(\hat{\varphi}(x))\}.$$

Note that $\tilde{\varphi}(F)_{h(\mathbf{V})}(\hat{\varphi}(x)) = \hat{\varphi}(F_{\mathbf{V}}(x))$, hence $\hat{\varphi}(y) \in \tilde{\varphi}(F)_{h(\mathbf{V})}(\hat{\varphi}(x))$ if and only if $y \in F_{\mathbf{V}}(x)$, and so we can write

$$X = \{(\hat{\varphi}(y), \text{Cl}_{h(\mathbf{V})}(\hat{\varphi}(y))) : y \in F_{\mathbf{V}}(x)\}.$$

Finally, $\text{Cl}_{h(\mathbf{V})}(\hat{\varphi}(y)) = h(\text{Cl}_{\mathbf{V}}(y))$, hence indeed

$$X = \tilde{\varphi}(F_{\mathbf{B}}((x, \mathbf{V}))).$$

□

Lemma 4.6.4. *Assume that \mathbf{B}_0 is a coherent EPPA-witness for \mathbf{A}^- and thus $\hat{\varphi}$ can be chosen to be coherent. Let φ_1, φ_2 and φ be partial automorphisms of \mathbf{A} such that $\varphi = \varphi_2 \circ \varphi_1$, and let $\tilde{\varphi}_1, \tilde{\varphi}_2$ and $\tilde{\varphi}$ be their corresponding extensions as above. Then $\tilde{\varphi} = \tilde{\varphi}_2 \circ \tilde{\varphi}_1$.*

Proof. Pick an arbitrary $(x, \mathbf{V}) \in B$. We know that

$$\tilde{\varphi}((x, \mathbf{V})) = (\hat{\varphi}(x), (\varphi_L, \hat{\varphi})(\mathbf{V})).$$

Similarly,

$$\tilde{\varphi}_2 \circ \tilde{\varphi}_1((x, \mathbf{V})) = (\hat{\varphi}_2(\hat{\varphi}_1(x)), (\varphi_L, \text{id}_{\mathbf{V}})((\text{id}_L, \hat{\varphi}_2)((\text{id}_L, \hat{\varphi}_1)(\mathbf{V}))))$$

(as by the assumption, the language parts of φ_1 and φ_2 compose to φ_L and moreover the language permutation commutes with applying $\hat{\varphi}_i$ on the vertex set of \mathbf{V}). By coherence of \mathbf{B}_0 we know that $\hat{\varphi}_2(\hat{\varphi}_1(x)) = \hat{\varphi}(x)$, and so

$$(\varphi_L, \text{id}_{\mathbf{V}})((\text{id}_L, \hat{\varphi}_2)((\text{id}_L, \hat{\varphi}_1)(\mathbf{V})) = (\varphi_L, \hat{\varphi})(\mathbf{V}),$$

hence indeed $\tilde{\varphi} = \tilde{\varphi}_2 \circ \tilde{\varphi}_1$. □

This proves Proposition 4.6.1.

Observation 4.6.5. *The number of vertices of the EPPA-witness \mathbf{B} constructed in this section can be bounded from above by a function which depends only on the number of vertices of \mathbf{B}_0 , the number of vertices of \mathbf{A} and the size of the orbit of the action of Γ_L by relabelling in which \mathbf{A} lies.*

Proof. We know that the vertex set of \mathbf{B} consists of pairs (x, \mathbf{V}) , where $x \in B_0$ and \mathbf{V} is a valuation structure for x . Thus, it is enough to bound the number of valuation structures for any vertex of \mathbf{B}_0 . The vertex set of any valuation structure is a subset of B_0 , hence there are at most $2^{|B_0|}$ different vertex sets of valuation structures. Hence it remains to bound the number of different valuation structures on a given subset $V \subseteq B_0$.

Let $\mathbf{A}_1, \dots, \mathbf{A}_o$ be an enumeration of the orbit of the action of Γ_L by relabeling in which \mathbf{A} lies and let $g_1, \dots, g_o \in \Gamma_L$ such that $\mathbf{A}_i = (g_i, \text{id}_A)(\mathbf{A})$. Note that for every Γ_L -structure \mathbf{U} and every embedding $(f_L, f_U): \mathbf{U} \rightarrow \mathbf{A}$, there is i such that $(f_L, f_U)(\mathbf{U}) = (g_i^{-1}, f_U)(\mathbf{U})$. Indeed, by the way embeddings compose, we have that $(f_L, f_U) = (f_L, \text{id}_A) \circ (\text{id}_L, f_U)$. Composing with (f_L^{-1}, id_A) on the left, we get that (id_L, f_U) is an embedding $\mathbf{U} \rightarrow (f_L^{-1}, \text{id}_A)(\mathbf{A})$. This means that there is i such that $(f_L^{-1}, \text{id}_A)(\mathbf{A}) = \mathbf{A}_i = (g_i, \text{id}_A)(\mathbf{A})$ and hence indeed $(f_L, f_U)(\mathbf{U}) = (g_i^{-1}, f_U)(\mathbf{U})$.

Fix $V \subseteq B_0$. For every valuation structure \mathbf{V} on this vertex set, there is an isomorphism $\iota = (\iota_L, \iota_V): V \rightarrow \text{Cl}_{\mathbf{A}}(y)$, where $y \in A$. Since ι is in particular an embedding $\mathbf{V} \rightarrow \mathbf{A}$, by the previous paragraph we get $1 \leq i \leq o$ such that $(\iota_L, \iota_V)(\mathbf{V}) = (g_i^{-1}, \iota_V)(\mathbf{V})$. Hence there are at most as many different structures \mathbf{V} as there are pairs (g_i^{-1}, ι_V) . Since there are at most o choices for g_i^{-1} and at most $|A|^{|V|} \leq |A|^{|A|}$ choices for ι_V , the claim is proved. \square

From now on, our structures may contain unary functions. To some extent, the unary functions do not interfere too much with the properties which we are going to ensure and thus it is possible to treat them “separately”. Namely, we will always first introduce a notion of a valuation function (in order to get the desired property) and then wrap the valuation functions in a variant of the valuation structures.

4.7 Irreducible structure faithful EPPA

In this section we prove the following proposition, which is a strengthening of [EHN21, Theorem 1.7], which in turn extends [HO03, Theorem 9].

Proposition 4.7.1. *Let L be a language consisting of relation and unary function symbols equipped with a permutation group Γ_L and let \mathbf{A} be a finite Γ_L -structure. Let \mathbf{B}_0 be a finite Γ_L -structure which is an EPPA-witness for \mathbf{A} . Then there is a finite Γ_L -structure \mathbf{B} which is an irreducible structure faithful EPPA-witness for \mathbf{A} , and a homomorphism-embedding $\mathbf{B} \rightarrow \mathbf{B}_0$.*

Moreover, if \mathbf{B}_0 is coherent then \mathbf{B} is coherent, too.

Remark 4.7.1. Note that up to this point, the permutation group Γ_L was not very relevant. In Section 4.4, the constructed EPPA-witnesses worked for Γ_L being the symmetric group, and in Section 4.6, it didn’t play an important role either.

However, in this section Γ_L plays a central role because it restricts which irreducible substructures can be sent to \mathbf{A} by an automorphism. For example, let L be the language consisting of n unary relation R^1, \dots, R^n and let \mathbf{A} be an L -structure consisting of one vertex which is in $R^1_{\mathbf{A}}$ and in no other relations.

For this \mathbf{A} , Section 4.4 will produce an EPPA-witness \mathbf{B}_0 , where $B_0 = \{v_S : S \subseteq \{1, \dots, n\}\}$ such that $v_S \in R^i_{\mathbf{B}_0}$ if and only if $i \in S$. Fix now a permutation group Γ_L on the language L and consider \mathbf{A} as a Γ_L -structure. An irreducible structure faithful EPPA-witness \mathbf{B} for \mathbf{A} will contain only vertices, which are in precisely one unary relation $R^i_{\mathbf{B}}$ such that moreover R^1 and R^i are in the same orbit of Γ_L . In particular, if $\Gamma_L = \{\text{id}_L\}$, then \mathbf{A} is an irreducible structure faithful EPPA-witness for itself. If $\Gamma_L = \text{Sym}(L)$, then a possible irreducible structure faithful EPPA-witness for \mathbf{A} has n vertices v_1, \dots, v_n such that v_i has precisely one unary mark R^i .

Fix a language L consisting of relation symbols and unary function symbols equipped with a permutation group Γ_L . Fix also a finite Γ_L -structure \mathbf{A} and its EPPA-witness \mathbf{B}_0 . Without loss of generality, assume that $\mathbf{A} \subseteq \mathbf{B}_0$. We now present a construction of an irreducible structure faithful EPPA-witness \mathbf{B} with a homomorphism-embedding (projection) to \mathbf{B}_0 , such that every extension of a partial automorphism in \mathbf{B} is induced by the extension of its projection to \mathbf{B}_0 .

Witness construction

Let \mathbf{I} be an irreducible substructure of \mathbf{B}_0 . We say that \mathbf{I} is *bad* if there is no automorphism $f: \mathbf{B}_0 \rightarrow \mathbf{B}_0$ such that $f(\mathbf{I}) \subseteq \mathbf{A}$. Given a vertex $x \in B_0$, we denote by $U(x)$ the set of all bad irreducible substructures of \mathbf{B}_0 containing x .

For a vertex $x \in B_0$, we say that a function assigning to every $\mathbf{I} \in U(x)$ a value from $\{1, \dots, |I| - 1\}$ is a *valuation function for x* . Given vertices $x, y \in B_0$ and their valuation functions χ and χ' respectively, we say that the pairs (x, χ) and (y, χ') are *generic*, if either $(x, \chi) = (y, \chi')$, or $x \neq y$ and for every $\mathbf{I} \in U(x) \cap U(y)$ it holds that $\chi(\mathbf{I}) \neq \chi'(\mathbf{I})$. We say that a set S is *generic* if it consists of pairs (x, χ) where $x \in B_0$ and χ is a valuation function for x , and every pair $(x, \chi), (y, \chi') \in S$ is generic. In particular, the projection to the first coordinate is injective on every generic set.

A *valuation structure* for a vertex $x \in B_0$ is a Γ_L -structure \mathbf{V} such that:

1. The vertex set of \mathbf{V} is a generic set of pairs (y, χ) with $y \in \text{Cl}_{\mathbf{B}_0}(x)$ and χ being a valuation function for y , and
2. the pair $\iota = (\text{id}_L, \iota_V)$, where $\iota_V((y, \chi)) = y$, is an isomorphism of \mathbf{V} and $\text{Cl}_{\mathbf{B}_0}(x)$.

For a pair (x, \mathbf{V}) , where $x \in B_0$ and \mathbf{V} is a valuation structure for x , we denote by $\chi(x, \mathbf{V})$ the (unique) valuation function for x such that $(x, \chi(x, \mathbf{V})) \in V$ and we put $\pi(x, \mathbf{V}) = x$ (π is again the *projection* from \mathbf{B} to \mathbf{B}_0). We say that a set S of pairs (x, \mathbf{V}) , such that $x \in B_0$ and \mathbf{V} is a valuation structure for x , is *generic*, if the union $\bigcup_{(x, \mathbf{V}) \in S} V$ is generic. Note that this implies that in particular $\{(x, \chi(x, \mathbf{V})) : (x, \mathbf{V}) \in S\}$ is generic and thus π is injective on every generic set.

Observe that if L contains no functions then every valuation structure \mathbf{V} for $x \in B_0$ contains exactly one vertex $(x, \chi(x, \mathbf{V}))$, and conversely, for every

valuation function χ for x there is exactly one valuation structure \mathbf{V} for x such that $\chi(x, \mathbf{V}) = \chi$.

Now we construct a Γ_L -structure \mathbf{B} :

1. The vertices of \mathbf{B} are all pairs (x, \mathbf{V}) , where $x \in B_0$ and \mathbf{V} is a valuation structure for x .
2. For every relation symbol $R \in L_{\mathcal{R}}$, we put

$$((x_1, \mathbf{V}_1), \dots, (x_{a(R)}, \mathbf{V}_{a(R)})) \in R_{\mathbf{B}},$$

if and only if $(x_1, \dots, x_{a(R)}) \in R_{\mathbf{B}_0}$, and $\{(x_1, \mathbf{V}_1), \dots, (x_{a(R)}, \mathbf{V}_{a(R)})\}$ is generic.

3. For every (unary) function symbol $F \in L_{\mathcal{F}}$, we put

$$F_{\mathbf{B}}((x, \mathbf{V})) = \{(y, \text{Cl}_{\mathbf{V}}((y, \chi))) : (y, \chi) \in F_{\mathbf{V}}((x, \chi(x, \mathbf{V})))\}.$$

Note that in the definition of $F_{\mathbf{B}}((x, \mathbf{V}))$ it holds that $\text{Cl}_{\mathbf{V}}((y, \chi))$ is isomorphic to $\text{Cl}_{\mathbf{B}_0}(y)$, so it is indeed a valuation structure for y . Also observe that \mathbf{B} is finite, because B_0 is finite, $U(x)$ is finite for every $x \in B_0$ (hence there are only finitely many candidate vertex sets for valuations structures), and there is at most one valuation structure on any candidate vertex set.

The following claim (whose proof is quite technical and will be given at the end of this section) justifies our definition of genericity and the construction of \mathbf{B} .

Claim 4.7.2. *Let $\mathbf{D} \subseteq \mathbf{B}$ be irreducible. Then D is generic.*

We also have this complementary fact to Claim 4.7.2, which will be useful several times in this section.

Claim 4.7.3. *Let $\mathbf{D} \subseteq \mathbf{B}$ be such that D is a generic set. Then the restriction of (id_L, π) to \mathbf{D} is an embedding $\mathbf{D} \rightarrow \mathbf{B}_0$.*

Next we define an embedding $\psi: \mathbf{A} \rightarrow \mathbf{B}$ with $\psi_L = \text{id}$. For every bad irreducible $\mathbf{I} \subseteq \mathbf{B}_0$, we fix an arbitrary injective function $u_{\mathbf{I}}: I \cap A \rightarrow \{1, 2, \dots, |I| - 1\}$. Such a function exists, because $A \cap I$ is a proper subset of I (otherwise \mathbf{I} would not be bad). For every $x \in A$ we define a valuation function χ_x for x such that $\chi_x(\mathbf{I}) = u_{\mathbf{I}}(x)$.

Given $x \in A$, we also define a valuation structure \mathbf{V}_x for x such that $V_x = \{(y, \chi_y) : y \in \text{Cl}_{\mathbf{A}}(x)\}$, and the structure on V_x is chosen such that the pair $(\text{id}_L, (y, \chi_y) \mapsto y)$ is an isomorphism $\mathbf{V}_x \rightarrow \text{Cl}_{\mathbf{A}}(x)$. We put $\psi(x) = (x, \mathbf{V}_x)$.

Claim 4.7.4. *ψ is an embedding $\mathbf{A} \rightarrow \mathbf{B}$ and $\mathbf{A}' = \psi(\mathbf{A})$ is generic.*

Constructing the extension

At some point, we will also need to prove irreducible structure faithfulness. And in that proof, we are going to need to construct some automorphisms of \mathbf{B} based on some automorphisms of \mathbf{B}_0 and partial automorphisms of \mathbf{B} . Because of it, we will prove a more general statement

Lemma 4.7.5. *Let φ be a partial automorphism of \mathbf{B} satisfying the following conditions:*

1. *Both the domain and the range of φ are generic, and*
2. *there is an automorphism $\hat{\varphi}$ of \mathbf{B}_0 which extends the projection of φ via π .*

Then there is an automorphism $\tilde{\varphi}$ of \mathbf{B} extending φ .

Note that if φ is a partial automorphism of \mathbf{A}' , then it satisfies both conditions (as \mathbf{A}' is generic) and therefore it can be extended to an automorphism of \mathbf{B} . Note also that in condition 2, the projection of φ via π is a partial automorphism of \mathbf{B}_0 by Claim 4.7.3.

Proof of Lemma 4.7.5. Put

$$D = \{(x, \chi(x, \mathbf{V})) : (x, \mathbf{V}) \in \text{Dom}(\varphi)\}$$

and

$$R = \{(x, \chi(x, \mathbf{V})) : (x, \mathbf{V}) \in \text{Range}(\varphi)\}.$$

Because both $\text{Dom}(\varphi)$ and $\text{Range}(\varphi)$ are generic, we get that $|D| = |R| = |\pi(D)| = |\pi(R)|$, so in particular no $x \in B_0$ appears in D or R with more than one valuation structure. Therefore, φ defines a bijection $q: D \rightarrow R$.

For a bad irreducible substructure $\mathbf{I} \subseteq \mathbf{B}_0$, we can define a partial permutation $\tau_{\mathbf{I}}^{\varphi}$ of $\{1, \dots, |I| - 1\}$, such that for every $(y, \chi) \in D$ with $y \in I$ and for $q((y, \chi)) = (\hat{\varphi}(y), \chi')$, we put

$$\tau_{\mathbf{I}}^{\varphi}(\chi(\mathbf{I})) = \chi'(\hat{\varphi}(\mathbf{I})).$$

This is indeed a partial permutation of $\{1, \dots, |I| - 1\}$, because both D and R are generic. Let $\hat{\tau}_{\mathbf{I}}^{\varphi}$ be the order-preserving extension of $\tau_{\mathbf{I}}^{\varphi}$.

Put

$$\mathcal{V} = \bigcup_{(x, \mathbf{V}) \in \mathbf{B}} V.$$

Having $\hat{\tau}_{\mathbf{I}}^{\varphi}$ for every bad \mathbf{I} , we can define $\hat{q}: \mathcal{V} \rightarrow \mathcal{V}$ as

$$\hat{q}((x, \chi)) = (\hat{\varphi}(x), \chi'),$$

where $\chi'(\hat{\varphi}(\mathbf{I})) = \hat{\tau}_{\mathbf{I}}^{\varphi}(\chi(\mathbf{I}))$. Since $\hat{\varphi}$ is an automorphism of \mathbf{B}_0 and each $\hat{\tau}_{\mathbf{I}}^{\varphi}$ is a permutation of $\{1, \dots, |I| - 1\}$, it follows that \hat{q} is a permutation of \mathcal{V} . It is easy to check that \hat{q} extends q .

Finally, we define $\tilde{\varphi}: B \rightarrow B$ by putting $\tilde{\varphi}_L = \varphi_L$ and

$$\tilde{\varphi}((x, \mathbf{V})) = (\hat{\varphi}(x), (\varphi_L, \hat{q})(\mathbf{V})).$$

The proof of the following claim, which will be given at the end of this section, is simply a mechanical verification that our constructions are well-defined.

Claim 4.7.6. *$\tilde{\varphi}$ is an automorphism of \mathbf{B} extending φ .*

□

Proofs

We again proceed analogously to Section 4.3.

Lemma 4.7.7. *\mathbf{B} is an EPPA-witness for \mathbf{A} . Moreover, if \mathbf{B}_0 is coherent then so is \mathbf{B} . If \mathbf{B}_0 is a coherent EPPA-witness for \mathbf{A} , then \mathbf{B} is a coherent EPPA-witness for \mathbf{A}' .*

Proof. Lemma 4.7.5 implies that \mathbf{B} is indeed an EPPA-witness for \mathbf{A} , because A' is generic. We thus focus on proving coherence. Let φ_1, φ_2 and φ be a coherent triple of partial automorphisms of \mathbf{A}' and let $\hat{\varphi}_1, \hat{\varphi}_2, \hat{\varphi}$ be automorphisms of \mathbf{B}_0 which are the coherent extensions of the projections of φ_1, φ_2 and φ by π .

Denote by $\hat{q}_1, \hat{q}_2, \hat{q}$ and $\tau_{\mathbf{I}}^\varphi, \tau_{\mathbf{I}}^{\varphi_1}$ and $\tau_{\mathbf{I}}^{\varphi_2}$ for every \mathbf{I} the corresponding functions from proof of Lemma 4.7.5. Coherence on the first coordinate follows from coherence of $\hat{\varphi}_1, \hat{\varphi}_2, \hat{\varphi}$, to get coherence on the second coordinate, we need to prove that $\hat{q} = \hat{q}_2 \circ \hat{q}_1$. To see that, one only needs to prove that $\tau_{\mathbf{I}}^\varphi = \tau_{\mathbf{I}}^{\varphi_2} \circ \tau_{\mathbf{I}}^{\varphi_1}$, which is true as all of them are extended in an order-preserving way. \square

Proof of Proposition 4.7.1. First we prove that (id_L, π) is a homomorphism-embedding from \mathbf{B} to \mathbf{B}_0 . From the construction it directly follows that it is a homomorphism which preserves functions. Let \mathbf{I} be an irreducible substructure of \mathbf{B} . By Claim 4.7.2 we get that \mathbf{I} is generic and hence by Claim 4.7.3 we get that (id_L, π) is an embedding on \mathbf{I} .

It only remains to prove irreducible structure faithfulness of \mathbf{B} . Let \mathbf{D} be an irreducible substructure of \mathbf{B} . By Claim 4.7.2 we get that \mathbf{D} is generic and hence $\pi(\mathbf{D})$ is not a bad substructure of \mathbf{B}_0 . It is, however, irreducible, because π is a homomorphism-embedding, and thus there is $\hat{\varphi} \in \text{Aut}(\mathbf{B}_0)$ such that $\hat{\varphi}(\pi(\mathbf{D})) \subseteq A$. Define $\varphi: \mathbf{D} \rightarrow \mathbf{A}'$ by $\varphi_D((x, \mathbf{V})) = \psi(\hat{\varphi}(x))$ and $\varphi_L = \hat{\varphi}_L$. This is a partial automorphism of \mathbf{B} with generic domain and range, whose projection extends to $\hat{\varphi}$. Lemma 4.7.5 then gives an automorphism of \mathbf{B} sending D to A' , which is what we wanted. \square

We are now ready to prove Theorem 4.1.1.

Proof of Theorem 4.1.1. Section 4.6 gives a finite coherent EPPA-witness \mathbf{B}_0 for \mathbf{A} , Proposition 4.7.1 ensures irreducible structure faithfulness and preserves coherence. The “consequently” part is immediate. \square

Remark 4.7.2. Note that Theorem 4.1.1 can be used to prove EPPA for classes where the relations are, for example, symmetric (because a non-symmetric relation is witnessed on a tuple which is irreducible), and similarly it implies EPPA for classes with unary functions whose range has a given size (for example, size 1, which means that we can prove EPPA for the standard model-theoretic unary functions).

Remark 4.7.3. Note that our partial automorphism extension actually has some functorial properties. Taking isomorphic copies, we can assume that $\mathbf{A} \subseteq \mathbf{B}_0$ and $\mathbf{A} \subseteq \mathbf{B}$ (then in particular $\pi \upharpoonright_A = \text{id}_A$). Given a partial automorphism φ of \mathbf{A} let $\hat{\varphi}$ be its (coherent) extension to an automorphism of \mathbf{B}_0 and let $\tilde{\varphi}$ be the constructed extension to an automorphism of \mathbf{B} using $\hat{\varphi}$. Let \sim be an equivalence relation on B given by $x \sim y \iff \pi(x) = \pi(y)$. Then \sim is a congruence with respect to $\tilde{\varphi}$ and the natural actions of $\hat{\varphi}$ and $\tilde{\varphi}$ on the equivalence classes of \sim coincide.

Observation 4.7.8. *The number of vertices of the EPPA-witness \mathbf{B} constructed in this section can be bounded from above by a function which depends only on the number of vertices of \mathbf{B}_0 and the number of vertices of \mathbf{A} .*

Proof. Put $m = |B_0|$ and $n = |A|$. There are at most 2^m bad substructures of \mathbf{B}_0 and hence at most $(m-1)^{2^m}$ valuation functions for a given $x \in B_0$ (this is a very rough estimate). Let P be the set of all pairs (x, χ) , where $x \in B_0$ and χ is a valuation function for x . We get that $|P| \leq m(m-1)^{2^m}$.

The vertices of \mathbf{B} are pairs (x, \mathbf{V}) , where $x \in B_0$ and \mathbf{V} is a valuation structure for x . The vertex set of every valuation structure is a subset of P (and hence there are at most $2^{|P|}$ of them) and the structure on \mathbf{V} is determined by an embedding $(\text{id}_L, \iota_V): \mathbf{V} \rightarrow \mathbf{B}_0$. There are at most $m^{|\mathbf{V}|} \leq m^m$ such embeddings. This finishes the proof. \square

Proof of Claim 4.7.2. For a contradiction, suppose that it is not the case, that is, there are $(u, \chi), (u', \chi') \in \bigcup_{(x, \mathbf{V}) \in D} V$ which form a non-generic pair. This implies that there are $(x, \mathbf{X}), (y, \mathbf{Y}) \in D$ such that $(u, \chi) \in X$ and $(u', \chi') \in Y$ (they cannot both lie in the same valuation structure, because the vertex sets of valuation structures are generic), and hence the set $\{(x, \mathbf{X}), (y, \mathbf{Y})\}$ is non-generic. Put $\mathbf{E}_x = \{(a, \mathbf{U}) \in D : (x, \mathbf{X}) \notin \text{Cl}_{\mathbf{D}}((a, \mathbf{U}))\}$ and similarly define \mathbf{E}_y . Since closures are unary, these are substructures of \mathbf{D} . Note that as closures in \mathbf{B} are generic, we also know that $(y, \mathbf{Y}) \in \mathbf{E}_x$ and $(x, \mathbf{X}) \in \mathbf{E}_y$. This means that $\mathbf{E}_x, \mathbf{E}_y$ are both non-empty and neither is a substructure of the other.

We first prove $\mathbf{E}_x \cup \mathbf{E}_y = \mathbf{D}$. Suppose for a contradiction that there is $(z, \mathbf{Z}) \in \mathbf{D}$ with $\{(x, \mathbf{X}), (y, \mathbf{Y})\} \subseteq \text{Cl}_{\mathbf{D}}((z, \mathbf{Z}))$. Then (by the construction of \mathbf{B}) we have $\mathbf{X}, \mathbf{Y} \subseteq \mathbf{Z}$, which is a contradiction with $(x, \mathbf{X}), (y, \mathbf{Y})$ forming a non-generic pair.

Fix $(a, \mathbf{U}) \in \mathbf{E}_x \setminus \mathbf{E}_y$ and $(b, \mathbf{W}) \in \mathbf{E}_y \setminus \mathbf{E}_x$. Because we know that $\mathbf{Y} \subseteq \mathbf{U}$ and $\mathbf{X} \subseteq \mathbf{W}$, we get that $(a, \mathbf{U}), (b, \mathbf{W})$ is not a generic pair and therefore no relation of \mathbf{D} contains both (a, \mathbf{U}) and (b, \mathbf{W}) . Thus \mathbf{D} is a free amalgam of \mathbf{E}_x and \mathbf{E}_y over their intersection, which is a contradiction with its irreducibility. Therefore D is indeed generic. \square

Proof of Claim 4.7.3. We have already observed that π is injective. The fact that π preserves relations and non-relations follows directly from the construction of \mathbf{B} . Let $F \in L$ be a function and fix $(x, \mathbf{V}) \in D$. We need to prove that $F_{\mathbf{B}_0}(x)$ is equal to $\pi(F_{\mathbf{B}}((x, \mathbf{V})))$.

By definition,

$$\pi(F_{\mathbf{B}}((x, \mathbf{V}))) = \{y : (y, \chi) \in F_{\mathbf{V}}((x, \chi(x, \mathbf{V})))\}.$$

Moreover, we know that the pair $\iota = (\text{id}_L, \iota_V)$, where $\iota_V((y, \chi)) = y$, is an isomorphism of \mathbf{V} and $\text{Cl}_{\mathbf{B}_0}(x)$. Hence if $(y, \chi) \in F_{\mathbf{V}}((x, \chi(x, \mathbf{V})))$ then $y \in F_{\mathbf{B}_0}(x)$ and conversely, whenever $y \in F_{\mathbf{B}_0}(x)$ then there is χ such that $(y, \chi) \in F_{\mathbf{V}}((x, \chi(x, \mathbf{V})))$ (and since \mathbf{V} is generic, such χ is uniquely determined). So

$$\pi(F_{\mathbf{B}}((x, \mathbf{V}))) = \{y : y \in F_{\mathbf{B}_0}(x)\} = F_{\mathbf{B}_0}(x).$$

This concludes the proof. \square

Proof of Claim 4.7.4. First we will prove that A' is a generic set. By definition this happens if $V = \bigcup_{(x, \mathbf{V}_x) \in A'} V_x = \bigcup_{x \in A} V_x$ is a generic set. Note that $V =$

$\{(x, \chi_x) : x \in A\}$. Let $x \neq y \in A$ be arbitrary and pick $\mathbf{I} \in U(x) \cap U(y)$. By the construction we have $\chi_x(\mathbf{I}) = u_{\mathbf{I}}(x) \neq u_{\mathbf{I}}(y) = \chi_y(\mathbf{I})$, hence (x, χ_x) and (y, χ_y) form a generic pair which implies that V is indeed a generic set. From this it follows that in particular every V_x is a generic set and hence ψ is a function $A \rightarrow B$.

Next fix a relation $R \in L$ and a tuple $\bar{x} \in A^n$. We will prove that $\psi(\bar{x}) \in R_{\mathbf{B}}$ if and only if $\bar{x} \in R_{\mathbf{A}}$. By the definition of \mathbf{B} the “only if” part is immediate, to prove the other implication, we need to prove that $\psi(\bar{x})$ (understood as a set) is generic, but clearly $\psi(\bar{x})$ is a subset of a generic set V , which concludes the proof.

Finally we prove that for every (unary) function $F \in L$ and every vertex $x \in \mathbf{A}$ we have $\psi(F_{\mathbf{A}}(x)) = F_{\mathbf{B}}(\psi(x))$ (remember that $\psi_L = \text{id}_L$). Clearly

$$\psi(F_{\mathbf{A}}(x)) = \{(y, \mathbf{V}_y) : y \in F_{\mathbf{A}}(x)\}.$$

Put $X = F_{\mathbf{B}}(\psi(x))$. By the construction we have

$$X = F_{\mathbf{B}}((x, \mathbf{V}_x)) = \{(y, \text{Cl}_{\mathbf{V}_x}((y, \chi))) : (y, \chi) \in F_{\mathbf{V}_x}((x, \chi(x, \mathbf{V}_x)))\}.$$

Note that $\chi(x, \mathbf{V}_x) = \chi_x$ and that if $(y, \chi) \in F_{\mathbf{V}_x}((x, \chi(x, \mathbf{V}_x)))$, then $\chi = \chi_y$. Hence, in particular, π is injective on $F_{\mathbf{V}_x}((x, \chi_x))$. So we can write

$$X = \{(y, \text{Cl}_{\mathbf{V}_x}((y, \chi_y))) : (y, \chi_y) \in F_{\mathbf{V}_x}((x, \chi_x))\}.$$

Since $(\text{id}_L, (y, \chi_y) \mapsto y)$ is an isomorphism $\mathbf{V}_x \rightarrow \text{Cl}_{\mathbf{A}}(x)$, we get that $(y, \chi_y) \in F_{\mathbf{V}_x}((x, \chi_x))$ if and only if $y \in F_{\text{Cl}_{\mathbf{A}}(x)}(x) = F_{\mathbf{A}}(x)$. For the same reason, $\text{Cl}_{\mathbf{V}_x}((y, \chi_y))$ is isomorphic to $\text{Cl}_{\mathbf{A}}(y)$ by projecting to the first coordinate and hence in fact $\text{Cl}_{\mathbf{V}_x}((y, \chi_y)) = \mathbf{V}_y$. Putting this together, we get that

$$F_{\mathbf{B}}(\psi(x)) = X = \{(y, \mathbf{V}_y) : y \in F_{\mathbf{A}}(x)\} = \psi(F_{\mathbf{A}}(x)).$$

□

Proof of Claim 4.7.6. It is easy to see that $\tilde{\varphi}$ is a bijection which maps generic sets to generic sets. Fix a relation $R \in L$ and a tuple $((x_1, \mathbf{V}_1), \dots, (x_n, \mathbf{V}_n)) \in B^n$. Note that

$$(x_1, \dots, x_n) \in R_{\mathbf{B}_0} \iff (\hat{\varphi}(x_1), \dots, \hat{\varphi}(x_n)) \in \tilde{\varphi}(R)_{\mathbf{B}_0},$$

because $\hat{\varphi}$ is an automorphism of \mathbf{B}_0 . Together with the fact that $\tilde{\varphi}$ maps generic sets to generic sets it follows that $((x_1, \mathbf{V}_1), \dots, (x_n, \mathbf{V}_n)) \in R_{\mathbf{B}}$ if and only if $(\tilde{\varphi}((x_1, \mathbf{V}_1)), \dots, \tilde{\varphi}((x_n, \mathbf{V}_n))) \in \tilde{\varphi}(R)_{\mathbf{B}}$.

It remains to prove that for every function $F \in L$ and every $(x, \mathbf{V}) \in B$ we have $\tilde{\varphi}(F_{\mathbf{B}}((x, \mathbf{V}))) = \varphi_L(F)_{\mathbf{B}}(\tilde{\varphi}((x, \mathbf{V})))$. To simplify the notation we put $h(\mathbf{V}) = (\varphi_L, \hat{q})(\mathbf{V})$ for every valuation structure \mathbf{V} . Fix an arbitrary $(x, \mathbf{V}) \in \mathbf{B}$ and put $\chi_0 = \chi(x, \mathbf{V})$. By the definition of \mathbf{B} we know that

$$\tilde{\varphi}(F_{\mathbf{B}}((x, \mathbf{V}))) = \{(\hat{\varphi}(y), h(\text{Cl}_{\mathbf{V}}((y, \chi)))) : (y, \chi) \in F_{\mathbf{V}}((x, \chi_0))\}.$$

Denote $X = \tilde{\varphi}(F)_{\mathbf{B}}(\tilde{\varphi}((x, \mathbf{V})))$. By the definition of \mathbf{B} , we have

$$X = \{(y, \text{Cl}_{h(\mathbf{V})}((y, \chi))) : (y, \chi) \in \tilde{\varphi}(F)_{h(\mathbf{V})}(\hat{q}((x, \chi_0)))\}.$$

Since \hat{q} is a bijection $\mathcal{V} \rightarrow \mathcal{V}$ which agrees with $\hat{\varphi}$ on the first coordinate, we can write

$$X = \{(\hat{\varphi}(y), \text{Cl}_{h(\mathbf{V})}(\hat{q}((y, \chi)))) : \hat{q}((y, \chi)) \in \tilde{\varphi}(F)_{h(\mathbf{V})}(\hat{q}((x, \chi_0)))\}.$$

Note that

$$\tilde{\varphi}(F)_{h(\mathbf{V})}(\hat{q}((x, \chi_0))) = \hat{q}(F_{\mathbf{V}}((x, \chi_0))),$$

hence $\hat{q}((y, \chi)) \in \tilde{\varphi}(F)_{h(\mathbf{V})}(\hat{q}((x, \chi_0)))$ if and only if $(y, \chi) \in F_{\mathbf{V}}((x, \chi_0))$, and so we have

$$X = \{(\hat{\varphi}(y), \text{Cl}_{h(\mathbf{V})}(\hat{q}((y, \chi)))) : (y, \chi) \in F_{\mathbf{V}}((x, \chi_0))\}.$$

Finally, $\text{Cl}_{h(\mathbf{V})}(\hat{q}((y, \chi))) = h(\text{Cl}_{\mathbf{V}}((y, \chi)))$, hence indeed

$$X = \tilde{\varphi}(F_{\mathbf{B}}((x, \mathbf{V}))).$$

□

4.8 Unwinding induced cycles

In this section we give a key ingredient for proving Theorem 4.1.2:

Lemma 4.8.1. *Let L be a language consisting of relations and unary functions equipped with a permutation group Γ_L . Let \mathbf{A} be a finite irreducible Γ_L -structure and let \mathbf{B}_0 be its (finite) irreducible structure faithful EPPA-witness. Assume that L contains a binary relation E which is fixed by every permutation in Γ_L and assume that $E_{\mathbf{A}}$ is a complete graph. (Note that by irreducible structure faithfulness $E_{\mathbf{B}_0}$ is an undirected graph without loops.)*

There is a finite Γ_L -structure \mathbf{B} which is an irreducible structure faithful EPPA-witness for \mathbf{A} satisfying the following:

1. *There is a homomorphism-embedding $f: \mathbf{B} \rightarrow \mathbf{B}_0$.*
2. *Let C be a subset of B . Then at least one of the following holds:*
 - (a) *$E_{\mathbf{B}} \cap C^2$ contains no (induced) cycle of length ≥ 4 ,*
 - (b) *$|f(C)| < |C|$, or*
 - (c) *$|E_{\mathbf{B}_0} \cap f(C)^2| > |E_{\mathbf{B}} \cap C^2|$.*

Moreover, if \mathbf{B}_0 is a coherent EPPA-witness for \mathbf{A} , then \mathbf{B} is also coherent.

Note that since f is a homomorphism-embedding, we get that if $|f(C)| = |C|$, then $|E_{\mathbf{B}_0} \cap f(C)^2| \geq |E_{\mathbf{B}} \cap C^2|$ and f induces an injective mapping from $E_{\mathbf{B}} \cap C^2$ to $E_{\mathbf{B}_0} \cap f(C)^2$.

In the rest of the section, we will prove Lemma 4.8.1. The construction is inspired by a similar construction for EPPA for metric spaces by the authors [HKN19a].

For the rest of the section, fix L , Γ_L , \mathbf{A} and \mathbf{B}_0 as in the statement of Lemma 4.8.1. Assume without loss of generality that $\mathbf{A} \subseteq \mathbf{B}_0$.

Valuations

A sequence (c_1, \dots, c_k) of distinct vertices of \mathbf{B}_0 is a *bad cycle sequence* if $k \geq 4$ and the structure induced by $E_{\mathbf{B}_0}$ on $\{c_1, \dots, c_k\}$ is a graph cycle containing precisely the edges connecting c_i and c_{i+1} for every $1 \leq i \leq k$ (where we identify $c_{k+1} = c_1$).

Given a vertex $x \in B_0$, we denote by $U(x)$ the set of all bad cycle sequences containing x . We call functions $U(x) \rightarrow \{0, 1\}$ *valuation functions for x* . Given vertices $x, y \in B_0$ and their valuation functions χ and χ' , we say that the pairs (x, χ) and (y, χ') are *generic*, if either $(x, \chi) = (y, \chi')$, or $x \neq y$ and for every bad cycle sequence $\vec{c} = (c_1, \dots, c_k) \in U(x) \cap U(y)$, one of the following holds:

1. There is $1 \leq i < k$ such that $\{c_i, c_{i+1}\} = \{x, y\}$ and $\chi(\vec{c}) = \chi'(\vec{c})$, or
2. $\{c_1, c_k\} = \{x, y\}$ and $\chi(\vec{c}) \neq \chi'(\vec{c})$.

A set S of pairs (x, χ) is *generic* if every pair $(x, \chi), (y, \chi') \in S$ is generic.

Let $x \in B_0$ be a vertex of \mathbf{B}_0 . A *valuation structure for x* is a Γ_L -structure \mathbf{V} such that:

1. The vertex set V is a generic set of pairs (y, χ) where $y \in \text{Cl}_{\mathbf{B}_0}(x)$ and χ is a valuation function for y .
2. The pair $(\text{id}_L, (y, \chi) \mapsto y)$ is an isomorphism of \mathbf{V} and $\text{Cl}_{\mathbf{B}_0}(x)$.

Let \mathbf{V} be a valuation structure for x . We denote by $\chi(x, \mathbf{V})$ the valuation function for x such that $(x, \chi(x, \mathbf{V})) \in V$. Similarly as in Section 4.7, if L contains no functions then every valuation structure \mathbf{V} for $x \in B_0$ contains exactly one vertex $(x, \chi(x, \mathbf{V}))$ and conversely, for every valuation function χ for x there is exactly one valuation structure \mathbf{V} for x such that $\chi(x, \mathbf{V}) = \chi$.

A set S of pairs (x, \mathbf{V}) , where \mathbf{V} is a valuation structure for x , is *generic*, if the union $\bigcup_{(x, \mathbf{V}) \in S} V$ is generic.

Witness construction

Now we construct a Γ_L -structure \mathbf{B} :

1. The vertices of \mathbf{B} are all pairs (x, \mathbf{V}) where $x \in B_0$ and \mathbf{V} is a valuation structure for x .
2. For every relation symbol $R \in L_{\mathcal{R}}$, we put

$$((x_1, \mathbf{V}_1), \dots, (x_{a(R)}, \mathbf{V}_{a(R)})) \in R_{\mathbf{B}},$$

if and only if $(x_1, \dots, x_{a(R)}) \in R_{\mathbf{B}_0}$, and $\{(x_1, \mathbf{V}_1), \dots, (x_{a(R)}, \mathbf{V}_{a(R)})\}$ is generic.

3. for every (unary) function symbol $F \in L_{\mathcal{F}}$ and every vertex $(x, \mathbf{V}) \in B$, we put

$$F_{\mathbf{B}}((x, \mathbf{V})) = \{(y, \text{Cl}_{\mathbf{V}}((y, \chi))) : (y, \chi) \in F_{\mathbf{V}}((x, \chi(x, \mathbf{V})))\}.$$

Claim 4.8.2. *If \mathbf{D} is an irreducible substructure of \mathbf{B} , then \mathbf{D} is generic.*

Define $\pi_B(x, \mathbf{V}) = x$ and $\pi_L = \text{id}_L$. We then have the following:

Claim 4.8.3. *\mathbf{B} is a finite Γ_L -structure and π is a homomorphism-embedding from \mathbf{B} to \mathbf{B}_0 which is an embedding on every generic $D \subseteq B$.*

Observe that since every pair of distinct vertices $x, y \in A$ is in $R_{\mathbf{A}}^E$, it follows that every bad cycle sequence contains at most two vertices of \mathbf{A} , and if it contains precisely two, then they are adjacent in $E_{\mathbf{B}_0}$. For every bad cycle sequence $\vec{c} = (c_1, \dots, c_k)$ containing at least one vertex of \mathbf{A} , we define a function $\chi_{\vec{c}}: A \cap \{c_1, \dots, c_k\} \rightarrow \{0, 1\}$ as follows.

$$\chi_{\vec{c}}(x) = \begin{cases} 1 & \text{if } x = c_1 \text{ and } c_k \in A, \\ 0 & \text{otherwise.} \end{cases}$$

Next we give an embedding $\psi: \mathbf{A} \rightarrow \mathbf{B}$. Given a vertex $x \in A$, we define a valuation function χ_x for x , putting $\chi_x(\vec{c}) = \chi_{\vec{c}}(x)$ for every $\vec{c} \in U(x)$, and we define a valuation structure \mathbf{V}_x for x with $V_x = \{(y, \chi_y) : y \in \text{Cl}_{\mathbf{A}}(x)\}$ such that π restricted to V_x is an isomorphism $\mathbf{V}_x \rightarrow \text{Cl}_{\mathbf{A}}(x)$ (which is a substructure of \mathbf{A}). We put $\psi_A(x) = (x, \mathbf{V}_x)$ and $\psi_L = \text{id}$.

Claim 4.8.4. *ψ is an embedding $\mathbf{A} \rightarrow \mathbf{B}$ and $\mathbf{A}' = \psi(\mathbf{A})$ is generic.*

Constructing the extension

Similarly as in the last section, we will need to prove irreducible structure faithfulness and the following slightly more general extension lemma will be useful in proving that.

Lemma 4.8.5. *Let φ be a partial automorphism of \mathbf{B} satisfying the following conditions:*

1. *Both the domain and the range of φ are generic, and*
2. *there is an automorphism $\hat{\varphi}$ of \mathbf{B}_0 which extends the projection of φ via π .*

Then there is an automorphism $\tilde{\varphi}$ of \mathbf{B} extending φ .

Note that by Claim 4.8.3, π is an embedding on generic sets, hence the projection of φ via π is a partial automorphism of \mathbf{B}_0 .

Proof. Let F be the set consisting of all bad cycle sequences \vec{c} for which there is a vertex $(x, \mathbf{V}) \in \text{Dom}(\varphi)$ such that $\chi(x, \mathbf{V})(\vec{c}) \neq \chi(\varphi((x, \mathbf{V}))) (\hat{\varphi}(\vec{c}))$.

We define a function f such that if χ is a valuation function for x then $f(\chi)$ is a valuation for x satisfying

$$f(\chi)(\hat{\varphi}(\vec{c})) = \begin{cases} \chi(\vec{c}) & \text{if } \vec{c} \notin F \\ 1 - \chi(\vec{c}) & \text{if } \vec{c} \in F. \end{cases}$$

Put $\mathcal{V} = \bigcup_{(x, \mathbf{V}) \in \mathbf{B}} V$. Next we define a function $\hat{q}: \mathcal{V} \rightarrow \mathcal{V}$ putting

$$\hat{q}((x, \chi)) = (\hat{\varphi}(x), f(\chi)),$$

and using it we construct the extension $\tilde{\varphi}$ such that $\tilde{\varphi}_L = \varphi_L$ and $\tilde{\varphi}((x, \mathbf{V})) = (\hat{\varphi}(x), (\varphi_L, \hat{q})(\mathbf{V}))$.

The proof of the following claim, which will be given at the end of this section, is simply a mechanical verification that our constructions are well-defined.

Claim 4.8.6. $\tilde{\varphi}$ is an automorphism of \mathbf{B} extending φ .

□

Proofs

Lemma 4.8.7. *The following statements about \mathbf{B} are true:*

1. \mathbf{B} is irreducible structure faithful.
2. If \mathbf{B}_0 is a coherent EPPA-witness for \mathbf{A} , then \mathbf{B} is a coherent EPPA-witness for \mathbf{A}' .

Proof. To prove irreducible structure faithfulness, let \mathbf{I} be an irreducible substructure of \mathbf{B} . By Claim 4.8.2, I is generic and hence π is an embedding on \mathbf{I} (Claim 4.8.3). This means that $\pi(\mathbf{I})$ is an irreducible substructure of \mathbf{B}_0 and thus there is an automorphism $\hat{\varphi}$ of \mathbf{B}_0 sending $\pi(\mathbf{I})$ to A . Put φ to be the partial automorphism of \mathbf{B} sending \mathbf{I} to $\psi(\hat{\varphi}(\pi(\mathbf{I})))$ with $\varphi_L = \hat{\varphi}_L$. This is a partial automorphism of \mathbf{B} with generic domain and range (using Claim 4.8.4) and $\hat{\varphi}$ extends φ . By Lemma 4.8.5 we get an automorphism $\tilde{\varphi}$ of \mathbf{B} extending φ , that is, $\tilde{\varphi}(I) \subseteq A'$. Therefore \mathbf{B} is indeed irreducible structure faithful.

To prove coherence, let φ_1, φ_2 and φ be a coherent triple of partial automorphisms of \mathbf{A}' and let $\hat{\varphi}_1, \hat{\varphi}_2, \hat{\varphi}$ be automorphisms of \mathbf{B}_0 which are the coherent extensions of their projections by π .

Denote by $\hat{q}_1, \hat{q}_2, \hat{q}$ and F, F_1 and F_2 the corresponding functions and sets from proof of Lemma 4.7.5. Coherence on the first coordinate follows from coherence of $\hat{\varphi}_1, \hat{\varphi}_2, \hat{\varphi}$. To get coherence on the second coordinate, we need to prove that F is the symmetric difference of F_1 and F_2 . This follows by the same argument as in the proof of Lemma 4.4.4: Since $\varphi = \varphi_2 \circ \varphi_1$ and $\hat{\varphi} = \hat{\varphi}_2 \circ \hat{\varphi}_1$, we have that

$$\chi(x, \mathbf{V})(\vec{c}) \neq \chi(\varphi((x, \mathbf{V}))) (\hat{\varphi}(\vec{c}))$$

if and only if exactly one of

$$\chi(x, \mathbf{V})(\vec{c}) \neq \chi(\varphi_1((x, \mathbf{V}))) (\hat{\varphi}_1(\vec{c}))$$

and

$$\chi(\varphi_1((x, \mathbf{V}))) (\hat{\varphi}_1(\vec{c})) \neq \chi(\varphi_2(\varphi_1((x, \mathbf{V})))) (\hat{\varphi}_2(\hat{\varphi}_1(\vec{c})))$$

happens. □

To finish the proof of Lemma 4.8.1, we now prove that for every $C \subseteq B$ such that $E_{\mathbf{B}} \cap C^2$ is a cycle of length ≥ 4 , it holds that $\pi|_C$ is not an embedding (of the reducts to relation E). This would imply that whenever $C \subseteq B$ contains an induced graph cycle of length ≥ 4 , one of (2b) and (2c) holds.

Fix a set $C \subseteq B$ such that $E_{\mathbf{B}}|_C$ is an induced graph cycle of length ≥ 4 and, for a contradiction, assume that its projection $\pi(C)$ is again an induced graph cycle of the same length in the relation $E_{\mathbf{B}_0}$. This means that we can enumerate C as $(x_1, \mathbf{V}_1), \dots, (x_k, \mathbf{V}_k)$ such that $\vec{c} = (x_1, \dots, x_k)$ is bad cycle sequence. For every $1 \leq i \leq k$, we have $\{(x_i, \mathbf{V}_i), (x_{i+1}, \mathbf{V}_{i+1})\} \in E_{\mathbf{B}}$ (identifying

$(x_{k+1}, \mathbf{V}_{k+1}) = (x_1, \mathbf{V}_1)$), so in particular the set $\{(x_i, \mathbf{V}_i), (x_{i+1}, \mathbf{V}_{i+1})\}$ is generic. By definition, this implies that for every $1 \leq i < k$, we have

$$\chi(x_i, \mathbf{V}_i)(\vec{c}) = \chi(x_{i+1}, \mathbf{V}_{i+1})(\vec{c}),$$

but

$$\chi(x_1, \mathbf{V}_1)(\vec{c}) \neq \chi(x_k, \mathbf{V}_k)(\vec{c}),$$

which is a contradiction.

Remark 4.8.1. Exactly as in the previous section, our partial automorphism extension has some functorial properties. Taking isomorphic copies, we can assume that $\mathbf{A} \subseteq \mathbf{B}_0$ and $\mathbf{A} \subseteq \mathbf{B}$ (then in particular $\pi \upharpoonright_{\mathbf{A}} = \text{id}_{\mathbf{A}}$). Given a partial automorphism φ of \mathbf{A} let $\hat{\varphi}$ be its (coherent) extension to an automorphism of \mathbf{B}_0 and let $\tilde{\varphi}$ be the constructed extension to an automorphism of \mathbf{B} using $\hat{\varphi}$. Let \sim be an equivalence relation on B given by $x \sim y \iff \pi(x) = \pi(y)$. Then \sim is a congruence with respect to $\tilde{\varphi}$ and the natural actions of $\hat{\varphi}$ and $\tilde{\varphi}$ on the equivalence classes of \sim coincide.

Observation 4.8.8. *The number of vertices of the EPPA-witness \mathbf{B} constructed in this section can be bounded from above by a function which depends only on the number of vertices of \mathbf{B}_0 .*

Proof. Let m be the number of vertices of \mathbf{B}_0 . There are at most $(m+1)^m$ (rough estimate) bad cycle sequences and hence at most $2^{(m+1)^m}$ valuation functions for any given vertex $x \in B_0$. This means that there are at most $m2^{(m+1)^m}$ different pairs (x, χ) , where $x \in B_0$ and χ is a valuation function for x , and thus at most $2^{m2^{(m+1)^m}}$ different generic sets. Given a generic set V and a vertex $x \in B_0$, there is at most one valuation structure for x with vertex set V . Since the vertex set of \mathbf{B} consists of all pairs (x, \mathbf{V}) , where $x \in B_0$ and \mathbf{V} is a valuation structure for x , the claim then follows. \square

Proof of Claim 4.8.2. This is a word-to-word copy of the proof of Claim 4.7.2.

For a contradiction, suppose that it is not the case, that is, there are $(u, \chi), (u', \chi') \in \bigcup_{(x, \mathbf{V}) \in D} V$ which form a non-generic pair. This implies that there are $(x, \mathbf{X}), (y, \mathbf{Y}) \in D$ such that $(u, \chi) \in X$ and $(u', \chi') \in Y$ (they cannot both lie in the same valuation structure, because the vertex sets of valuation structures are generic), and hence the set $\{(x, \mathbf{X}), (y, \mathbf{Y})\}$ is non-generic. Put $\mathbf{E}_x = \{(a, \mathbf{U}) \in D : (x, \mathbf{X}) \notin \text{Cl}_{\mathbf{D}}((a, \mathbf{U}))\}$ and similarly define \mathbf{E}_y . Since closures are unary, these are substructures of \mathbf{D} . Note that as closures in \mathbf{B} are generic, we also know that $(y, \mathbf{Y}) \in \mathbf{E}_x$ and $(x, \mathbf{X}) \in \mathbf{E}_y$. This means that $\mathbf{E}_x, \mathbf{E}_y$ are both non-empty and neither is a substructure of the other.

We first prove $\mathbf{E}_x \cup \mathbf{E}_y = \mathbf{D}$. Suppose for a contradiction that there is $(z, \mathbf{Z}) \in \mathbf{D}$ with $\{(x, \mathbf{X}), (y, \mathbf{Y})\} \subseteq \text{Cl}_{\mathbf{D}}((z, \mathbf{Z}))$. Then (by the construction of \mathbf{B}) we have $\mathbf{X}, \mathbf{Y} \subseteq \mathbf{Z}$, which is a contradiction with $(x, \mathbf{X}), (y, \mathbf{Y})$ forming a non-generic pair.

Fix $(a, \mathbf{U}) \in \mathbf{E}_x \setminus \mathbf{E}_y$ and $(b, \mathbf{W}) \in \mathbf{E}_y \setminus \mathbf{E}_x$. Because we know that $\mathbf{Y} \subseteq \mathbf{U}$ and $\mathbf{X} \subseteq \mathbf{W}$, we get that $(a, \mathbf{U}), (b, \mathbf{W})$ is not a generic pair and therefore no relation of \mathbf{D} contains both (a, \mathbf{U}) and (b, \mathbf{W}) . Thus \mathbf{D} is a free amalgam of \mathbf{E}_x and \mathbf{E}_y over their intersection, which is a contradiction with its irreducibility. Therefore D is indeed generic. \square

Proof of Claim 4.8.3. Finiteness of \mathbf{B} follows from Observation 4.8.8. Given $(x, \mathbf{V}) \in B$, we have that $\pi(\mathbf{V}) = \text{Cl}_{\mathbf{B}_0}(x)$ (since \mathbf{V} is a valuation structure). By definition of \mathbf{B} ,

$$\pi(F_{\mathbf{B}}((x, \mathbf{V}))) = \{y : (y, \chi) \in F_{\mathbf{V}}((x, \chi(x, \mathbf{V})))\}.$$

As V is generic, we have that for every y there is at most one χ such that $(y, \chi) \in V$. Moreover, since π is an isomorphism $\mathbf{V} \rightarrow \text{Cl}_{\mathbf{B}_0}(x)$, we can write

$$\pi(F_{\mathbf{B}}((x, \mathbf{V}))) = \{y : y \in F_{\text{Cl}_{\mathbf{B}_0}}(x)\},$$

and because $F_{\text{Cl}_{\mathbf{B}_0}}(x) = F_{\mathbf{B}_0}(x)$, we indeed have

$$\pi(F_{\mathbf{B}}((x, \mathbf{V}))) = F_{\mathbf{B}_0}(x),$$

that is, π preserves functions.

Let $R \in L$ be a relation and let $((x_1, \mathbf{V}_1), \dots, (x_n, \mathbf{V}_n))$ be a tuple of vertices of \mathbf{B} . Clearly, if $((x_1, \mathbf{V}_1), \dots, (x_n, \mathbf{V}_n)) \in R_{\mathbf{B}}$ then $\pi(((x_1, \mathbf{V}_1), \dots, (x_n, \mathbf{V}_n))) = (x_1, \dots, x_n) \in R_{\mathbf{B}_0}$, hence π is a homomorphism. If $\{(x_1, \mathbf{V}_1), \dots, (x_n, \mathbf{V}_n)\}$ is generic, the definition of \mathbf{B} gives us that $((x_1, \mathbf{V}_1), \dots, (x_n, \mathbf{V}_n)) \in \mathbf{B}$ if and only if $\pi(((x_1, \mathbf{V}_1), \dots, (x_n, \mathbf{V}_n))) = (x_1, \dots, x_n) \in R_{\mathbf{B}_0}$, hence π is an embedding on every generic set.

The fact that π is a homomorphism-embedding now follows from Claim 4.8.2. \square

Proof of Claim 4.8.4. First we will prove that A' is a generic set. By definition this happens if $V = \bigcup_{(x, \mathbf{V}_x) \in A'} V_x = \bigcup_{x \in A} V_x$ is a generic set. Note that $V = \{(x, \chi_x) : x \in A\}$. Let $x \neq y \in A$ be arbitrary and pick $\vec{c} \in U(x) \cap U(y)$. Remember that since $E_{\mathbf{A}}$ is a complete graph, there are at most two vertices of \mathbf{A} in every bad cycle sequence, and if there are two, then they are connected by an edge of the cycle. Hence we have $\chi_x(\vec{c}) = \chi_{\vec{c}}(x)$ and $\chi_{\vec{c}}(y) = \chi_y(\vec{c})$. By the choice of $\chi_{\vec{c}}$ we get that (x, χ_x) and (y, χ_y) form a generic pair which implies that V is indeed a generic set. From this it follows that in particular every V_x is a generic set and hence ψ is a function $A \rightarrow B$. What follows is a word-to-word copy of the proof of Claim 4.7.4.

Next fix a relation $R \in L$ and a tuple $\bar{x} \in A^n$. We will prove that $\psi(\bar{x}) \in R_{\mathbf{B}}$ if and only if $\bar{x} \in R_{\mathbf{A}}$. By the definition of \mathbf{B} the “only if” part is immediate, to prove the other implication, we need to prove that $\psi(\bar{x})$ (understood as a set) is generic, but clearly $\psi(\bar{x})$ is a subset of a generic set V , which concludes the proof.

Finally we prove that for every (unary) function $F \in L$ and every vertex $x \in \mathbf{A}$ we have $\psi(F_{\mathbf{A}}(x)) = F_{\mathbf{B}}(\psi(x))$ (remember that $\psi_L = \text{id}_L$). Clearly

$$\psi(F_{\mathbf{A}}(x)) = \{(y, \mathbf{V}_y) : y \in F_{\mathbf{A}}(x)\}.$$

Put $X = F_{\mathbf{B}}(\psi(x))$. By the construction we have

$$X = F_{\mathbf{B}}((x, \mathbf{V}_x)) = \{(y, \text{Cl}_{\mathbf{V}_x}((y, \chi))) : (y, \chi) \in F_{\mathbf{V}_x}((x, \chi(x, \mathbf{V}_x)))\}.$$

Note that $\chi(x, \mathbf{V}_x) = \chi_x$ and that if $(y, \chi) \in F_{\mathbf{V}_x}((x, \chi(x, \mathbf{V}_x)))$, then $\chi = \chi_y$. Hence, in particular, π is injective on $F_{\mathbf{V}_x}((x, \chi_x))$. So we can write

$$X = \{(y, \text{Cl}_{\mathbf{V}_x}((y, \chi_y))) : (y, \chi_y) \in F_{\mathbf{V}_x}((x, \chi_x))\}.$$

Since $(\text{id}_L, (y, \chi_y) \mapsto y)$ is an isomorphism $\mathbf{V}_x \rightarrow \text{Cl}_{\mathbf{A}}(x)$, we get that $(y, \chi_y) \in F_{\mathbf{V}_x}((x, \chi_x))$ if and only if $y \in F_{\text{Cl}_{\mathbf{A}}(x)}(x) = F_{\mathbf{A}}(x)$. For the same reason, $\text{Cl}_{\mathbf{V}_x}((y, \chi_y))$ is isomorphic to $\text{Cl}_{\mathbf{A}}(y)$ by projecting to the first coordinate and hence in fact $\text{Cl}_{\mathbf{V}_x}((y, \chi_y)) = \mathbf{V}_y$. Putting this together, we get that

$$F_{\mathbf{B}}(\psi(x)) = X = \{(y, \mathbf{V}_y) : y \in F_{\mathbf{A}}(x)\} = \psi(F_{\mathbf{A}}(x)).$$

□

Proof of Claim 4.8.6. It is easy to see that $\tilde{\varphi}$ is a bijection which maps generic sets to generic sets. Since the domain of φ is generic, it follows that for every bad cycle sequence \vec{c} there are at most two vertices from \vec{c} in $\pi(\text{Dom}(\varphi))$, and if there are two of them, they are connected by an edge of the cycle. The same holds for $\text{Range}(\varphi)$.

Fix a bad cycle sequence \vec{c} of length k and suppose that there are distinct vertices $x, y \in \vec{c}$ and valuation structures \mathbf{U}, \mathbf{V} such that $(x, \mathbf{U}), (y, \mathbf{V}) \in \text{Dom}(\varphi)$ and denote $\varphi((x, \mathbf{U})) = (\hat{\varphi}(x), \mathbf{U}')$ and $\varphi((y, \mathbf{V})) = (\hat{\varphi}(y), \mathbf{V}')$. We know that there is i such that (without loss of generality) $x = \vec{c}_i$ and $y = \vec{c}_{i+1}$ (with $\vec{c}_{k+1} = \vec{c}_1$). By genericity of $\text{Dom}(\varphi)$ we know that $\chi(x, \mathbf{U})(\vec{c}) = \chi(y, \mathbf{V})(\vec{c})$ if and only if $i \neq k$. Because $\hat{\varphi}$ is a bijection $B_0 \rightarrow B_0$ we have that $\hat{\varphi}(x) = \hat{\varphi}(\vec{c})_i$ and $\hat{\varphi}(y) = \hat{\varphi}(\vec{c})_{i+1}$. And as $\text{Range}(\varphi)$ is also generic, we get that $\chi(\hat{\varphi}(x), \mathbf{U}')(\hat{\varphi}(\vec{c})) = \chi(\hat{\varphi}(y), \mathbf{V}')(\hat{\varphi}(\vec{c}))$ if and only if $i \neq k$. Hence

$$\chi(x, \mathbf{U})(\vec{c}) = \chi(y, \mathbf{V})(\vec{c}) \iff \chi(\hat{\varphi}(x), \mathbf{U}')(\hat{\varphi}(\vec{c})) = \chi(\hat{\varphi}(y), \mathbf{V}')(\hat{\varphi}(\vec{c})).$$

Moreover, this happens if and only if $\vec{c} \in F$, and thus $\tilde{\varphi}$ extends φ . In the remaining paragraphs we prove that $\tilde{\varphi}$ is an automorphism of \mathbf{B} . The proof is in fact a word-to-word copy of the analogous argument from Claim 4.7.6.

Fix a relation $R \in L$ and a tuple $((x_1, \mathbf{V}_1), \dots, (x_n, \mathbf{V}_n)) \in B^n$. Note that

$$(x_1, \dots, x_n) \in R_{\mathbf{B}_0} \iff (\hat{\varphi}(x_1), \dots, \hat{\varphi}(x_n)) \in \tilde{\varphi}(R)_{\mathbf{B}_0},$$

because $\hat{\varphi}$ is an automorphism of \mathbf{B}_0 . Together with the fact that $\tilde{\varphi}$ maps generic sets to generic sets it follows that $((x_1, \mathbf{V}_1), \dots, (x_n, \mathbf{V}_n)) \in R_{\mathbf{B}}$ if and only if $(\tilde{\varphi}((x_1, \mathbf{V}_1)), \dots, \tilde{\varphi}((x_n, \mathbf{V}_n))) \in \tilde{\varphi}(R)_{\mathbf{B}}$.

It remains to prove that for every function $F \in L$ and every $(x, \mathbf{V}) \in B$ we have $\tilde{\varphi}(F_{\mathbf{B}}((x, \mathbf{V}))) = \varphi_L(F)_{\mathbf{B}}(\tilde{\varphi}((x, \mathbf{V})))$. To simplify the notation we put $h(\mathbf{V}) = (\varphi_L, \hat{q})(\mathbf{V})$ for every valuation structure \mathbf{V} . Fix an arbitrary $(x, \mathbf{V}) \in \mathbf{B}$ and put $\chi_0 = \chi(x, \mathbf{V})$. By the definition of \mathbf{B} we know that

$$\tilde{\varphi}(F_{\mathbf{B}}((x, \mathbf{V}))) = \{(\hat{\varphi}(y), h(\text{Cl}_{\mathbf{V}}((y, \chi)))) : (y, \chi) \in F_{\mathbf{V}}((x, \chi_0))\}.$$

Denote $X = \tilde{\varphi}(F)_{\mathbf{B}}(\tilde{\varphi}((x, \mathbf{V})))$. By the definition of \mathbf{B} , we have

$$X = \{(y, \text{Cl}_{h(\mathbf{V})}((y, \chi))) : (y, \chi) \in \tilde{\varphi}(F)_{h(\mathbf{V})}(\hat{q}((x, \chi_0)))\}.$$

Since \hat{q} is a bijection $\mathcal{V} \rightarrow \mathcal{V}$ which agrees with $\hat{\varphi}$ on the first coordinate, we can write

$$X = \{(\hat{\varphi}(y), \text{Cl}_{h(\mathbf{V})}(\hat{q}((y, \chi)))) : \hat{q}((y, \chi)) \in \tilde{\varphi}(F)_{h(\mathbf{V})}(\hat{q}((x, \chi_0)))\}.$$

Note that

$$\tilde{\varphi}(F)_{h(\mathbf{V})}(\hat{q}((x, \chi_0))) = \hat{q}(F_{\mathbf{V}}((x, \chi_0))),$$

hence $\hat{q}((y, \chi)) \in \tilde{\varphi}(F)_{h(\mathbf{V})}(\hat{q}((x, \chi_0)))$ if and only if $(y, \chi) \in F_{\mathbf{V}}((x, \chi_0))$, and so we have

$$X = \{(\hat{\varphi}(y), \text{Cl}_{h(\mathbf{V})}(\hat{q}((y, \chi)))) : (y, \chi) \in F_{\mathbf{V}}((x, \chi_0))\}.$$

Finally, $\text{Cl}_{h(\mathbf{V})}(\hat{q}((y, \chi))) = h(\text{Cl}_{\mathbf{V}}((y, \chi)))$, hence indeed

$$X = \tilde{\varphi}(F_{\mathbf{B}}((x, \mathbf{V}))).$$

□

4.9 Locally tree-like EPPA-witnesses: Proof of Theorem 4.1.2

The goal of this section is to prove Theorem 4.1.2 using Lemma 4.8.1.

Definition 4.9.1. Let L be a language equipped with a permutation group Γ_L and let \mathbf{A} be a finite Γ_L -structure. We recursively define what a *tree amalgamation of copies of \mathbf{A}* is.

1. If \mathbf{D} is isomorphic to \mathbf{A} then \mathbf{D} is a tree amalgamation of copies of \mathbf{A} .
2. If \mathbf{B}_1 and \mathbf{B}_2 are tree amalgamations of copies of \mathbf{A} , \mathbf{D} is a Γ_L -structure and $\alpha_1: \mathbf{A} \rightarrow \mathbf{B}_1$, $\alpha_2: \mathbf{A} \rightarrow \mathbf{B}_2$ and $\delta_1, \delta_2: \mathbf{D} \rightarrow \mathbf{A}$ are embeddings then the free amalgamation of \mathbf{B}_1 and \mathbf{B}_2 over \mathbf{D} with respect to $\alpha_1 \circ \delta_1$ and $\alpha_2 \circ \delta_2$ is also a tree amalgamation of copies of \mathbf{A} .

The following proposition gives an alternative way of viewing tree amalgamations.

Proposition 4.9.1. *Let L be a language equipped with a permutation group Γ_L , let \mathbf{A} be a finite Γ_L -structure and let \mathbf{C} be a finite Γ_L -structure. The following statements are equivalent:*

1. \mathbf{C} is a tree amalgamation of copies of \mathbf{A} .
2. *There exists a sequence $\mathbf{A} = \mathbf{C}_1, \dots, \mathbf{C}_n = \mathbf{C}$ of finite Γ_L -structures such that for every $1 \leq i < n$ there is a Γ_L -structure \mathbf{D} , embeddings $\delta_1, \delta_2: \mathbf{D} \rightarrow \mathbf{A}$ and an embedding $\alpha: \mathbf{A} \rightarrow \mathbf{D}_i$ such that \mathbf{D}_{i+1} is a free amalgamation of \mathbf{A} and \mathbf{D}_i with respect to δ_1 and $\alpha \circ \delta_2$.*

Note that if \mathbf{A} is irreducible (which will always be the case in this paper), the process in point 2 can be understood as having a graph tree \mathbf{T} whose vertices precisely correspond to copies of \mathbf{A} in \mathbf{C} and each edge determines, how the neighbouring copies of \mathbf{A} overlap.

Proof of Proposition 4.9.1. The direction (2) \Rightarrow (1) is trivial, as (2) is just a special case of the recursive definition of tree amalgamation of copies of \mathbf{A} .

To obtain the other direction, we will use induction on the recursive construction of \mathbf{C} to prove an even stronger statement, namely that for every copy of $\mathbf{A} \in \mathbf{C}$, we can pick \mathbf{C}_1 to correspond to the given copy. Clearly, this holds if \mathbf{C} is isomorphic to \mathbf{A} . Suppose now that \mathbf{C} is the free amalgamation of \mathbf{B}_1 and

\mathbf{B}_2 with respect to $\alpha_1 \circ \delta_1$ and $\alpha_2 \circ \delta_2$ as in Definition 4.9.1 and without loss of generality assume that the chosen copy of \mathbf{A} lies in \mathbf{B}_1 .

By the induction hypothesis, we get $\mathbf{A} = \mathbf{C}_1, \dots, \mathbf{C}_n = \mathbf{B}_1$ such that \mathbf{C}_1 corresponds to the chosen copy. Also by the induction hypothesis, we get $\mathbf{A} = \mathbf{C}'_1, \dots, \mathbf{C}'_m = \mathbf{B}_2$ such that \mathbf{C}'_1 corresponds to the copy of \mathbf{A} given by α_2 . It is easy to see that if, for $1 \leq i \leq m$, we put \mathbf{C}_{n+1} to be the free amalgamation of \mathbf{C}_n and \mathbf{C}'_i with respect to $\alpha_1 \circ \delta_1$ and $\alpha_2 \circ \delta_2$, then $\mathbf{C}_1, \dots, \mathbf{C}_{m+n}$ witnesses that \mathbf{C} satisfies point 2. \square

Note that in both equivalent definitions, we require that we only amalgamate over copies of \mathbf{D} which lie in a copy of \mathbf{A} . The reason for it is that when \mathbf{A} is irreducible, it allows us to prove the following two observations about tree amalgamations of copies of \mathbf{A} .

Observation 4.9.2. *Let \mathbf{A} be a finite irreducible Γ_L -structure, let \mathbf{C} be a tree amalgamation of copies of \mathbf{A} witnessed by the sequence $\mathbf{A} = \mathbf{C}_1, \dots, \mathbf{C}_n = \mathbf{C}$ of Γ_L -structures and let $\mathbf{I} \subseteq \mathbf{C}$ be an irreducible structure. Then either $\mathbf{I} \subseteq \mathbf{C}_1$, or there is $1 < i \leq n$ such that $\mathbf{I} \not\subseteq \mathbf{C}_{i-1}$ and \mathbf{I} lies fully in the copy of \mathbf{A} which together with \mathbf{C}_{i-1} forms \mathbf{C}_i . Consequently, every embedding of an irreducible structure to \mathbf{C} extends to an embedding of \mathbf{A} to \mathbf{C}*

Proof. This follows from the fact that if \mathbf{U} is a free amalgamation of \mathbf{U}_1 and \mathbf{U}_2 (without loss of generality we can assume that all the embeddings are inclusions and $U = U_1 \cup U_2$) and \mathbf{V} is an irreducible substructure of \mathbf{U} , then $\mathbf{V} \subseteq \mathbf{U}_1$ or $\mathbf{V} \subseteq \mathbf{U}_2$. The “consequently” part is immediate. \square

Note that this in particular implies that the only copies of \mathbf{A} in \mathbf{C} are those which we added in some step of the construction of \mathbf{C} .

Observation 4.9.3. *Let \mathcal{C} be a hereditary amalgamation class of finite Γ_L -structures and let $\mathbf{A} \in \mathcal{C}$ be an irreducible structure. Suppose that \mathbf{C} is a tree amalgamation of copies of \mathbf{A} . Then there is $\mathbf{E} \in \mathcal{C}$ and a homomorphism-embedding $e: \mathbf{C} \rightarrow \mathbf{E}$.*

Proof. We will proceed by induction on the recursive definition of \mathbf{C} . If \mathbf{C} is isomorphic to \mathbf{A} , then the statement clearly holds with e being the identity. Otherwise we get $\mathbf{B}_1, \mathbf{B}_2, \mathbf{D}, \delta_1, \delta_2, \alpha_1$ and α_2 as in Definition 4.9.1. By the induction hypothesis, we get $\mathbf{E}_1, \mathbf{E}_2 \in \mathcal{C}$ and homomorphism-embeddings $e_1: \mathbf{B}_1 \rightarrow \mathbf{E}_1$ and $e_2: \mathbf{B}_2 \rightarrow \mathbf{E}_2$. Since \mathbf{A} is irreducible, we get that $e_i \circ \alpha_i$ is an embedding $\mathbf{A} \rightarrow \mathbf{E}_i$ for $i \in \{1, 2\}$, hence in particular the structure induced by \mathbf{E}_i on $e_i(\alpha_i(\delta_i(D)))$ is isomorphic to \mathbf{D} for $i \in \{1, 2\}$. Therefore, we can put \mathbf{E} to be the amalgamation of \mathbf{E}_1 and \mathbf{E}_2 with respect to $e_1 \circ \alpha_1 \circ \delta_1$ and $e_2 \circ \alpha_2 \circ \delta_2$. \square

Note that the fact that we are only amalgamating over structures which lie in a copy of \mathbf{A} was crucial, because “being irreducible” is not a hereditary property (for example, if L is a language containing one ternary relation R and \mathbf{X} is an L -structure such that $X = \{a, b, c\}$ and $R_{\mathbf{X}} = \{(a, b, c)\}$, then \mathbf{X} is irreducible, but the substructure of \mathbf{X} induced on $\{a, b\}$ is not irreducible).

We will make use of the following lemma which has a graph-theoretic proof:

Lemma 4.9.4. *Let L be a language equipped with a permutation group Γ_L , assume that L contains a binary symmetric relation E , and let \mathbf{A} be a finite irreducible Γ_L -structure such that $E_{\mathbf{A}}$ is a complete graph. Let \mathbf{B} be a Γ_L -structure satisfying the following:*

1. *Every irreducible substructure of \mathbf{B} is isomorphic to a substructure of \mathbf{A} , and*
2. *\mathbf{B} contains no induced graph cycles (of length ≥ 4) in the relation $E_{\mathbf{B}}$.*

Then \mathbf{B} is a substructure of a tree amalgamation of copies of \mathbf{A} .

Proof. We proceed by induction on $|B|$. If \mathbf{B} is irreducible then the statement follows trivially, hence we can assume that \mathbf{B} is reducible.

Note that condition 1 implies that if \mathbf{C} is an irreducible substructure of \mathbf{B} , then $E_{\mathbf{C}}$ is a clique. And conversely, whenever $E_{\mathbf{B}}$ induces a clique on $C \subseteq B$ then $\text{Cl}_{\mathbf{B}}(C)$ is irreducible: Indeed, suppose for a contradiction that $\text{Cl}_{\mathbf{B}}(C)$ is the free amalgamation of some \mathbf{U}_1 and \mathbf{U}_2 over \mathbf{V} such that $\mathbf{U}_1, \mathbf{U}_2 \neq \text{Cl}_{\mathbf{B}}(C)$. If $C \subseteq U_1$, then we would get that $\text{Cl}_{\mathbf{B}}(C) \subseteq \mathbf{U}_1$ (because \mathbf{U}_1 is a substructure of $\text{Cl}_{\mathbf{B}}(C)$) which is a contradiction, similarly for U_2 . Hence there are $x_1, x_2 \in C$ such that $x_1 \in U_1 \setminus V$ and $x_2 \in U_2 \setminus V$. But this implies that $(x_1, x_2) \notin E_{\mathbf{B}}$, which is a contradiction.

For the following paragraphs, we will mainly consider the graph relation $E_{\mathbf{B}}$ and we will treat subsets of B as (induced) subgraphs of the graph $(B, E_{\mathbf{B}})$. We will use the standard terminology of graph theory.

Let \mathbf{C} be an inclusion minimal substructure of \mathbf{B} such that C forms a vertex cut of B (i.e. $B \setminus C$ is not connected in $E_{\mathbf{B}}$) and let $C' \subseteq C$ be an inclusion minimal vertex cut of B . Such a \mathbf{C} exists, because \mathbf{B} is reducible. Note that from the minimality of \mathbf{C} it follows that $\mathbf{C} = \text{Cl}_{\mathbf{B}}(C')$.

First observe that from the minimality of C' it follows that for every pair of distinct vertices $x, y \in C'$ there are two distinct nonempty connected components $B_1, B_2 \subset B \setminus C'$ such that both B_1, B_2 contain a vertex adjacent to x as well as a vertex adjacent to y . Now observe that C' is a clique: If there was a pair of vertices $x \neq y \in C'$ such that $(x, y) \notin E_{\mathbf{B}}$, we could construct an induced cycle of length ≥ 4 using x and y and vertices of B_1, B_2 from the previous paragraph. This implies that \mathbf{C} is irreducible, because it is the closure of a clique.

From the condition on \mathbf{C} we get that $B \setminus C$ is not connected, that is, it can be split into two non-empty disjoint parts $B_1 \cup B_2 = B \setminus C$ such that there are no edges between B_1 and B_2 (and therefore no relations or functions at all thanks to condition 1). This means that \mathbf{B} is the free amalgamation of (its substructures induced on) $B_1 \cup C$ and $B_2 \cup C$ over \mathbf{C} .

Using the induction hypothesis, we get Γ_L -structures \mathbf{D}_1 and \mathbf{D}_2 which are tree amalgamations of copies of \mathbf{A} such that the substructures induced by \mathbf{B} on $B_i \cup C$ are substructures of \mathbf{D}_i for $i \in \{1, 2\}$. Since \mathbf{C} is irreducible, it follows that there are embeddings $\alpha_1: \mathbf{A} \rightarrow \mathbf{D}_1$, $\alpha_2: \mathbf{A} \rightarrow \mathbf{D}_2$ and $\delta_1, \delta_2: \mathbf{C} \rightarrow \mathbf{A}$ (by Observation 4.9.2) and hence we can put \mathbf{D} to be the free amalgamation of \mathbf{B}_1 and \mathbf{B}_2 over \mathbf{C} with respect to $\alpha_1 \circ \delta_1$ and $\alpha_2 \circ \delta_2$. Clearly $\mathbf{B} \subseteq \mathbf{D}$ (up to an isomorphism) and \mathbf{D} is a tree amalgamation of copies of \mathbf{A} . \square

Now we are ready to prove Theorem 4.1.2.

Proof of Theorem 4.1.2. We intend to use Lemma 4.8.1 as the main ingredient to this proof. However, Lemma 4.8.1 expects that there is a graph edge relation E in the language and that $E_{\mathbf{A}}$ is a complete graph, which is not guaranteed by the assumptions of Theorem 4.1.2. For this reason, we extend the language L to L^+ , adding a binary symmetric relation E fixed by every permutation of the language (assuming without loss of generality that $E \notin L$), put \mathbf{A}' to be the Γ_{L^+} -structure obtained from \mathbf{A} by putting $E_{\mathbf{A}'}$ to be the complete graph on A , and put \mathbf{B}'_0 to be the Γ_{L^+} -structure obtained from \mathbf{B}_0 by putting $E_{\mathbf{B}'_0}$ to be the complete graph on \mathbf{B}'_0 .

Observe that partial automorphisms of \mathbf{A}' are precisely the partial automorphisms of \mathbf{A} (barring the new relation E fixed by every permutation of the language) and $\text{Aut}(\mathbf{B}_0) = \text{Aut}(\mathbf{B}'_0)$. Therefore, \mathbf{B}'_0 is an EPPA-witness for \mathbf{A}' and it is coherent if \mathbf{B}_0 is.

Put $N = (n-1)\binom{n}{2} + 1$. Use Proposition 4.7.1 on \mathbf{A}' and \mathbf{B}'_0 to get \mathbf{B}'_1 , an irreducible structure faithful (coherent) EPPA-witness for \mathbf{A}' with a homomorphism-embedding $f_1: \mathbf{B}'_1 \rightarrow \mathbf{B}'_0$. Next, by applying Lemma 4.8.1 iteratively N times, we construct a sequence of Γ_{L^+} -structures $\mathbf{B}'_2, \dots, \mathbf{B}'_{N+1}$ and a sequence of maps f_2, \dots, f_{N+1} , such that for every $1 \leq i \leq N+1$ it holds that \mathbf{B}'_i is an irreducible structure faithful EPPA-witness for \mathbf{A}' , f_i is a homomorphism-embedding $\mathbf{B}'_i \rightarrow \mathbf{B}'_{i-1}$, and if \mathbf{B}'_{i-1} is coherent then so is \mathbf{B}'_i .

Put $\mathbf{B}' = \mathbf{B}'_{N+1}$ and put \mathbf{B} to be the Γ_L -reduct of \mathbf{B}' forgetting the relation E . Since every automorphism of \mathbf{B}' is also an automorphism of \mathbf{B} , we get that \mathbf{B} is an EPPA-witness for \mathbf{A} , and if \mathbf{B}_0 was coherent, then so is \mathbf{B} . To see that \mathbf{B} is irreducible structure faithful, note that an irreducible substructure of \mathbf{B} is also an irreducible substructure of \mathbf{B}' .

Let \mathbf{C}_{N+1} be a substructure of \mathbf{B} on at most n vertices and let \mathbf{C}'_{N+1} be a substructure of \mathbf{B}' on the same vertices as \mathbf{C}_{N+1} . Denote by $\mathbf{C}'_1, \dots, \mathbf{C}'_N$ the structures such that for every $1 \leq i \leq N$ it holds that $\mathbf{C}'_i = f_{i+1}(\mathbf{C}'_{i+1})$.

Since we used Lemma 4.8.1 N times, let us count how many times one of (2b) and (2c) from Lemma 4.8.1 has happened. Clearly, possibility (2b) could have happened at most $n-1$ times, because $|\mathbf{C}'_{N+1}| \leq n$ and $|\mathbf{C}'_1| \geq 1$. And for every fixed $m = |\mathbf{C}'_i|$, possibility (2c) could have happened at most $\binom{m}{2} \leq \binom{n}{2}$ times. Therefore, (2b) or (2c) have together happened at most $N-1$ times, which means that there is $2 \leq i \leq N+1$ such that possibility (2a) happened in the i -th step. This then means that \mathbf{C}'_i contains no induced cycles of length ≥ 4 .

Because \mathbf{B}'_i is an irreducible structure faithful EPPA-witness for \mathbf{A}' , we get that every irreducible substructure of \mathbf{B}'_i is isomorphic to a substructure of \mathbf{A}' , so in particular, this holds for irreducible substructures of \mathbf{C}'_i . Hence, we can apply Lemma 4.9.4 on \mathbf{C}'_i to obtain a tree amalgamation \mathbf{D}' of copies of \mathbf{A}' and a homomorphism-embedding $f: \mathbf{C}'_{N+1} \rightarrow \mathbf{D}'$ (obtained by composing the output of Lemma 4.9.4 with some of the f_i 's).

Let \mathbf{D} be the Γ_L -reduct obtained from \mathbf{D}' by forgetting the relation E . It is easy to check that f is a homomorphism-embedding $\mathbf{C}_{N+1} \rightarrow \mathbf{D}$ and that \mathbf{D} is a tree amalgamation of copies of \mathbf{A} , which concludes the proof. \square

Observation 4.9.5. *The number of vertices of the EPPA-witness \mathbf{B} provided by Theorem 4.1.2 can be bounded from above by a function which depends only on the number of vertices of \mathbf{B}_0 and on n .*

Proof. \mathbf{B} was obtained from \mathbf{B}_0 by iteratively applying Lemma 4.8.1 $N = (n - 1) \binom{n}{2} + 1$ times. We know that in each step, the number of vertices of the constructed structure can be bounded by a function of the number of vertices of the original structure (by Observation 4.8.8). The claim then follows. \square

4.10 A generalisation of the Herwig–Lascar theorem: Proof of Theorem 4.1.5

Next, we show how Theorem 4.1.2 implies Theorem 4.1.5. Note that unlike Theorem 4.1.5, Theorem 4.1.2 assumes that \mathbf{A} is irreducible, because otherwise one can not define what tree amalgamation is. In order to deal with it, we extend the language L to L^+ , adding a binary symmetric relation E fixed by every permutation of the language (assuming without loss of generality that $E \notin L$), and consider the class consisting of finite Γ_{L^+} -structures \mathbf{A} where $E_{\mathbf{A}}$ is a complete graph. Moreover, for such \mathbf{A} , we will denote by \mathbf{A}^- its Γ_L -reduct forgetting the relation E .

We will need the following technical lemma.

Lemma 4.10.1. *Let L be a language consisting of relations and unary functions equipped with a permutation group Γ_L and fix a finite Γ_{L^+} -structure \mathbf{A} such that $E_{\mathbf{A}}$ is a complete graph. Assume that there is a (not necessary finite) Γ_L -structure \mathbf{M} containing \mathbf{A}^- as a substructure such that every partial automorphism of \mathbf{A}^- extends to an automorphism of \mathbf{M} . Then, for every tree amalgamation \mathbf{D} of copies of \mathbf{A} , there is a homomorphism-embedding $h: \mathbf{D}^- \rightarrow \mathbf{M}$. Moreover, for every embedding $\alpha: \mathbf{A} \rightarrow \mathbf{D}$ there is an automorphism f of \mathbf{M} such that $f(h(\alpha(A))) = A$.*

Proof. We proceed by induction on the tree construction of \mathbf{D} (Definition 4.9.1). The claim clearly holds if \mathbf{D} is isomorphic to \mathbf{A} . Suppose now that \mathbf{B}_1 and \mathbf{B}_2 are tree amalgamations of copies of \mathbf{A} and \mathbf{E} is a substructure of \mathbf{A} with embeddings $\delta_1, \delta_2: \mathbf{E} \rightarrow \mathbf{A}$, $\alpha_1: \mathbf{A} \rightarrow \mathbf{B}_1$ and $\alpha_2: \mathbf{A} \rightarrow \mathbf{B}_2$ such that \mathbf{D} is the free amalgamation of \mathbf{B}_1 and \mathbf{B}_2 over \mathbf{E} with respect to $\alpha_1 \circ \delta_1$ and $\alpha_2 \circ \delta_2$.

By the induction hypothesis, we have homomorphism-embeddings $h_1: \mathbf{B}_1^- \rightarrow \mathbf{M}$ and $h_2: \mathbf{B}_2^- \rightarrow \mathbf{M}$ and automorphisms f_1, f_2 of \mathbf{M} such that $f_i(h_i(\alpha_i(A))) = A$ for $i \in \{1, 2\}$. Let φ be a partial automorphism of \mathbf{A}^- sending

$$f_1(h_1(\alpha_1(\delta_1(\mathbf{E}^-)))) \mapsto f_2(h_2(\alpha_2(\delta_2(\mathbf{E}^-))))$$

and let $\hat{\varphi}$ be its extension to a partial automorphism of \mathbf{M} . It is easy to check that the function $h: D \rightarrow M$ defined by

$$h(x) = \begin{cases} \hat{\varphi}(f_1(h_1(x))) & \text{if } x \in B_1 \\ f_2(h_2(x)) & \text{otherwise} \end{cases}$$

is a homomorphism embedding $\mathbf{D}^- \rightarrow \mathbf{M}$. The moreover part follows straightforwardly as \mathbf{A} is irreducible and therefore every copy of \mathbf{A} in \mathbf{D} is either in \mathbf{B}_1 or in \mathbf{B}_2 . \square

Now we are ready to prove Theorem 4.1.5.

Proof of Theorem 4.1.5. Let \mathbf{A}^+ be the Γ_{L^+} -expansion of \mathbf{A} adding a clique in the relation E . Clearly, \mathbf{A}^+ is in a finite orbit of the action of Γ_{L^+} by relabelling, hence we can use Theorem 4.1.1 to get a Γ_{L^+} -structure \mathbf{B}_0 which is an irreducible structure faithful coherent EPPA-witness for \mathbf{A}^+ . Let n be the number of vertices of the largest structure in \mathcal{F} and let \mathbf{B} be given by Theorem 4.1.2. We will show that \mathbf{B}^- satisfies the statement. Clearly, it is a coherent EPPA-witness for \mathbf{A} . Since every irreducible substructure of \mathbf{B}^- is all the more so an irreducible substructure of \mathbf{B} , we get that \mathbf{B}^- is irreducible structure faithful. To finish the proof, it remains to show that $\mathbf{B}^- \in \text{Forb}_{\text{he}}(\mathcal{F})$.

For a contradiction, suppose that there is $\mathbf{F} \in \mathcal{F}$ with a homomorphism-embedding $g: \mathbf{F} \rightarrow \mathbf{B}^-$. We have that $|g(F)| \leq |F| \leq n$. Let \mathbf{C} be the substructure of \mathbf{B} induced on $g(F)$. From Theorem 4.1.2, we get a tree amalgamation \mathbf{D} of copies of \mathbf{A}^+ and a homomorphism-embedding $f: \mathbf{C} \rightarrow \mathbf{D}$. Composing $f \circ g$, we get that \mathbf{F} has a homomorphism-embedding to \mathbf{D}^- . However, Lemma 4.10.1 gives a homomorphism-embedding $\mathbf{D}^- \rightarrow \mathbf{M}$, hence we get a homomorphism-embedding $\mathbf{F} \rightarrow \mathbf{M}$, which is a contradiction with $\mathbf{M} \in \text{Forb}_{\text{he}}(\mathcal{F})$. \square

4.11 Connections to the structural Ramsey theory: Proof of Theorem 4.1.6

Most of the applications of the Herwig–Lascar theorem proceed similarly to applications of a theorem developed independently in the context of the structural Ramsey theory [HN19]. Both EPPA and the Ramsey property imply the amalgamation property (cf. Observation 4.2.3 and [Neš05]), however, the amalgamation property is not enough to imply either of them. This motivates the following strengthening of (strong) amalgamation introduced in [HN19]:

Definition 4.11.1. Let \mathbf{C} be a structure. An irreducible structure \mathbf{C}' is a *completion* of \mathbf{C} if there is a homomorphism-embedding $\mathbf{C} \rightarrow \mathbf{C}'$. It is a *strong completion* if the homomorphism-embedding is injective. A completion is *automorphism-preserving* if it is strong and for every $\alpha \in \text{Aut}(\mathbf{C})$ there is $\alpha' \in \text{Aut}(\mathbf{C}')$ such that $\alpha \subseteq \alpha'$ and moreover the map $\alpha \mapsto \alpha'$ is a group homomorphism $\text{Aut}(\mathbf{C}) \rightarrow \text{Aut}(\mathbf{C}')$.

To see that completion is a strengthening of amalgamation, let \mathcal{K} be a class of irreducible structures. The amalgamation property for \mathcal{K} can be equivalently formulated as follows: For $\mathbf{A}, \mathbf{B}_1, \mathbf{B}_2 \in \mathcal{K}$ embeddings $\alpha_1: \mathbf{A} \rightarrow \mathbf{B}_1$ and $\alpha_2: \mathbf{A} \rightarrow \mathbf{B}_2$, there is $\mathbf{C} \in \mathcal{K}$ which is a completion of the free amalgamation of \mathbf{B}_1 and \mathbf{B}_2 over \mathbf{A} with respect to α_1 and α_2 (which itself need not be in \mathcal{K}). In the same way, strong completion strengthens strong amalgamation and automorphism-preserving completion strengthens the so-called amalgamation property with automorphisms.

Definition 4.11.2. Let L be a language equipped with a permutation group Γ_L . Let \mathcal{E} be a class of finite Γ_L -structures and let \mathcal{K} be a subclass of \mathcal{E} consisting of irreducible structures. We say that \mathcal{K} is a *locally finite subclass* of \mathcal{E} if for every $\mathbf{A} \in \mathcal{K}$ and every $\mathbf{B}_0 \in \mathcal{E}$ there is a finite integer $n = n(\mathbf{A}, \mathbf{B}_0)$ such that every Γ_L -structure \mathbf{B} has a completion $\mathbf{B}' \in \mathcal{K}$, provided that it satisfies the following:




Figure 4.3: An example of a decomposition of a structure \mathbf{A} containing one unary function constructed in the proof of Proposition 4.11.1.

1. Every irreducible substructure of \mathbf{B} has an embedding to \mathbf{A} ,
2. there is a homomorphism-embedding from \mathbf{B} to \mathbf{B}_0 , and
3. every substructure of \mathbf{B} on at most n vertices has a completion in \mathcal{K} .

We say that \mathcal{K} is a *locally finite automorphism-preserving subclass of \mathcal{E}* if \mathbf{B}' can always be chosen to be automorphism-preserving.

Note that if \mathcal{K} is hereditary, point 1 implies that every irreducible substructure of \mathbf{B} is in \mathcal{K} . Note also that we are only promised that every substructure on at most n vertices has a completion in \mathcal{K} , even though we are asking for an automorphism-preserving (hence, in particular, strong) completion.

Luckily, for languages where all functions are unary, one can prove that if a structure has a completion in a strong amalgamation class then it has in fact a strong completion, which makes verifying local finiteness much easier. This was first proved in [HN19] as Proposition 2.6, we include a proof for completeness.

Proposition 4.11.1. *Let L be a language equipped with a permutation group Γ_L such that all function symbols of L are unary and let \mathcal{K} be a hereditary class of finite irreducible Γ_L -structures with the strong amalgamation property. For every finite Γ_L -structure \mathbf{A} , it holds that it has a completion in \mathcal{K} if and only if it has a strong completion in \mathcal{K} .*

Proof. One implication is trivial. To prove the other, assume to the contrary that there is a Γ_L -structure \mathbf{A} with no strong completion in \mathcal{K} , a Γ_L -structure $\mathbf{B} \in \mathcal{K}$ and a homomorphism-embedding $f: \mathbf{A} \rightarrow \mathbf{B}$ (that is, \mathbf{B} is a completion of \mathbf{A}). Among all such examples, choose one with $|\{\{u, v\} \subseteq A : f(u) = f(v)\}|$ minimal. Note that this implies that whenever there is a Γ_L -structure \mathbf{A}' and homomorphism-embeddings $g_1: \mathbf{A} \rightarrow \mathbf{A}'$ and $g_2: \mathbf{A}' \rightarrow \mathbf{B}$ such that $f = g_2 \circ g_1$ and g_1 is surjective, we have that either g_1 is injective, or g_2 is injective (as otherwise $g_2: \mathbf{A}' \rightarrow \mathbf{B}$ contradicts the minimality).

We decompose the vertex set of \mathbf{A} into five parts denoted by L_1 , L_2 , R_1 , R_2 , and C as depicted in Figure 4.3 by the following procedure.

Because f is not a strong completion in \mathcal{K} , we know that there is a pair of vertices $l \neq r \in A$ such that $f(l) = f(r)$. Now observe that, by the non-existence of \mathbf{A}' , for every other pair of vertices $v_1 \neq v_2 \in A$ satisfying $f(v_1) = f(v_2)$ it holds

that one vertex is in $\text{Cl}_{\mathbf{A}}(l)$ and the other is in $\text{Cl}_{\mathbf{A}}(r)$: Indeed, otherwise we could first identify only vertices from $\text{Cl}_{\mathbf{A}}(l)$ with vertices from $\text{Cl}_{\mathbf{A}}(r)$, yielding such a structure \mathbf{A}' .

Because vertex closures are irreducible substructures, we know that f identifies two irreducible substructures $\mathbf{U} = \text{Cl}_{\mathbf{A}}(l)$ and $\mathbf{V} = \text{Cl}_{\mathbf{A}}(r)$ of \mathbf{A} to one and is injective otherwise.

Put $L_1 = U \setminus V$ and $R_1 = V \setminus U$. Observe that because l and r can be chosen arbitrarily, if a substructure of \mathbf{A} contains a vertex of L_1 then it contains all vertices of L_1 (otherwise we would again get a contradiction with the non-existence of \mathbf{A}'). By symmetry, the same holds for R_1 . Denote by L_2 the set of all vertices $v \in A \setminus L_1$ such that $L_1 \subseteq \text{Cl}_{\mathbf{A}}(v)$. Analogously denote by R_2 the set of all vertices $v \in A \setminus R_1$ such that $R_1 \subseteq \text{Cl}_{\mathbf{A}}(v)$. L_2 and R_2 are disjoint, because f is an embedding on irreducible substructures, and thus no vertex closure (which is an irreducible substructure) can contain both L_1 and R_1 (as $f(L_1) = f(R_1)$). By a similar irreducibility argument, we get that there is no tuple $\bar{t} \in R_{\mathbf{A}}$, $R \in L$, containing both a vertex from $L_1 \cup L_2$ and a vertex from $R_1 \cup R_2$.

Let C be the set of all vertices whose vertex closure does not contain L_1 nor R_1 , that is, $C = A \setminus (L_1 \cup L_2 \cup R_1 \cup R_2)$. Because all functions are unary, \mathbf{A} induces a substructure \mathbf{C} on C . Similarly, denote by \mathbf{A}_l the substructure induced by \mathbf{A} on $C \cup L_1 \cup L_2$, and by \mathbf{A}_r the substructure induced by \mathbf{A} on $C \cup R_1 \cup R_2$.

Because \mathcal{K} is hereditary and f is injective on $A \setminus (L_1 \cup R_1)$, we know that $\mathbf{B} \in \mathcal{K}$ is a strong completion of all of \mathbf{A}_l , \mathbf{A}_r and \mathbf{C} . Applying the strong amalgamation property of \mathcal{K} , there is $\mathbf{D} \in \mathcal{K}$ which is a strong amalgamation of $f(\mathbf{A}_l)$ and $f(\mathbf{A}_r)$ over $f(\mathbf{C})$, hence a strong completion of \mathbf{A} , which is a contradiction. \square

Note that in [HN19] it is also observed that the unarity assumption of Proposition 4.11.1 cannot be omitted. We now prove Theorem 4.1.6.

Proof of Theorem 4.1.6. Given $\mathbf{A} \in \mathcal{K}$, use the fact that \mathcal{E} has (coherent) EPPA to obtain a (coherent) EPPA-witness $\mathbf{B}_0 \in \mathcal{E}$. Let $n = n(\mathbf{A}, \mathbf{B}_0)$ be as in the definition of a locally finite subclass and let \mathbf{B}_1 and a homomorphism-embedding $f: \mathbf{B}_1 \rightarrow \mathbf{B}_0$ be given by Theorem 4.1.2 for \mathbf{A} , \mathbf{B}_0 and n .

Because \mathbf{B}_1 is irreducible structure faithful, it follows that every irreducible structure of \mathbf{B}_1 can be sent by an automorphism to \mathbf{A} . We also get that every substructure $\mathbf{D} \subseteq \mathbf{B}_1$ on at most n vertices has a homomorphism-embedding to a tree amalgamation of copies of \mathbf{A} . Using Observation 4.9.3, we obtain $\mathbf{E} \in \mathcal{K}$ and a homomorphism-embedding $\mathbf{D} \rightarrow \mathbf{E}$, by composing these two homomorphism-embeddings, we get that every substructure of \mathbf{B}_1 on at most n vertices has a (not necessarily strong) completion in \mathcal{K} , and Proposition 4.11.1 gives us that it has a strong completion in \mathcal{K} .

Now we can use the fact that \mathcal{K} is a locally finite automorphism-preserving subclass of \mathcal{E} to get an automorphism-preserving completion \mathbf{B} of \mathbf{B}_1 . Finally, if \mathbf{B}_0 was coherent, then \mathbf{B}_1 and consequently \mathbf{B} are coherent, too, thanks to the moreover part of Definition 4.11.1. \square

4.12 Applications

In this section we present three applications of our general results.

4.12.1 Free amalgamation classes

We characterize free amalgamation classes of finite Γ_L -structures with relations and unary functions which have EPPA. We start with an easy observation.

Observation 4.12.1 ([HO03, EHN21, Sin17]). *Let \mathcal{K} be a free amalgamation class, let $\mathbf{A} \in \mathcal{K}$ be a finite structure and let \mathbf{B} be an irreducible structure faithful EPPA-witness for \mathbf{A} . Then $\mathbf{B} \in \mathcal{K}$.*

Proof. Assume for a contradiction that $\mathbf{B} \notin \mathcal{K}$. Let \mathbf{B}_0 be an inclusion minimal substructure of \mathbf{B} such that $\mathbf{B}_0 \notin \mathcal{K}$. Because \mathcal{K} is a free amalgamation class it follows that \mathbf{B}_0 is irreducible. However, this is a contradiction with the existence of an automorphism φ of \mathbf{B} such that $\varphi(\mathbf{B}_0) \subseteq \mathbf{A}$. \square

Now we can prove Corollary 4.1.4 which characterises free amalgamation classes with EPPA.

Proof. If there is $\mathbf{A} \in \mathcal{K}$ which lies in an infinite orbit of the action of Γ_L by relabelling then by Theorem 4.1.3 there is no finite EPPA-witness for \mathbf{A} , hence \mathcal{K} does not have EPPA.

If $\mathbf{A} \in \mathcal{K}$ lies in a finite orbit of the action of Γ_L by relabelling then by Theorem 4.1.1 there is a finite irreducible structure faithful coherent EPPA-witness \mathbf{B} for \mathbf{A} . By Observation 4.12.1 \mathbf{B} lies in \mathcal{K} . \square

4.12.2 Metric spaces without large cliques

We continue with an example of an application of Theorem 4.1.6, which was first proved by Conant [Con19, Theorem 3.9] (see also [ABWH⁺17b]).

Proposition 4.12.2. *Let \mathbf{K}_n denote the metric space on n vertices where all distances are 1. The class \mathcal{M}_n of all finite integer-valued metric spaces which do not contain a copy of \mathbf{K}_n has coherent EPPA for every $n \geq 2$.*

Proof. We will consider integer-valued metric spaces to be relational structures in the language $L = \{R^1, R^2, \dots\}$ (with trivial Γ_L), where $(x, y) \in R^a$ if and only if $d(x, y) = a$. We do not explicitly represent $d(x, x) = 0$. Let \mathcal{E}_n be the class of all L -structures \mathbf{A} such that $R_{\mathbf{A}}^i$ is symmetric and irreflexive for every $R^i \in L$, for every pair of vertices $x, y \in A$ it holds that $\{x, y\}$ is in at most one of $R_{\mathbf{A}}^i$ and $\mathbf{K}_n \not\subseteq \mathbf{A}$.

Clearly, \mathcal{E}_n is a free amalgamation class, and since Γ_L is trivial, we get that every orbit of the action of Γ_L by relabelling has size 1. Therefore, by Corollary 4.1.4, \mathcal{E}_n has irreducible structure faithful coherent EPPA. \mathcal{M}_n is a hereditary subclass of \mathcal{E}_n and consists of irreducible substructures. We need to verify that \mathcal{M}_n is a locally finite automorphism-preserving subclass of \mathcal{E}_n and that it has the strong amalgamation property in order to use Theorem 4.1.6 and thus finish the proof.

Note that if we have $\mathbf{B}_0 \in \mathcal{E}_n$ and a finite Γ_L -structure \mathbf{B} with a homomorphism-embedding $f: \mathbf{B} \rightarrow \mathbf{B}_0$, the following holds for \mathbf{B} :

1. $\mathbf{K}_n \not\subseteq \mathbf{B}$,
2. the relation $R_{\mathbf{B}}^i$ is symmetric and irreflexive for every $i \geq 1$,

3. every pair of vertices $x, y \in B$ is in at most one $R_{\mathbf{B}}^i$ relation, and
4. there is a finite set $S \subset \{1, 2, \dots\}$ such that for every $i \in \{1, 2, \dots\} \setminus S$ we have $R_{\mathbf{B}}^i = \emptyset$ (i.e. \mathbf{B} uses only distances from S).

Note also that whenever we have a structure \mathbf{B} satisfying conditions 1–4, we can equivalently view it as an S -edge-labelled graph, that is, a triple (B, E, d) such that $\{x, y\} \in E$ if and only if there is $i \in S$ such that $\{x, y\} \in R_{\mathbf{B}}^i$ and $d: E \rightarrow S$ is such that $d(x, y) = i$ if and only if $\{x, y\} \in R_{\mathbf{B}}^i$ (note that we write $d(x, y)$ instead of $d(\{x, y\})$).

Let $\mathbf{C} = (C, E, d)$ be an \mathbb{N}^+ -edge-labelled cycle (that is, (C, E) is a graph cycle) and enumerate the vertices as $C = \{c_1, \dots, c_n\}$ such that c_i and c_{i+1} are adjacent for every $1 \leq i \leq n$ (we identify c_{n+1} with c_1) and $d(c_1, c_n)$ is maximal. We say that \mathbf{C} is a *non-metric cycle* if

$$d(c_1, c_n) > \sum_{i=1}^{n-1} d(c_i, c_{i+1}).$$

The following claim is standard and was used many times (e.g. [Sol05, Neš07, Con19, HN19]). For a proof, see for example Observation 2.1 of [HKN19a].

Claim 4.12.3. *Let $S \subset \mathbb{N}^+$ be a finite set of distances and let $\mathbf{B} = (B, E, d)$ be a finite S -edge-labelled graph. There is a metric space \mathbf{M} on the same vertex set B such that the identity is a homomorphism-embedding $\mathbf{B} \rightarrow \mathbf{M}$ if and only if there is no non-metric cycle \mathbf{C} with a homomorphism-embedding $\mathbf{C} \rightarrow \mathbf{B}$. Moreover, $\text{Aut}(\mathbf{M}) = \text{Aut}(\mathbf{B})$, and if $\mathbf{K}_n \not\subseteq \mathbf{B}$, then $\mathbf{K}_n \not\subseteq \mathbf{M}$.*

In other words, we have a characterization of edge-labelled graphs with a completion to a metric space. Let's first see how this claim implies both strong amalgamation and local finiteness. For strong amalgamation, it is enough to observe that free amalgamations of metric spaces contain no non-metric cycles (indeed, if there was one, then we could find one in \mathbf{B}_1 or \mathbf{B}_2 , which would be a contradiction). For local finiteness observe that there are only finitely many non-metric cycles with distances from a finite set S , hence there is an upper bound n on the number of their vertices (which only depends on S) and we are done.

To conclude, we give a sketch of proof of the claim. Put $m = \max(2, \max S)$ and define function $d': B^2 \rightarrow \mathbb{N}$ as

$$d'(x, y) = \min(m, \min_{\mathbf{P} \text{ a path } x \rightarrow y \text{ in } \mathbf{B}} \|\mathbf{P}\|),$$

where by $\|\mathbf{P}\|$ we mean the sum of distances of \mathbf{P} . It is easy to check that (B, d') is a metric space, that it preserves automorphisms and that $d'|_E = d$ if and only if \mathbf{B} contains no (homomorphism-embedding of a) non-metric cycle. We remark that (B, d') is called the *shortest path completion* of \mathbf{B} in [HN19]. \square

Remark 4.12.1. The fact that we used \mathcal{E}_n as the base class in the proof of Proposition 4.12.2 was a matter of choice. We could also, for example, start with the class of all L -structures; the condition that every small enough substructure of \mathbf{B} has a completion in \mathcal{M}_n would also ensure that $R_{\mathbf{B}}^i$ are symmetric and irreflexive, that every pair of vertices is in at most one relation and that \mathbf{B} does not contain \mathbf{K}_n .

4.12.3 Structures with constants

We show how languages equipped with a permutation group can help us reduce EPPA for languages with constants (nulary functions) to languages without constants. Since the goal of this section is to illustrate applications of our main theorems, we will only construct EPPA-witnesses for structures where the constants behave in a special way.

To simplify the notation, if \mathbf{A} is a Γ_L -structure and c is a constant symbol of Γ_L , we will write $c_{\mathbf{A}}$ instead of $c_{\mathbf{A}}()$. Moreover, if the image of $c_{\mathbf{A}}$ is a singleton x (recall that, in general, functions go to the powerset of \mathbf{A}), we will write $c_{\mathbf{A}} = x$ instead of $c_{\mathbf{A}} = \{x\}$.

We first give a definition.

Definition 4.12.1. Let L be a language equipped with a permutation group Γ_L and let \mathbf{A} be a Γ_L -structure. We define the *constant trace* of \mathbf{A} , denoted by $\text{ctr}(\mathbf{A})$, as

$$\text{ctr}(\mathbf{A}) = \text{Cl}_{\mathbf{A}}(\emptyset).$$

In particular, $\text{ctr}(\mathbf{A})$ is a (possibly empty) Γ_L -structure.

For example, if Γ_L contains no constants, then the constant traces of all Γ_L -structures are empty. If Γ_L contains, say, two constants a and b and a binary relation E and \mathbf{A} is a Γ_L structure such that $a_{\mathbf{A}}$ and $b_{\mathbf{A}}$ are singletons, $a_{\mathbf{A}} \neq b_{\mathbf{A}}$, and moreover $(a_{\mathbf{A}}, b_{\mathbf{A}}) \in E_{\mathbf{A}}$, then $\text{ctr}(\mathbf{A})$ is the two-vertex Γ_L -structure with the corresponding relation E .

If Γ_L contains one constant symbol c and one unary function symbol F and \mathbf{A} is a Γ_L -structure containing a vertex x such that $c_{\mathbf{A}}$ is a singleton, $c_{\mathbf{A}} \neq x$ and $x \in F_{\mathbf{A}}(c_{\mathbf{A}})$, then $\text{ctr}(\mathbf{A})$ also contains x .

Theorem 4.12.4. *Let L be a language equipped with a permutation group Γ_L where the arity of every function is at most 1 and let \mathbf{A} be a finite Γ_L -structure. Let $L_{\mathcal{F}}^0$ be the set of all constant symbols of L . Assume the following:*

1. *For every $g \in \Gamma_L$ and every $c \in L_{\mathcal{F}}^0$ it holds that $g(c) = c$.*
2. *$L_{\mathcal{F}}^0$ is finite.*
3. *For every $c \in L_{\mathcal{F}}^0$ it holds that $c_{\mathbf{A}}$ is a singleton.*
4. *For every $c \neq c' \in L_{\mathcal{F}}^0$ it holds that $c_{\mathbf{A}} \neq c'_{\mathbf{A}}$.*
5. *For every $c \in L_{\mathcal{F}}^0$ and for every unary function $F \in L$ it holds that $F_{\mathbf{A}}(c_{\mathbf{A}}) = \emptyset$.*
6. *\mathbf{A} lies in a finite orbit of the action of Γ_L by relabelling.*

Then there is a finite Γ_L -structure \mathbf{B} which is an irreducible structure faithful coherent EPPA-witness for \mathbf{A} .

We again remark that our goal here was to keep the proof as simple as possible, a similar theorem can be proved with much weaker assumptions. In fact, one can obtain a category theory-like theorem which then makes it possible to lift the

main theorems of this paper to work for languages with constants. These results will appear elsewhere.

The structure of the proof will be similar to that of Proposition 4.5.1. That is, we will define a new language without constants and we will reduce the question to the question of EPPA in that language.

Proof of Theorem 4.12.4. Without loss of generality we will assume that L does not contain the symbol \star . Given a function $f: \{1, \dots, n\} \rightarrow L_{\mathcal{F}}^0 \cup \{\star\}$, we put $|f| = |\{i \in n : f(i) = \star\}|$. Observe that from assumptions 3 and 5 it follows that the vertex set of $\text{ctr}(\mathbf{A})$ is precisely $\{c_{\mathbf{A}} : c \in L_{\mathcal{F}}^0\}$.

Now, we define a language M without constant symbols. Let $R \in L$ be an n -ary relation symbol. For every function $f: \{1, \dots, n\} \rightarrow L_{\mathcal{F}}^0 \cup \{\star\}$ such that $|f| > 0$, we put an $|f|$ -ary relation symbol $R^{R,f}$ in M . Let $F \in L$ be a unary function symbol. For every $c \in S$, we put a unary relation symbol $R^{F,c}$ in M . We also put all unary function symbols of L into M .

Given $g \in \Gamma_L$, we define $\pi_g: M \rightarrow M$ as

$$\pi_g(T) = \begin{cases} g(T) & \text{if } T \text{ is a unary function symbol,} \\ R^{g(F),c} & \text{if } T = R^{F,c}, \text{ where } F \in L \text{ is a unary function symbol,} \\ R^{g(R),f} & \text{if } T = R^{R,f}, \text{ where } R \in L \text{ is a relation symbol.} \end{cases}$$

We put $\Gamma_M = \{\pi_g : g \in \Gamma_L\}$. Observe that Γ_M is a permutation group on M ($\pi_{gh} = \pi_g \pi_h$). We claim that $g \mapsto \pi_g$ is a group isomorphism: Clearly it is a surjective homomorphism, injectivity follows from the fact that M contains all unary function symbols of L , for every relation symbol $R \in L$ we have $R^{R,\star} \in M$ (where by \star we mean the constant \star function), and every $g \in \Gamma_L$ fixes $L_{\mathcal{F}}^0$ pointwise.

Given an m -tuple $(x_1, \dots, x_m) = \bar{x} \in A^m$ and a function $f: \{1, \dots, n\} \rightarrow L_{\mathcal{F}}^0 \cup \{\star\}$ such that $|f| = m$, we define \bar{x}^f to be the n -tuple (y_1, \dots, y_n) , where

$$y_i = \begin{cases} f(i)_{\mathbf{A}} & \text{if } f(i) \in L_{\mathcal{F}}^0, \\ x_j & \text{if } f(i) = \star \text{ and } |\{k < i : f(k) = \star\}| = j - 1. \end{cases}$$

Put $D = A \setminus \text{ctr}(\mathbf{A})$ (that is, the members of D are precisely the non-constant vertices of \mathbf{A}). We claim that for every n -tuple $(y_1, \dots, y_n) = \bar{y} \in A^n$, there is precisely one triple (m, \bar{x}, f) , where $m \in \mathbb{N}$, $\bar{x} \in D^m$ and f is a function $\{1, \dots, n\} \rightarrow L_{\mathcal{F}}^0 \cup \{\star\}$ with $|f| = m$, such that $\bar{y} = \bar{x}^f$. Indeed, put

$$f(i) = \begin{cases} c & \text{if } y_i = c_{\mathbf{A}} \text{ for some } c \in L_{\mathcal{F}}^0, \\ \star & \text{otherwise,} \end{cases}$$

$m = |f|$ and $x_i = y_j$, where j is chosen such that $f(j) = \star$ and $|\{k < j : f(k) = \star\}| = i - 1$.

Let \mathbf{C} be a Γ_M -structure such that C is disjoint from $K = \text{ctr}(\mathbf{A})$. We define a Γ_L -structure $T(\mathbf{C})$ as follows:

1. The vertex set of $T(\mathbf{C})$ is $C \cup K$.

2. The identity on K is an isomorphism between $\text{ctr}(\mathbf{A})$ and the structure induced by $T(\mathbf{C})$ on K (in particular, the constants are defined on K in $T(\mathbf{C})$ in the same way as in \mathbf{A}).

3. For every unary function $F \in L$ and every $x \in C$, we put

$$F_{T(\mathbf{C})}(x) = F_{\mathbf{C}}(x) \cup \{c_{T(\mathbf{C})} : c \in L_{\mathcal{F}}^0 \text{ and } x \in R_{\mathbf{C}}^{F,c}\}.$$

4. For every relation $R^{R,f} \in M$ and every $\bar{x} \in R_{\mathbf{C}}^{R,f}$, we put $\bar{x}^f \in R_{T(\mathbf{C})}$.

Note that (π_g, α) is an embedding of Γ_M -structures $\mathbf{E} \rightarrow \mathbf{F}$, if and only if $(g, \alpha \cup \text{id}_K)$ is an embedding $T(\mathbf{E}) \rightarrow T(\mathbf{F})$. This follows directly from the construction. It also implies that \mathbf{E} lies in a finite orbit of the action of Γ_M by relabelling if and only if $T(\mathbf{E})$ lies in a finite orbit of the action of Γ_L by relabelling.

Next, we define a Γ_M -structure \mathbf{D} such that $T(\mathbf{D}) = \mathbf{A}$. We put the vertex set of \mathbf{D} to be D , the relations and functions are defined as follows:

1. For every unary function $F \in L$ and every vertex $x \in D$, we put $F_{\mathbf{D}}(x) = F_{\mathbf{A}}(x) \setminus \text{ctr}(\mathbf{A})$.
2. For every unary function $F \in L$, every vertex $x \in D$ and every constant $c \in L_{\mathcal{F}}^0$, we put $x \in R_{\mathbf{D}}^{F,c}$ if and only if $c_{\mathbf{A}} \in F_{\mathbf{A}}(x)$.
3. For every n -ary relation $R \in L$ and every $\bar{y} \in R_{\mathbf{A}}$ such that $\bar{y} = \bar{x}^f$, where $\bar{x} \in D^m$, $f: \{1, \dots, n\} \rightarrow L_{\mathcal{F}}^0 \cup \{\star\}$ and $m \geq 1$, we put $\bar{x} \in R_{\mathbf{D}}^{R,f}$.

It is straightforward to verify that indeed $T(\mathbf{D}) = \mathbf{A}$.

Since Γ_M is a language where all functions are unary, by Theorem 4.1.1 we get an irreducible structure faithful coherent EPPA-witness \mathbf{C} for \mathbf{D} . Without loss of generality we can assume that C is disjoint from K . We claim that $\mathbf{B} = T(\mathbf{C})$ is an irreducible structure faithful coherent EPPA-witness for \mathbf{A} .

Let (g, α) be a partial automorphism of \mathbf{A} . This implies that $(\pi_g, \alpha \upharpoonright_D)$ is a partial automorphism of \mathbf{D} , which by the assumption extends to an automorphism (π_g, θ) of \mathbf{C} . This implies that $(g, \theta \cup \text{id}_K)$ is an automorphism of \mathbf{B} extending (g, α) . Since the extensions in \mathbf{C} can be chosen to be coherent, by the construction we get coherence also for \mathbf{B} .

To get irreducible structure faithfulness of \mathbf{B} , observe that if $\mathbf{P} \subseteq \mathbf{C}$ is the free amalgamation of \mathbf{P}_1 and \mathbf{P}_2 over \mathbf{Q} , then $T(\mathbf{P})$ is the free amalgamation of $T(\mathbf{P}_1)$ and $T(\mathbf{P}_2)$ over $T(\mathbf{Q})$. This follows from the fact that functions in a Γ_M -structure \mathbf{X} are subsets of the corresponding functions in $T(\mathbf{X})$ and if, for $n \geq 2$, an n -tuple is in a relation in \mathbf{X} , then it is a sub-tuple of a tuple in a relation of $T(\mathbf{X})$.

Taking the contrapositive, this means that if \mathbf{I} is an irreducible substructure of \mathbf{B} , then \mathbf{C} induces an irreducible substructure on $I \setminus K$. Hence, there is an automorphism $(\pi_g, \alpha): \mathbf{C} \rightarrow \mathbf{C}$ sending $I \setminus K$ to A and thus $(g, \alpha \cup \text{id}_K)$ is an automorphism of \mathbf{B} such that $(g, \alpha \cup \text{id}_K)(I) \subseteq A$. \square

4.12.4 EPPA for special non-unary functions

One of our motivations for introducing languages equipped with a permutation group was that it gives a nice formalism to stack several EPPA constructions on top of each other, thereby allowing to prove coherent EPPA for certain classes with non-unary functions. We conclude this paper with two examples of this. This section can be seen as an introduction to Section 4.12.5.

The following theorem is a variant of Ivanov's observation that permomorphisms of Herwig [Her98, Lemma 1] can be used to prove EPPA of equivalence relations on n -tuples [Iva15]:

Theorem 4.12.5. *Let L be a finite language consisting of two unary relations U, V and functions F^1, \dots, F^n , each of arity at least 1. Let \mathcal{C} be the class of all finite L -structures \mathbf{A} satisfying the following:*

1. $U_{\mathbf{A}} \cap V_{\mathbf{A}} = \emptyset$ and $U_{\mathbf{A}} \cup V_{\mathbf{A}} = A$,
2. for every $1 \leq i \leq n$ it holds that $\text{Dom}(F_{\mathbf{A}}^i) \subseteq (U_{\mathbf{A}})^{a(F^i)}$ and $\text{Range}(F_{\mathbf{A}}^i) \subseteq V_{\mathbf{A}}$.

(Equivalently, structures in \mathcal{C} can be viewed as 2-sorted structures where all the functions go from the first sort to the other.) Then \mathcal{C} has irreducible structure faithful coherent EPPA.

Proof. Fix $\mathbf{A} \in \mathcal{C}$. We will construct $\mathbf{B} \in \mathcal{C}$ such that \mathbf{B} is the desired EPPA-witness. Towards that, we define a language L^* consisting of an $a(F^i)$ -ary relation $R^{i,v}$ for every $1 \leq i \leq n$ and every $v \in V_{\mathbf{A}}$. Let Γ_{L^*} be the permutation group obtained by the natural action of $\text{Sym}(V_{\mathbf{A}})$ on L^* . Next we define an Γ_{L^*} -structure \mathbf{A}_0 such that the vertex set of \mathbf{A}_0 is precisely $U_{\mathbf{A}}$ and for every tuple \bar{x} of vertices of \mathbf{A}_0 and every relation $R^{i,v} \in L^*$ we put $\bar{x} \in R_{\mathbf{A}_0}^{i,v}$ if and only if $v \in F_{\mathbf{A}}^i(\bar{x})$. Let \mathbf{B}_0 be an irreducible structure faithful coherent EPPA-witness for \mathbf{A}_0 (obtained for example using Theorem 4.1.1). Without loss of generality we can assume that $\mathbf{A}_0 \subseteq \mathbf{B}_0$.

Next we reconstruct an L -structure \mathbf{B} using \mathbf{B}_0 as a template as follows:

1. The vertex set of \mathbf{B} is the disjoint union $B_0 \cup V_{\mathbf{A}}$.
2. $U_{\mathbf{B}} = B_0$ and $V_{\mathbf{B}} = V_{\mathbf{A}}$.
3. For every $1 \leq i \leq n$, every $v \in V_{\mathbf{A}}$ and every tuple \bar{x} from B_0 we put $v \in F_{\mathbf{B}}^i(\bar{x})$ if and only if $\bar{x} \in R_{\mathbf{B}_0}^{i,v}$.

Clearly, $\mathbf{B} \in \mathcal{C}$. Since $\mathbf{A}_0 \subseteq \mathbf{B}_0$, we get that $A \subseteq B$. To see that \mathbf{A} is in fact a substructure of \mathbf{B} , observe that $U_{\mathbf{A}} = U_{\mathbf{B}} \cap A$, $V_{\mathbf{A}} = V_{\mathbf{B}}$ and whenever \bar{x} is a tuple of vertices from $U_{\mathbf{A}}$, $v \in V_{\mathbf{A}}$ and $1 \leq i \leq n$, then $v \in F_{\mathbf{A}}^i(\bar{x})$ if and only if $\bar{x} \in R_{\mathbf{A}_0}^{i,v}$ (by the construction of \mathbf{A}), which happens if and only if $\bar{x} \in R_{\mathbf{B}_0}^{i,v}$ (since $\mathbf{A}_0 \subseteq \mathbf{B}_0$) and this is true if and only if $v \in F_{\mathbf{B}}^i(\bar{x})$ (by the construction of \mathbf{B}). Hence indeed $\mathbf{A} \subseteq \mathbf{B}$.

Now we show how to construct an automorphism of \mathbf{B} from an automorphism of \mathbf{B}_0 and a permutation of $V_{\mathbf{A}}$. Let f' be a permutation of $V_{\mathbf{A}}$ and let $f = (f_{L^*}, f_{B_0})$ be an automorphism of \mathbf{B} such that f_{L^*} is induced by f' . Put $\theta = f_{B_0} \cup f'$. We claim that θ is an automorphism of \mathbf{B} . Clearly, θ is a bijection

$B \rightarrow B$ which preserves the unary relations. Given an arbitrary $1 \leq i \leq n$, an arbitrary tuple \bar{x} of vertices from B_0 and an arbitrary $v \in V_{\mathbf{A}}$, we know that $v \in F_{\mathbf{B}}^i(\bar{x})$ if and only if $\bar{x} \in R_{\mathbf{B}_0}^{i,v}$ (by the construction of \mathbf{B}), which happens if and only if $f_{B_0}(\bar{x}) \in f_{L^*}(R^{i,v})_{\mathbf{B}_0} = R_{\mathbf{B}_0}^{i,f'(v)}$ (as f is an automorphism and f_{L^*} is induced by f'), and by the construction of \mathbf{B} it is equivalent to $f'(v) \in F_{\mathbf{B}}^i(f'_{B_0}(\bar{x}))$. Hence θ is an automorphism of \mathbf{B} .

To see that \mathbf{B} is irreducible structure faithful, it is enough to observe that if $\mathbf{C} \subseteq \mathbf{B}$ is irreducible, then either C consists of a single vertex of $V_{\mathbf{A}}$, or $\mathbf{C} = \text{Cl}_{\mathbf{B}}(C \cap U_{\mathbf{B}})$ and for every pair $x \neq y \in C \cap U_{\mathbf{B}}$ there is a tuple \bar{x} of vertices of \mathbf{C} containing both x and y , and $1 \leq i \leq n$ such that $F_{\mathbf{C}}^i(\bar{x}) \neq \emptyset$. Consequently, \mathbf{B}_0 induces an irreducible substructure on $C \cap U_{\mathbf{B}}$. By irreducible structure faithfulness there is an automorphism $f = (f_{L^*}, f_{B_0})$ of \mathbf{B}_0 such that $f(C \cap U_{\mathbf{B}}) \subseteq A_0$. Let f' be an arbitrary permutation of $V_{\mathbf{A}}$ inducing f_{L^*} and let θ be the automorphism of \mathbf{B} constructed from f and f' in the previous paragraph. Clearly, $\theta(C \cap U_{\mathbf{B}}) \subseteq A_0$, and therefore $\theta(\mathbf{C}) = \theta(\text{Cl}_{\mathbf{B}}(C \cap U_{\mathbf{B}})) \subseteq \mathbf{A}$. This finishes the proof of irreducible structure faithfulness.

Next we prove that \mathbf{B} is an EPPA-witness for \mathbf{A} . Let φ be a partial automorphism of \mathbf{A} . Remember that φ preserves the unary relations. Let f' be the coherent extension of $\varphi \upharpoonright_{V_{\mathbf{A}}}$ to a permutation of $V_{\mathbf{A}}$ obtained using Proposition 4.2.8. Let $f_{L^*} \in \Gamma_{L^*}$ be induced by f' and put $\varphi_0 = (f_{L^*}, \varphi \upharpoonright_{A_0})$. Observe that φ_0 is a partial automorphism of \mathbf{A}_0 and extend it to an automorphism $f = (f_{L^*}, f_{B_0})$ of \mathbf{B}_0 (in a coherent way). Put $\tilde{\varphi} = f_{B_0} \cup f'$. By the previous paragraphs, $\tilde{\varphi}$ is an automorphism of \mathbf{B}_0 . Moreover, since $\varphi \upharpoonright_{V_{\mathbf{A}}} \subseteq f'$ and $\varphi \upharpoonright_{A_0} \subseteq f_{B_0}$, we get that $\tilde{\varphi}$ extends φ .

To finish the proof, note that since both f and f' were chosen to be coherent, $\tilde{\varphi}$ is coherent as well. \square

Remark 4.12.2. This construction can be carried out more generally for infinitely many functions, more than 2 unary marks (as long as all functions go in one direction) and more complicated structures living on each unary mark (as long as the whole multi-sorted structure still lies in a finite orbit of the relabelling action). This will appear elsewhere. In the next section, we adapt this construction for a class which does not a priori look multi-sorted.

4.12.5 EPPA for k -orientations with d -closures

In this section we extend the construction from Section 4.12.4 and prove EPPA for the class of all k -orientations with d -closures, thereby confirming a conjecture from [EHN19]. We only define the relevant classes and prove EPPA for them here, to get more context (for example the connection with Hrushovski's predimension constructions and the importance for the structural Ramsey theory), see [EHN19].

Let \mathbf{G} be an oriented graph (that is, if there is an edge from vertex u to vertex v then there is no edge from v to u). We say that it is a k -orientation if the out-degree of every vertex is at most k . We say that a vertex $x \in G$ is a *root* if its out-degree is strictly smaller than k . Let \mathcal{D}^k be the class of all finite k -orientations. While \mathcal{D}^k is not an amalgamation class, there are two natural expansions which do have the free amalgamation property:

Definition 4.12.2. Let L be the graph language with a single binary relation E and let L_s be its expansion by a unary function symbol F .

Let \mathbf{G} be a k -orientation. By $s(\mathbf{G})$ we denote the L_s -expansion of \mathbf{G} putting

$$F_{s(\mathbf{G})}(x) = \{y \in G : y \text{ is reachable from } x \text{ by an oriented path}\}.$$

Here, an oriented path from x to y is a sequence $x = v_1, v_2, \dots, v_m = y$ with $m \geq 1$ such that for every $1 \leq i < m$ it holds that $(v_i, v_{i+1}) \in E_{\mathbf{G}}$. Put $\mathcal{D}_s^k = \{s(\mathbf{G}) : \mathbf{G} \in \mathcal{D}^k\}$.

Recall that $\text{Cl}_{\mathbf{A}}(x)$ denotes the smallest substructure of \mathbf{A} containing x and is called the closure of x in \mathbf{A} . For $\mathbf{G} \in \mathcal{D}_s^k$ and $y \in G$, we denote by $\text{roots}(y)$ the set of all roots of \mathbf{G} which are in $\text{Cl}_{\mathbf{G}}(y)$. Define $\mathcal{D}_{s^+}^k$ to be the subclass of \mathcal{D}_s^k such that $\mathbf{G} \in \mathcal{D}_{s^+}^k$ if and only if for every $y \in G$ it holds that $\text{roots}(y) \neq \emptyset$.

Definition 4.12.3. Let L_d be an expansion of L_s adding an n -ary function symbol F^n for every $n \geq 1$.

Given $\mathbf{G} \in \mathcal{D}_s^k$, we denote by $d(\mathbf{G})$ the L_d -expansion of \mathbf{G} putting $F_{d(\mathbf{G})}^n(x_1, \dots, x_n) = \emptyset$ if (x_1, \dots, x_n) is not a tuple of distinct roots and

$$F_{d(\mathbf{G})}^n(x_1, \dots, x_n) = \{y \in G : \text{roots}(y) = \{x_1, \dots, x_n\}\}$$

if (x_1, \dots, x_n) is a tuple of distinct roots. Put $\mathcal{D}_d^k = \{d(\mathbf{G}) : \mathbf{G} \in \mathcal{D}_{s^+}^k\}$.

Note that in the definition of \mathcal{D}_d^k we are only considering members of $\mathcal{D}_{s^+}^k$. The reason is that if there was a vertex with $\text{roots}(y) = \emptyset$, it would be in the closure of the empty set, i.e. we would need to add constants. It is possible to do so, but it would make the construction a bit complicated and for the applications we have in mind it does not make any difference.

It is easy to see that \mathcal{D}_s^k is a free amalgamation class. Combining with Corollary 4.1.4, we get the following theorem proved by Evans, Hubička and Nešetřil [EHN19, EHN21].

Theorem 4.12.6. \mathcal{D}_s^k has irreducible structure faithful coherent EPPA for every $k \geq 1$.

It is again straightforward to verify (and it was done in [EHN19]) that \mathcal{D}_d^k is a free amalgamation class. Since it contains non-unary functions, the results of this paper cannot be applied directly to prove that \mathcal{D}_d^k has irreducible structure faithful coherent EPPA. However, we can use the fact that the non-unary functions go from root vertices to non-root vertices and show the following theorem, which was conjectured to hold in [EHN19, Conjecture 7.5].

Theorem 4.12.7. \mathcal{D}_d^k has irreducible structure faithful coherent EPPA for every $k \geq 1$.

In the rest of this section, we will prove this theorem. The proof is based on the following observation: Let S be a set consisting of root vertices only, let S_1 be the L_s -closure of S (i.e. we ignore the F^n functions) and let S_2 be the L_d -closure of S (i.e. we also consider the F^n functions). Then the root vertices in S_1 are precisely the root vertices in S_2 . Consequently, if one is interested in root vertices

only, all closures are unary, even in the presence of higher-arity functions. Thus, we can view structures from \mathcal{D}_d^k as two-sorted structures (one sort being the roots and the other being the non-roots) in which all non-unary functions go from one sort to the other, which allows us to use a similar structure of arguments as in Section 4.12.4.

Fix $\mathbf{A} \in \mathcal{D}_d^k$ and denote by \mathbf{A}_0 its L_s -reduct (so $\mathbf{A}_0 \in \mathcal{D}_s^k$). Let $\mathbf{B}_0 \in \mathcal{D}_s^k$ be an irreducible structure faithful coherent EPPA-witness for \mathbf{A}_0 given by Theorem 4.12.6.

Let \mathfrak{P} be the set of all pairs $(x, (x_1, \dots, x_n))$ such that x is a non-root vertex of \mathbf{B}_0 , (x_1, \dots, x_n) is a tuple of distinct root vertices of \mathbf{B}_0 and $\text{roots}_{\mathbf{B}_0}(x) = \{x_1, \dots, x_n\}$. Note that we have such a pair for each possible permutation of $\{x_1, \dots, x_n\}$. Given $P = (x, (x_1, \dots, x_n)) \in \mathfrak{P}$, we define $\pi(P) = x$ to be the *projection* and put $|P| = n$.

Denote by L^+ the expansion of L_s adding a $|P|$ -ary relation symbol R^P for every $P \in \mathfrak{P}$ and a $(|P| + 1)$ -ary relation symbol E^P for every $P \in \mathfrak{P}$. Let Γ_{L^+} be the permutation group on L^+ consisting of all permutations of the R^P and E^P symbols induced by the natural action of $\text{Aut}(\mathbf{B}_0)$ on \mathfrak{P} . In particular, E and F are fixed by Γ_{L^+} .

Denote by \mathbf{A}_1 the Γ_{L^+} -structure created from \mathbf{A}_0 by removing all non-root vertices, keeping the edges between root vertices, putting $F_{\mathbf{A}_1}(v) = F_{\mathbf{A}_0}(v) \cap A_1$, adding $(x_1, \dots, x_n) \in R_{\mathbf{A}_1}^{(x, (x_1, \dots, x_n))}$ if and only if x is a non-root vertex of \mathbf{A}_0 and $\text{roots}_{\mathbf{A}_0}(x) = \{x_1, \dots, x_n\}$, and adding $(a, x_1, \dots, x_n) \in E_{\mathbf{A}_1}^{(x, (x_1, \dots, x_n))}$ if and only if $(x_1, \dots, x_n) \in R_{\mathbf{A}_1}^{(x, (x_1, \dots, x_n))}$, a is a root vertex of \mathbf{A}_0 and $(a, x) \in E_{\mathbf{A}_0}$. Let \mathbf{B}_1 be an irreducible structure faithful coherent EPPA-witness for \mathbf{A}_1 given by Theorem 4.1.1.

We will now reconstruct an L_d -structure $\mathbf{B} \in \mathcal{D}_d^k$ from \mathbf{B}_1 such that \mathbf{B} will be an irreducible structure faithful coherent EPPA-witness for \mathbf{A} . The general idea is to put back the non-root vertices according to the R^P and E^P relations using \mathbf{B}_0 as a template.

Let \mathcal{T}_0 be the set consisting of all pairs (P, \bar{x}) such that $P \in \mathfrak{P}$, \bar{x} is a tuple of vertices of \mathbf{B}_1 and $\bar{x} \in R_{\mathbf{B}_1}^P$. We say that $(P, \bar{x}) \sim (P', \bar{x}')$ if $\pi(P) = \pi(P')$ and \bar{x} and \bar{x}' are different permutations of the same set. Let \mathcal{T} consist of exactly one (arbitrary) member of each equivalence class of \sim on \mathcal{T}_0 .

Put $B = B_1 \cup \mathcal{T}$. For $u, v \in B$, we put $(u, v) \in E_{\mathbf{B}}$ if and only if one of the following holds:

- C1 $u, v \in B_1$ and $(u, v) \in E_{\mathbf{B}_1}$,
- C2 $u \in B_1$, $v = ((x, (x_1, \dots, x_n)), (w_1, \dots, w_n)) \in \mathcal{T}$ and $(u, w_1, \dots, w_n) \in E_{\mathbf{B}_1}^{(x, (x_1, \dots, x_n))}$,
- C3 $u = ((x, (x_1, \dots, x_n)), (w_1, \dots, w_n)) \in \mathcal{T}$, $v \in B_1$, there is $1 \leq i \leq n$ such that $v = w_i$ and $(x, x_i) \in E_{\mathbf{B}_0}$, or
- C4 $u = ((x, (x_1, \dots, x_n)), (w_1, \dots, w_n)) \in \mathcal{T}$, $v = ((y, (y_1, \dots, y_m)), (t_1, \dots, t_m)) \in \mathcal{T}$, $\{t_1, \dots, t_m\} \subseteq \{w_1, \dots, w_n\}$ and $(x, y) \in E_{\mathbf{B}_0}$.

For every $x \in B$ we put

$$F_{\mathbf{B}}(x) = \{y \in B : y \text{ is reachable from } x \text{ in } B \text{ by an oriented path}\}.$$

Finally, we put $F_{\mathbf{B}}^n(x_1, \dots, x_n) = \emptyset$ if (x_1, \dots, x_n) is not a tuple of distinct vertices of \mathbf{B}_1 and $F_{\mathbf{B}}^n(x_1, \dots, x_n) = \{y \in B : \text{support}(y) = \{x_1, \dots, x_n\}\}$ if (x_1, \dots, x_n) is a tuple of distinct vertices of \mathbf{B}_1 . Here, $\text{support}(v)$ is defined as follows:

1. If $v \in B_1$, we put $\text{support}(v) = \text{Cl}_{\mathbf{B}_1}(v)$.
2. Otherwise $v \in \mathcal{T}$ and thus $v = (P, \bar{x})$ for some choice of P and \bar{x} . In this case we put $\text{support}(v) = \text{Cl}_{\mathbf{B}_1}(\bar{x})$ (where by $\text{Cl}_{\mathbf{B}_1}(\bar{x})$ we mean the smallest substructure of \mathbf{B}_1 containing all vertices from \bar{x}).

Depending on the context, we may consider $\text{support}(v)$ to be a substructure of \mathbf{B}_1 or just a subset of B_1 .

Lemma 4.12.8. *Let $(w_1, \dots, w_n) \in R_{\mathbf{B}_1}^{(x, (x_1, \dots, x_n))}$. There is automorphism f of \mathbf{B}_1 such that $f(\{w_1, \dots, w_n\}) \subseteq A_1$. If there is also $u \in B_1$ such that $(u, w_1, \dots, w_n) \in E_{\mathbf{B}_1}^{(x, (x_1, \dots, x_n))}$ then f can be chosen so that also $f(u) \subseteq A_1$.*

Moreover, whenever f is an automorphism of \mathbf{B}_1 such that $f(\{w_1, \dots, w_n\}) \subseteq A_1$ and f' is an automorphism of \mathbf{B}_0 such that f_L is induced by f' then the following hold:

1. $(f(w_1), \dots, f(w_n)) \in R_{\mathbf{A}_1}^{(f'(x), (f'(x_1), \dots, f'(x_n)))}$,
2. for every $1 \leq i \leq n$ it holds that $f(w_i) = f'(x_i)$, and
3. $f'(\{x, x_1, \dots, x_n\}) \subseteq A_0$,
4. $\text{roots}_{\mathbf{B}_0}(f'(x)) = f'(\{x_1, \dots, x_n\})$.

If there is also $u \in B_1$ such that $f(u) \in A_1$, then $(u, w_1, \dots, w_n) \in E_{\mathbf{B}_1}^{(x, (x_1, \dots, x_n))}$ if and only if $(f(u), f'(x)) \in E_{\mathbf{A}_0}$.

Proof. The first part is straightforward: Since (w_1, \dots, w_n) (or (u, w_1, \dots, w_n) respectively) is in a relation of \mathbf{B}_1 , we get that $\text{Cl}_{\mathbf{B}_1}(\{w_1, \dots, w_n\})$ (or $\text{Cl}_{\mathbf{B}_1}(\{u, w_1, \dots, w_n\})$ respectively) is an irreducible substructure of \mathbf{B}_1 , and so there is an automorphism f of \mathbf{B}_1 with the desired properties by irreducible structure faithfulness of \mathbf{B}_1 .

Suppose now that we have such automorphisms f and f' . The first statement is just rephrasing that f is an automorphism with f_L induced by f' . From the construction of \mathbf{A}_1 it follows that whenever $(t_1, \dots, t_n) \in R_{\mathbf{A}_1}^{(y, (y_1, \dots, y_n))}$, then $t_i = y_i$ for every $1 \leq i \leq n$, which implies the second point. The third point is a direct consequence of the second point and the construction of \mathbf{A}_1 . To see the fourth point, note that \mathbf{A}_0 is a substructure of \mathbf{B}_0 , hence $\text{roots}_{\mathbf{B}_0}(f'(x)) = \text{roots}_{\mathbf{A}_0}(f'(x)) = f'(\{x_1, \dots, x_n\})$.

If there is also $u \in B_1$ such that $f(u) \subseteq A_1$, then directly from the definition of the relations on \mathbf{A}_1 it follows that $(u, w_1, \dots, w_n) \in E_{\mathbf{B}_1}^{(x, (x_1, \dots, x_n))}$ if and only if $(f(u), f'(x)) \in E_{\mathbf{A}_0}$. \square

Observation 4.12.9. *If $v = ((x, (x_1, \dots, x_n)), (w_1, \dots, w_n)) \in \mathcal{T}$ then*

$$\text{support}(v) = \{w_1, \dots, w_n\}.$$

Proof. From the definition of \mathcal{T} , we know that $(w_1, \dots, w_n) \in R_{\mathbf{B}_1}^{(x, (x_1, \dots, x_n))}$. So, by Lemma 4.12.8, we get automorphisms f and f' such that $f(w_i) = f'(x_i) \in A_1$ for every i and $f'(\{x_1, \dots, x_n\})$ are the only roots reachable from $f'(x)$ in \mathbf{B}_0 . Consequently, they are all the more so the only roots reachable from $f'(\{x_1, \dots, x_n\}) = f(\{w_1, \dots, w_n\})$ in \mathbf{A}_0 and hence $\text{Cl}_{\mathbf{B}_1}(f(\{w_1, \dots, w_n\})) = f(\{w_1, \dots, w_n\})$. Sending it back by f^{-1} then gives

$$\text{support}(v) = \text{Cl}_{\mathbf{B}_1}(\{w_1, \dots, w_n\}) = \{w_1, \dots, w_n\}.$$

□

The following observation follows directly from the construction of \mathbf{B} .

Observation 4.12.10. *Whenever $(u, v) \in E_{\mathbf{B}}$, we have that*

$$\text{support}(v) \subseteq \text{support}(u).$$

Proof. We have to distinguish four cases:

1. If $u, v \in B_1$, by C1 we know that $(u, v) \in E_{\mathbf{B}_1}$. This implies that $v \in \text{Cl}_{\mathbf{B}_1}(u)$, so $\text{Cl}_{\mathbf{B}_1}(v) \subseteq \text{Cl}_{\mathbf{B}_1}(u)$ and hence $\text{support}(v) \subseteq \text{support}(u)$.
2. If $u \in B_1$, $v = ((x, (x_1, \dots, x_n)), (w_1, \dots, w_n)) \in \mathcal{T}$, by C2 we know that $(u, w_1, \dots, w_n) \in E_{\mathbf{B}_1}^{(x, (x_1, \dots, x_n))}$. By definition, $\text{support}(u) = \text{Cl}_{\mathbf{B}_1}(u)$ and $\text{support}(v) = \{w_1, \dots, w_n\}$. Using Lemma 4.12.8 we get automorphisms f and f' such that $\text{roots}_{\mathbf{A}_0}(f'(x)) = f'(\{x_1, \dots, x_n\})$ and $(f(u), f'(x)) \in E_{\mathbf{A}_0}$. So

$$f(\{w_1, \dots, w_n\}) = f'(\{x_1, \dots, x_n\}) \subseteq F_{\mathbf{A}_0}(f'(x)) \subseteq F_{\mathbf{A}_0}(f(u)).$$

Consequently, $\{w_1, \dots, w_n\} \subseteq F_{\mathbf{A}_1}(u)$, and hence

$$\text{support}(v) = \{w_1, \dots, w_n\} \subseteq \text{Cl}_{\mathbf{B}_1}(u) = \text{support}(u).$$

3. If $u = ((x, (x_1, \dots, x_n)), (w_1, \dots, w_n)) \in \mathcal{T}$ and $v \in B_1$, by C3 we have $1 \leq i \leq n$ such that $v = w_i$. Then

$$\text{support}(v) = \text{Cl}_{\mathbf{B}_1}(w_i) \subseteq \text{Cl}_{\mathbf{B}_1}(\{w_1, \dots, w_n\}) = \text{support}(u).$$

4. If $u = ((x, (x_1, \dots, x_n)), (w_1, \dots, w_n)) \in \mathcal{T}$ and $v = ((y, (y_1, \dots, y_m)), (t_1, \dots, t_m)) \in \mathcal{T}$, by C4 we get immediately that $\text{support}(v) = \{t_1, \dots, t_m\} \subseteq \{w_1, \dots, w_n\} = \text{support}(u)$.

□

Our next goal is to show that $\mathbf{B} \in \mathcal{D}_d^k$. Towards that direction we define the following procedure to map portions of \mathbf{B} to substructures of \mathbf{A} . Given a vertex $v \in B$ and an automorphism $f = (f_L, f_{B_1})$ of \mathbf{B}_1 such that $f(\text{support}(v))$ is a substructure of \mathbf{A}_1 we define f -correspondence $c_f(v) \in A$ as follows:

1. If $v \in B_1$, we put $c_f(v) = f_{B_1}(v)$.

2. Otherwise $v \in \mathcal{T}$. Then $v = (P, \bar{x})$ (for some choice of P and \bar{x}) and we put $c_f(v) = \pi(f_L(P))$. (Here, by $f_L(P)$ we mean the P' such that $f_L(R^P) = R^{P'}$.)

Claim 4.12.11 (on correspondence). *Let $f = (f_L, f_{B_1})$ be an automorphism of \mathbf{B}_1 and $v \neq v' \in B$ such that both $f(\text{support}(v))$ and $f(\text{support}(v'))$ are substructures of \mathbf{A}_1 . Then*

P1 $c_f(v) \neq c_f(v')$.

P2 $(v, v') \in E_{\mathbf{B}}$ if and only if $(c_f(v), c_f(v')) \in E_{\mathbf{A}}$.

Proof. Let f' be an automorphism of \mathbf{B}_0 inducing f_L . If $v, v' \in B_1$ then we know that $c_f(v) = f(v) \neq f(v') = c_f(v')$ and P1 follows. If precisely one of v, v' is in $v \in B_1$ then it follows that precisely one of $c_f(v), c_f(v')$ is a root of \mathbf{A} and P1 follows as well.

So $v = (P, \bar{x}) \in \mathcal{T}$ and $v' = (P', \bar{x}') \in \mathcal{T}$. We will show that $\pi(P) \neq \pi(P')$, which would imply that $c_f(v) = f'(\pi(P)) \neq f'(\pi(P')) = c_f(v')$, hence P1 holds. For a contradiction, suppose that $\pi(P) = \pi(P')$. By the construction we have that $\bar{x} \in R_{\mathbf{B}_1}^P$ and $\bar{x}' \in R_{\mathbf{B}_1}^{P'}$ and Lemma 4.12.8 then implies that $f(\bar{x}) = \text{roots}_{\mathbf{B}_0}(f'(\pi(P))) = \text{roots}_{\mathbf{B}_0}(f'(\pi(P'))) = f(\bar{x}')$ (where we consider \bar{x} and \bar{x}' as sets), hence \bar{x} and \bar{x}' are different permutations of the same set. This is, however, in a contradiction with the definition of \mathcal{T} and the fact that $v \neq v'$, which finishes the proof of P1.

If $v, v' \in B_1$, P2 immediately follows from C1. If $v \in B_1$ and $v' = ((x, (x_1, \dots, x_n)), (w_1, \dots, w_n)) \in \mathcal{T}$, we know by C2 that $(v, v') \in E_{\mathbf{B}}$ if and only if $(v, w_1, \dots, w_n) \in E_{\mathbf{B}_1}^{(x, (x_1, \dots, x_n))}$. By Lemma 4.12.8 we know that this happens if and only if $(c_f(v), c_f(v')) = (f(v), f'(x)) \in E_{\mathbf{A}_0} = E_{\mathbf{A}}$.

Now suppose that $v = ((x, (x_1, \dots, x_n)), (w_1, \dots, w_n)) \in \mathcal{T}$ and $v' \in B_1$. If there is $1 \leq i \leq n$ such that $v' = w_i$, we know that $(v, v') \in E_{\mathbf{B}}$ if and only if $(x, x_i) \in E_{\mathbf{B}_0}$ by C3. Lemma 4.12.8 tells us that $f'(x_i) = f(w_i) = f(v')$, and since f' is an automorphism of \mathbf{B}_0 , we get that $(x, x_i) \in E_{\mathbf{B}_0}$ if and only if $(c_f(v), c_f(v')) = (f'(x), f'(x_i)) \in E_{\mathbf{A}_0} = E_{\mathbf{A}}$. So $v' \neq w_i$ for any i . In that case $(v, v') \notin E_{\mathbf{B}}$ by C3. But, again using Lemma 4.12.8, we know that $\text{roots}_{\mathbf{B}_0}(f'(x)) = f(\{w_1, \dots, w_n\})$, and as $f(v')$ is a root of \mathbf{B}_0 , we know that $f(v') \notin \text{roots}_{\mathbf{B}_0}(f'(x))$, so in particular $(c_f(v), c_f(v')) = (f'(x), f(v')) \notin E_{\mathbf{A}}$.

Finally, suppose that $v = ((x, \bar{x}), \bar{w}) \in \mathcal{T}$ and $v' = ((y, \bar{y}), \bar{t}) \in \mathcal{T}$. If $(c_f(v), c_f(v')) = (f'(x), f'(y)) \in E_{\mathbf{A}}$, then we know that (as sets)

$$f(\bar{t}) = f'(\bar{y}) = \text{roots}_{\mathbf{A}}(f'(y)) \subseteq \text{roots}_{\mathbf{A}}(f'(x)) = f'(\bar{x}) = f(\bar{w}),$$

so $\bar{t} \subseteq \bar{w}$ (as sets) and hence $(v, v') \in E_{\mathbf{B}}$ by C4. So $(c_f(v), c_f(v')) = (f'(x), f'(y)) \notin E_{\mathbf{A}}$. But then, as f' is an automorphism, we get that $(x, y) \notin E_{\mathbf{A}}$ and thus $(v, v') \notin E_{\mathbf{B}}$ by C4. \square

Corollary 4.12.12. *Let v_1, \dots, v_m be a sequence of vertices of \mathbf{B} such that $(v_i, v_{i+1}) \in E_{\mathbf{B}}$ for every $1 \leq i < m$ and let f be an automorphism of \mathbf{B}_1 such that $\text{support}(v_1) \subseteq \mathbf{A}_1$. Then $c_f(v_i) \in A$ for every $1 \leq i \leq m$ and $(c_f(v_i), c_f(v_{i+1})) \in E_{\mathbf{A}}$ for every $1 \leq i < m$. Moreover, such an automorphism always exists.*

Proof. Put $\mathbf{S} = \text{support}(v_1)$. First note that \mathbf{S} is an irreducible substructure of \mathbf{B}_1 (it is either the closure of a vertex, or a tuple in a relation), so irreducible structure faithfulness of \mathbf{B}_1 gives the moreover part. By Observation 4.12.10 we know that $\text{support}(v_i) \subseteq \mathbf{S}$ for every $1 \leq i \leq m$. Hence $c_f(v_i)$ is defined for every $1 \leq i \leq m$ and an application of Claim 4.12.11 finishes the proof. \square

Claim 4.12.13. *Let $v \in B$ and let f be an automorphism of \mathbf{B}_1 such that $f(\text{support}(v)) \subseteq \mathbf{A}_1$. Then c_f restricts to a bijection between the out-neighbours of v in \mathbf{B} and the out-neighbours of $c_f(v)$ in \mathbf{A} .*

Proof. Pick an arbitrary $v' \in B$ such that $(v, v') \in E_{\mathbf{B}}$. By Corollary 4.12.12 we get that $(c_f(v), c_f(v')) \in E_{\mathbf{A}}$ and moreover if $v'' \in B$ is a different out-neighbour of v then by Claim 4.12.11 we know that $c_f(v') \neq c_f(v'')$. So c_f restricts to an injective function from the out-neighbours of v in \mathbf{B} to the out-neighbours of $c_f(v)$ in \mathbf{A} .

To prove that it is surjective, pick an arbitrary $y \in A$ such that $(c_f(v), y) \in E_{\mathbf{A}}$. We will find $y' \in \mathbf{B}$ such that $(v, y') \in E_{\mathbf{B}}$ and $c_f(y') = y$. If $v \in B_1$, we know that $c_f(v) = f(v)$ and hence $c_f(v)$ is a root of \mathbf{A} . If y is also a root of \mathbf{A} , it follows that $(v, f^{-1}(y)) \in E_{\mathbf{B}_1}$ and so $(v, f^{-1}(y)) \in E_{\mathbf{B}}$ by C1, hence we can put $y' = f^{-1}(y)$. If y is a non-root of \mathbf{A} then, by the construction of \mathbf{A}_1 , there is a tuple $\bar{y} = (y_1, \dots, y_n)$ of vertices of \mathbf{A}_1 such that $\bar{y} \in R_{\mathbf{A}_1}^{(y, \bar{y})}$ and $(f(v), y_1, \dots, y_n) \in E_{\mathbf{A}_1}^{(y, \bar{y})}$ (without loss of generality the enumeration of \bar{y} is chosen so that $((y, \bar{y}), \bar{y}) \in \mathcal{T}$). Let f' be an automorphism of \mathbf{B}_0 which induces f_L . Putting $y'' = (((f')^{-1}(y), f^{-1}(\bar{y})), f^{-1}(\bar{y})) \in \mathcal{T}_0$ and picking $y' \in \mathcal{T}$ such that $y' \sim y''$, we can see that indeed $c_f(y') = y$ and $(v, y') \in E_{\mathbf{B}}$ (by C2).

So suppose $v = ((x, (x_1, \dots, x_n)), (w_1, \dots, w_n)) \in \mathcal{T}$. Let f' be an automorphism of \mathbf{B}_0 which induces f_L . We know that $c_f(v) = f'(x)$ and that $f'(x)$ is a non-root of \mathbf{A} . If y is a root of \mathbf{A} , we get that $y \in \text{roots}_{\mathbf{A}}(f'(x))$, and hence, by the construction of \mathbf{A}_1 , there is $1 \leq i \leq n$ such that $f'(x_i) = y$. By Lemma 4.12.8 we know that $f(w_i) = f'(x_i) = y$, hence $w_i = f^{-1}(y) \in B_1$. If we put $y' = f^{-1}(y)$, we get that $c_f(y') = y$ and, by C3, $(v, y') \in E_{\mathbf{B}}$.

The last case is that y is a non-root of \mathbf{A} . Since $(f'(x), y) \in E_{\mathbf{A}}$, we get that

$$\text{roots}_{\mathbf{A}}(y) \subseteq \text{roots}_{\mathbf{A}}(f'(x)) = f'(\{w_1, \dots, w_n\}) = f'(\{x_1, \dots, x_n\}).$$

Let \bar{y} be an enumeration of $\text{roots}_{\mathbf{A}}(y)$. By the construction of \mathbf{A}_1 we get that $\bar{y} \in R_{\mathbf{A}_1}^{(y, \bar{y})}$ and so $f^{-1}(\bar{y}) \in R_{\mathbf{B}_1}^{(f'^{-1}(y), f'^{-1}(\bar{y}))}$. Put $y' = ((f'^{-1}(y), f'^{-1}(\bar{y})), f^{-1}(\bar{y}))$ and assume that enumeration of \bar{y} was chosen so that $y' \in \mathcal{T}$. Then $c_f(y') = y$ and, by C4, $(v, y') \in E_{\mathbf{B}}$, which concludes the proof. \square

Corollary 4.12.14. *\mathbf{B} is a k -orientation and the roots of \mathbf{B} are precisely members of B_1 .*

Proof. Pick an arbitrary $v \in \mathbf{B}$. Since $\text{support}(v)$ is an irreducible substructure of \mathbf{B}_1 , there is an automorphism f of \mathbf{B}_1 sending $\text{support}(v)$ to A_1 . Hence, by Claim 4.12.13, the out-degree of v in \mathbf{B} is the same as the out-degree of $c_f(v)$ in \mathbf{A} . Consequently, the out-degree of v in \mathbf{B} is at most k (hence \mathbf{B} is a k -orientation) and it is less than k if and only if $c_f(v)$ is a root of \mathbf{A} which happens if and only if $c_f(v) = f(v)$, i.e. if $v \in B_1$. \square

Corollary 4.12.15. *Given $u \in \mathbf{B}$, an automorphism $f: \mathbf{B}_1 \rightarrow \mathbf{B}_1$ sending $\text{support}(u)$ into A_1 and a sequence $c_f(u) = v_1, \dots, v_m$ of vertices of \mathbf{A} such that $(v_i, v_{i+1}) \in E_{\mathbf{A}}$ for every $1 \leq i < m$, there is a sequence $u = v'_1, \dots, v'_m$ of vertices of \mathbf{B} such that $(v'_i, v'_{i+1}) \in E_{\mathbf{B}}$ for every $1 \leq i < m$ and $c_f(v'_i) = v_i$ for every $1 \leq i \leq m$.*

Proof. We will prove this by induction on m . For $m = 1$ the statement is trivial. For the induction step, suppose that the statement is true for $m - 1$. By the induction hypothesis we have a sequence v'_1, \dots, v'_{m-1} of vertices of \mathbf{B} such that $(v'_i, v'_{i+1}) \in E_{\mathbf{B}}$ for every $1 \leq i < m - 1$ and $c_f(v'_i) = v_i$ for every $1 \leq i \leq m - 1$. By Observation 4.12.10 we know that $\text{support}(v'_i) \subseteq \text{support}(u)$ for every $1 \leq i \leq m - 1$, hence Claim 4.12.13 for v'_{m-1} tells us that c_f is a bijection between the out-neighbours of v'_{m-1} and $c_f(v'_{m-1}) = v_{m-1}$. Therefore in particular, there is some $v'_m \in B$ such that $c_f(v'_m) = v_m$ and $(v'_{m-1}, v'_m) \in E_{\mathbf{B}}$ which concludes the proof. \square

Claim 4.12.16. *For every $u \in \mathbf{B}$ it holds that $\text{roots}_{\mathbf{B}}(u) = \text{support}(u)$.*

Proof. First we prove that $\text{roots}_{\mathbf{B}}(u) \subseteq \text{support}(u)$. Pick an arbitrary $v \in \text{roots}_{\mathbf{B}}(u)$. This means that $v \in B_1$ (by Corollary 4.12.14) and that there is a sequence $u = v_1, \dots, v_m = v$ of vertices of \mathbf{B} such that $(v_i, v_{i+1}) \in E_{\mathbf{B}}$ for every $1 \leq i < m$. By Corollary 4.12.12 we get an automorphism f of \mathbf{B}_1 such that $c_f(v_i) \in A$ for every $1 \leq i \leq m$ and $(c_f(v_i), c_f(v_{i+1})) \in E_{\mathbf{A}}$ for every $1 \leq i < m$. Consequently, $c_f(v) = f(v) \in \text{roots}_{\mathbf{A}}(c_f(u))$ and so, by the construction of \mathbf{B}_1 and \mathbf{B} , $v \in \text{support}(u)$.

To see that $\text{roots}_{\mathbf{B}}(u) \supseteq \text{support}(u)$, pick an arbitrary $v \in \text{support}(u)$ and let f be an automorphism of \mathbf{B} sending $\text{support}(u)$ to A_1 . By the definition of \mathbf{B}_1 and $\text{support}(u)$ this means that $f(v) \in \text{roots}_{\mathbf{A}}(c_f(u))$, so in particular $c_f(v) = f(v) \in B_1$ and there is a sequence $c_f(u) = v_1, \dots, v_m = c_f(v)$ of vertices of \mathbf{A} such that $(v_i, v_{i+1}) \in E_{\mathbf{A}}$ for every $1 \leq i < m$. Using Corollary 4.12.15 we get a sequence v'_1, \dots, v'_m of vertices of \mathbf{B} such that $(v'_i, v'_{i+1}) \in E_{\mathbf{B}}$ for every $1 \leq i < m$ and $c_f(v'_i) = v_i$ for every $1 \leq i \leq m$. In particular, $v'_1 = u$ and $v'_m = v \in B_1$, hence $v \in \text{roots}_{\mathbf{B}}(u)$ (by Corollary 4.12.14). \square

Claim 4.12.17. *Let \mathbf{D}_1 be a substructure of \mathbf{B}_1 and f an automorphism of \mathbf{B}_1 such that $f(\mathbf{D}_1)$ is a substructure of \mathbf{A}_1 . Put*

$$D = \{v \in B : \text{support}(v) \subseteq D_1\}.$$

Then \mathbf{B} induces a substructure \mathbf{D} on D and c_f is an isomorphism from \mathbf{D} to a substructure induced by \mathbf{A} on $\text{Cl}_{\mathbf{A}}(f(D_1))$.

Proof. Since \mathbf{D}_1 is a substructure of \mathbf{B}_1 , it follows that $D_1 \subseteq D$. Let $u \in D$ and $v \in B$ be vertices such that $(u, v) \in E_{\mathbf{B}}$. From Observation 4.12.10 it follows that $\text{support}(v) \subseteq \text{support}(u) \subseteq D_1$ and so $v \in D$. This means that there are no outgoing edges from D in \mathbf{B} and thus D is closed on the function $F_{\mathbf{B}}$. By Claim 4.12.16, D is also closed on all functions $F_{\mathbf{B}}^n$, hence indeed \mathbf{B} induces a substructure \mathbf{D} on D .

Next we will prove that c_f is a bijection $D \rightarrow \text{Cl}_{\mathbf{A}}(f(D_1))$. Clearly, $\text{Dom}(c_f) \supseteq D$. Fix $v \in D$. If $v \in D_1$ then $c_f(v) = f(v) \in f(D_1)$. Conversely, all roots in $\text{Cl}_{\mathbf{A}}(f(D_1))$ are from $f(D_1)$, because $f(\mathbf{D}_1) \subseteq \mathbf{A}_1$.

So suppose $v = ((x, (x_1, \dots, x_n)), (w_1, \dots, w_n)) \in \mathcal{T}$ with $\{w_1, \dots, w_n\} \subseteq D_1$. Let f' be an automorphism of \mathbf{B}_0 which induces f_L . By Lemma 4.12.8 we get that $f(w_i) = f'(x_i)$ for every $1 \leq i \leq n$ and $\text{roots}_{\mathbf{A}}(f'(x)) = f'(\{x_1, \dots, x_n\})$. Since $c_f(v) = f'(x)$, it follows that $\text{roots}_{\mathbf{A}}(c_f(v)) = f(\{w_1, \dots, w_n\}) \subseteq f(D_1)$, hence indeed $c_f(v) \in \text{Cl}_{\mathbf{A}}(f(D_1))$. Conversely, let v be a non-root vertex of $\text{Cl}_{\mathbf{A}}(f(D_1))$ and let $\bar{y} = \text{roots}_{\mathbf{A}}(v)$ be an arbitrary enumeration. We know that $\bar{y} \subseteq f(D_1)$ and by the construction we also get that $\bar{y} \in R_{\mathbf{A}_1}^{(v, \bar{y})}$. Hence we can reconstruct $v'' = (((f')^{-1}(v), f^{-1}(\bar{y})), f^{-1}(\bar{y})) \in \mathcal{T}_0$ and $v' \in \mathcal{T}$ with $v' \sim v''$ such that $c_f(v') = v$. So indeed $\text{Range}(c_f) = \text{Cl}_{\mathbf{A}}(f(D_1))$. From P1 of Claim 4.12.11 we get that c_f is a bijection $D \rightarrow \text{Cl}_{\mathbf{A}}(f(D_1))$.

Finally, from P2 of Claim 4.12.11 it follows that for every $u, v \in D$ we have $(u, v) \in E_{\mathbf{D}}$ if and only if $(c_f(u), c_f(v)) \in E_{\mathbf{A}}$. Since all functions in \mathbf{B} and \mathbf{A} are defined from the graph structure, it follows that c_f indeed is an isomorphism $\mathbf{D} \rightarrow \text{Cl}_{\mathbf{A}}(f(D_1))$. \square

Lemma 4.12.18. *Let $\mathbf{D} \subseteq \mathbf{B}$ be irreducible. Then \mathbf{B}_1 induces an irreducible substructure on $D \cap B_1$.*

Proof. First note that if \mathbf{D} is a substructure of \mathbf{B} , then \mathbf{B}_1 induces a substructure on $D \cap B_1$. Indeed, by Corollary 4.12.14 we know that $D \cap B_1$ are precisely the root vertices of \mathbf{D} . If there is $v \in \text{Cl}_{\mathbf{B}_1}(D \cap B_1) \setminus D$ then by Corollary 4.12.14 it is a root vertex of \mathbf{B} . This means that there is $u \in D \cap B_1$ such that $v \in \text{Cl}_{\mathbf{B}_1}(u) = \text{support}(u) = \text{roots}_{\mathbf{B}}(u)$ (by Claim 4.12.16). But this implies that $v \in \text{Cl}_{\mathbf{B}}(u)$, hence $v \in \mathbf{D}$, a contradiction.

Put $D_1 = D \cap B_1$ and let \mathbf{D}_1 be the substructure induced by \mathbf{B}_1 on D_1 . We know that D_1 are precisely the root vertices of \mathbf{D} . From the definition of \mathbf{B} and Claim 4.12.16 it follows that

$$D = \{v \in B : \text{support}(v) \subseteq D_1\} = \{v \in B : \text{roots}_{\mathbf{B}}(v) \subseteq D_1\}.$$

We will now prove that if \mathbf{D}_1 is reducible then \mathbf{D} is also reducible. Taking the contrapositive then proves the statement of this claim. Suppose that there are substructures $\mathbf{D}_1^b, \mathbf{D}_1^l, \mathbf{D}_1^r \subseteq \mathbf{D}_1$ such that \mathbf{D}_1 is the free amalgamation of \mathbf{D}_1^l and \mathbf{D}_1^r over \mathbf{D}_1^b (in particular, $\mathbf{D}_1^b \subseteq \mathbf{D}_1^l$ and $\mathbf{D}_1^b \subseteq \mathbf{D}_1^r$). Put

$$\begin{aligned} D^l &= \{v \in B : \text{support}(v) \subseteq D_1^l\}, \\ D^r &= \{v \in B : \text{support}(v) \subseteq D_1^r\}, \\ D^b &= \{v \in B : \text{support}(v) \subseteq D_1^b\} \end{aligned}$$

and let $\mathbf{D}^l, \mathbf{D}^r$ and \mathbf{D}^b be the substructures of \mathbf{B} induced on D^l, D^r and D^b respectively. Clearly, $\mathbf{D}^l, \mathbf{D}^r, \mathbf{D}^b \subseteq \mathbf{D}$, $\mathbf{D}^b \subseteq \mathbf{D}^l$ and $\mathbf{D}^b \subseteq \mathbf{D}^r$. We will prove that \mathbf{D} is the free amalgamation of \mathbf{D}^l and \mathbf{D}^r over \mathbf{D}^b .

Since both \mathbf{D}^l and \mathbf{D}^r are substructures of \mathbf{B} , they are in particular closed on functions F and F^n . If there were vertices $u \in D^l, v \in D^r$ such that $(u, v) \in E_{\mathbf{D}}$, then $v \in F_{\mathbf{D}}(u)$, which is a contradiction. Hence there are no edges spanning vertices of both \mathbf{D}^l and \mathbf{D}^r at the same time.

If $v \in D^b$ then, as \mathbf{D}^b is a substructure, it follows that $\text{roots}_{\mathbf{B}}(v) \subseteq D^b$, similarly for \mathbf{D}^l and \mathbf{D}^r . It follows that whenever \bar{x} contains vertices from both $D^l \setminus D^b$ and $D^r \setminus D^b$ then $F_{\mathbf{D}}^{|\bar{x}|}(\bar{x}) = \emptyset$. Consequently, \mathbf{D} is the free amalgamation of \mathbf{D}^l and \mathbf{D}^r over \mathbf{D}^b . \square

Corollary 4.12.19. $\mathbf{B} \in \mathcal{D}_d^k$.

Proof. Let \mathbf{D} be an irreducible substructure of \mathbf{B} . Since \mathcal{D}_d^k is a free amalgamation class, it suffices to prove that $\mathbf{D} \in \mathcal{D}_d^k$. By Lemma 4.12.18 we know that \mathbf{B}_1 induces an irreducible substructure on $D \cap B_1$ and by Claim 4.12.16 this substructure is non-empty. Now we can use irreducible structure faithfulness of \mathbf{B}_1 and Claim 4.12.17 to get an embedding $c_f: \mathbf{D} \rightarrow \mathbf{A}$. As $\mathbf{A} \in \mathcal{D}_d^k$, we get that $\mathbf{D} \in \mathcal{D}_d^k$, hence indeed $\mathbf{B} \in \mathcal{D}_d^k$. \square

Lemma 4.12.20. *Let $f = (f_L, f_{B_1})$ be an automorphism of \mathbf{B}_1 and let f' be an automorphism of \mathbf{B}_0 which induces f_L . Let ι_0 be the map on \mathcal{T}_0 given by $\iota_0((x, \bar{x}), \bar{w}) = ((f'(x), f'(\bar{x})), f_{B_1}(\bar{w}))$ and let ι be the map induced by ι_0 on \mathcal{T} (it is easy to see that \sim is a congruence with respect to ι_0). Define $\theta = f_{B_1} \cup \iota$. Then θ is an automorphism of \mathbf{B} .*

Proof. Since f' is an automorphism of \mathbf{B}_0 , it follows that ι_0 is a bijection $\mathcal{T}_0 \rightarrow \mathcal{T}_0$. Consequently, ι is a bijection $\mathcal{T} \rightarrow \mathcal{T}$. As f_{B_1} is a bijection $B_1 \rightarrow B_1$ and B is the disjoint union of B_1 and \mathcal{T} , it follows that θ is a bijection $B \rightarrow B$. By Corollary 4.12.19 we know that the functions $F_{\mathbf{B}}$ and $F_{\mathbf{B}}^n$ are defined in \mathbf{B} from the graph structure. To see that θ is an automorphism of \mathbf{B} , it remains to prove that for every $u, v \in B$ we have $(u, v) \in E_{\mathbf{B}}$ if and only if $(\theta(u), \theta(v)) \in E_{\mathbf{B}}$. We will distinguish four cases.

1. First suppose that $u, v \in B_1$. By C1 we know that $(u, v) \in E_{\mathbf{B}}$ if and only if $(u, v) \in E_{\mathbf{B}_1}$. Since f is an automorphism of \mathbf{B}_1 , we know that $(u, v) \in E_{\mathbf{B}_1}$ if and only if $(\theta(u), \theta(v)) = (f_{B_1}(u), f_{B_1}(v)) \in E_{\mathbf{B}_1}$, so indeed $(u, v) \in E_{\mathbf{B}}$ if and only if $(\theta(u), \theta(v)) \in E_{\mathbf{B}}$.
2. If $u \in B_1$ and $v = ((x, \bar{x}), (w_1, \dots, w_n)) \in \mathcal{T}$, we know by C2 that $(u, v) \in E_{\mathbf{B}}$ if and only if $(u, w_1, \dots, w_n) \in E_{\mathbf{B}_1}^{(x, \bar{x})}$. Again, since f is an automorphism of \mathbf{B}_1 , we get that

$$(u, w_1, \dots, w_n) \in E_{\mathbf{B}_1}^{(x, \bar{x})}$$

if and only if

$$f_{B_1}((u, w_1, \dots, w_n)) \in E_{\mathbf{B}_1}^{(f'(x), f'(\bar{x}))}.$$

As $\theta(u) = f_{B_1}(u)$ and $\theta(v) = ((f'(x), f'(\bar{x})), f_{B_1}((w_1, \dots, w_n)))$, we get from C2 that this happens if and only if $(\theta(u), \theta(v)) \in E_{\mathbf{B}}$.

3. If $u = ((x, (x_1, \dots, x_n)), (w_1, \dots, w_n)) \in \mathcal{T}$ and $v \in B_1$, we know by C3 that $(u, v) \in E_{\mathbf{B}}$ if and only if there is $1 \leq i \leq n$ such that $v = w_i$ and $(x, x_i) \in E_{\mathbf{B}_0}$. This is equivalent to $f_{B_1}(v) = f_{B_1}(w_i)$ and $(f'(x), f'(x_i)) \in E_{\mathbf{B}_0}$ which is in turn once again equivalent to $(\theta(u), \theta(v)) \in E_{\mathbf{B}}$.
4. Finally, if $u = ((x, \bar{x}), \bar{w}) \in \mathcal{T}$, $v = ((y, \bar{y}), \bar{t}) \in \mathcal{T}$ then (by C4) $(u, v) \in E_{\mathbf{B}}$ if and only if $\bar{t} \subseteq \bar{w}$ (as sets) and $(x, y) \in E_{\mathbf{B}_0}$. This is equivalent to $f_{B_1}(\bar{t}) \subseteq f_{B_1}(\bar{w})$ (as sets) and $(f'(x), f'(y)) \in E_{\mathbf{B}_0}$, or in other words, $(\theta(u), \theta(v)) \in E_{\mathbf{B}}$.

\square

Note that $\text{Cl}_{\mathbf{A}}(A_1) = A$. Therefore, Claim 4.12.17 for $\mathbf{D}_1 = \mathbf{A}_1$ and the identity automorphism gives us an isomorphism c_{id} . Put $\psi = c_{\text{id}}^{-1}: \mathbf{A} \rightarrow \mathbf{B}$ and denote $\mathbf{A}' = \psi(\mathbf{A})$. This will be the copy of \mathbf{A} in \mathbf{B} whose automorphisms we are going to extend. Note that for every $v \in A_1$ we have that $\psi(v) = v$ and for every $v \in A \setminus A_1$ we have that $\pi(\psi(v)) = v$. It follows that $A_1 \subseteq A'$ and

$$A' = \{v \in B : \text{support}(v) \subseteq A_1\}.$$

Proposition 4.12.21. *\mathbf{B} is an irreducible structure faithful coherent EPPA-witness for \mathbf{A}' .*

Proof. First we refine the proof of Corollary 4.12.19 to prove that \mathbf{B} is irreducible structure faithful. Let \mathbf{D} be an irreducible substructure of \mathbf{B} . By Lemma 4.12.18 we know that \mathbf{B}_1 induces an irreducible substructure \mathbf{D}_1 on $D \cap B_1$. Now we can use irreducible structure faithfulness of \mathbf{B}_1 to get an automorphism $f = (f_L, f_{B_1})$ of \mathbf{B}_1 such that $f(D_1) \subseteq A_1$. Let f' be an automorphism of \mathbf{B}_0 inducing f_L .

Use Lemma 4.12.20 to construct an automorphism θ of \mathbf{B} . Clearly, $\theta(D_1) \subseteq A_1 \subseteq A'$. Hence it remains to prove that $\theta(D \setminus D_1) \subseteq A'$ (where $D \setminus D_1$ are precisely the non-root vertices of \mathbf{D}). Pick an arbitrary $v \in D \setminus D_1$. We know that $\text{support}(v) \subseteq D_1$. By definition of θ we also know that $\text{support}(\theta(v)) = f_{B_1}(\text{support}(v)) \subseteq A_1$, and so $\theta(v) \in A'$ and thus indeed $\theta(D) \subseteq A'$.

To see that \mathbf{B} is an EPPA-witness for \mathbf{A}' , let φ be a partial automorphism of \mathbf{A}' . Consider $\varphi_0 = \psi^{-1} \circ \varphi \circ \psi$ as a partial automorphism of \mathbf{A}_0 (note that $\varphi_0(v) = \varphi(v)$ for every $v \in A_1$) and extend it to an automorphism $\tilde{\varphi}_0$ of \mathbf{B}_0 . Let $\phi_L \in \Gamma_{L^+}$ be the permutation of L^+ given by $\tilde{\varphi}_0$ and put $\phi_{A_1} = \varphi_0 \upharpoonright_{A_1} = \varphi \upharpoonright_{A_1}$. Note that ϕ_{A_1} is a bijection $A_1 \rightarrow A_1$, because φ_0 preserves whether a vertex is a root or not. Put $\phi = (\phi_L, \phi_{A_1})$. It is easy to verify that ϕ is a partial automorphism of \mathbf{A}_1 : It preserves the L_s relations and functions, because φ_0 does. Suppose that $\bar{w} \subseteq \text{Dom}(\phi)$ and $\bar{w} \in R_{\mathbf{A}_1}^{(x, \bar{w})}$, then this means that, in \mathbf{A} , there is a vertex x such that $\text{roots}_{\mathbf{A}}(x) = \bar{w}$. This implies that $x \in F_{\mathbf{A}}^n(\bar{w})$ and hence $x \in \text{Dom}(\varphi)$. Consequently, $\phi_L(R^{(x, \bar{w})}) = R^{(\varphi(x), \varphi(\bar{w}))}$ and so indeed $\phi_{A_1}(\bar{w}) \in \phi_L(R^{(x, \bar{w})})_{\mathbf{A}_1}$.

Let $\tilde{\phi}: \mathbf{B}_1 \rightarrow \mathbf{B}_1$ be the extension of ϕ and use Lemma 4.12.20 for $\tilde{\varphi}_0$ (as f') and $\tilde{\phi}$ (as f) to get an automorphism $\tilde{\varphi}$ of \mathbf{B} (called θ in the statement of Lemma 4.12.20). By the construction, we know that $\varphi \upharpoonright_{A_1} \subseteq \tilde{\varphi}$. In order to prove that $\tilde{\varphi}$ extends φ it thus remains to argue that $\tilde{\varphi}(v) = \varphi(v)$ for every $v \in \text{Dom}(\varphi) \cap \mathcal{T}$.

Pick an arbitrary $v = ((x, \bar{x}), \bar{w}) \in \text{Dom}(\varphi) \cap \mathcal{T}$. Since $v \in \text{Dom}(\varphi)$, we know that $v \in A'$, hence $\bar{w} = \text{support}(v) \subseteq A_1$, and consequently $\bar{x} = \bar{w}$ (as tuples). By the construction we know that (up to applying the same permutation on \bar{x} and \bar{w} to pick the correct member of the \sim -equivalence class),

$$\tilde{\varphi}(v) = ((\tilde{\varphi}_0(x), \tilde{\varphi}_0(\bar{x})), \tilde{\phi}(\bar{w})),$$

which is equal to

$$((\varphi_0(x), \varphi_0(\bar{x})), \varphi(\bar{w})) = ((\varphi_0(x), \varphi(\bar{w})), \varphi(\bar{w})).$$

In particular, $\pi(\tilde{\varphi}(v)) = \varphi_0(x)$. By the construction we know that $v = \psi(x)$ and consequently $\pi(v) = x$. By the definition of φ_0 we know that $\psi \circ \varphi_0 = \varphi \circ \psi$, so in particular $\varphi_0(x) = \pi(\psi(\varphi_0(x))) = \pi(\varphi(\psi(x)))$. So indeed, $\pi(\tilde{\varphi}(v)) = \pi(\varphi(v))$.

We know that $\text{roots}_{\mathbf{B}}(v) = \bar{w}$ (as sets), hence $\bar{w} \subseteq \text{Dom}(\varphi)$. Since φ is a partial automorphism, $\varphi(\text{roots}_{\mathbf{B}}(v)) = \text{roots}_{\mathbf{B}}(\varphi(v))$. But $\text{roots}_{\mathbf{B}}(\varphi(v)) = \text{support}(\varphi(v))$ (by Claim 4.12.16) and we know that $\text{support}(\varphi(v)) = \varphi(\bar{w})$. It follows that $\tilde{\varphi}(v) = \varphi(v)$ and hence $\tilde{\varphi}$ indeed extends φ , which concludes the argument.

Since both $\tilde{\varphi}_0$ and $\tilde{\phi}$ can be chosen coherently, it follows that $\tilde{\varphi}$ is also coherent and hence \mathbf{B} is a coherent EPPA-witness for \mathbf{A}' . \square

Proof of Theorem 4.12.7. In this section we constructed, for an arbitrary $\mathbf{A} \in \mathcal{D}_d^k$ a structure \mathbf{B} . By Corollary 4.12.19, $\mathbf{B} \in \mathcal{D}_d^k$ and by Proposition 4.12.21 \mathbf{B} is an irreducible structure faithful coherent EPPA-witness for an isomorphic copy of \mathbf{A} . Hence indeed \mathcal{D}_d^k has irreducible structure faithful coherent EPPA. \square

4.13 Conclusion

Comparing known EPPA classes and known Ramsey classes one can easily identify two main weaknesses of the state-of-the-art EPPA constructions.

1. The need for automorphism-preserving completion procedure is not necessary in the Ramsey context. The example of two-graphs [EHKN20] shows that there are classes with EPPA which do not admit automorphism-preserving completions (see [Kon20] for a more systematic treatment of certain classes of this kind). Understanding the situation better might lead to solutions of some of the long standing open problems in this area including the question whether the class of all finite tournaments has EPPA (see [HPSW19], [HJKS19] and [HJKS23] for recent progress on this problem).
2. There is a lack of general EPPA constructions for classes with non-unary function symbols. Again, there are known classes with non-unary function symbols that have EPPA (e.g. finite groups or classes from Section 4.12.4). It is however not known whether, for example, the class of all finite partial Steiner systems or the class of all finite equivalences on unordered pairs with two equivalence classes have EPPA.

On the other hand, in this paper we consider Γ_L -structures which, in the finite language case, reduce to the usual model-theoretic structures in the Ramsey context, because the action of Γ_L must be trivial in order for the class to be rigid. This has some additional applications including:

1. Elimination of imaginaries for classes having definable equivalence classes (see [Iva15, HN19]),
2. representation of special non-unary functions which map vertices of one type to vertices of different type (see Section 4.12.4 or Theorem 4.12.7), or
3. representation of antipodal structures and switching classes ([EHKN20, Kon20]).

We refer the reader to [HN19, ABWH⁺17b, Kon18, Kon19, HKN18] for various examples of (automorphism-preserving) locally finite subclasses.

One of the main weaknesses of Theorems 4.1.1, 4.1.5 and 4.1.6 is that they only allow unary functions. It would be interesting to know whether they hold without this restriction.

Question 4.13.1. *Do Theorems 4.1.1, 4.1.5 and 4.1.6 hold for languages with non-unary functions?*

A positive answer to Question 4.13.1 would have some applications which are interesting on their own and have been asked before. We present two of them as separate questions.

Question 4.13.2. *Let L be the language consisting of a single binary function and let \mathcal{C} be the class of all finite L -structures (say, such that the image of every pair of vertices has cardinality at most one). Does \mathcal{C} have EPPA?*

Question 4.13.3. *Does the class of all finite partial Steiner triple systems have EPPA, where one only wants to extend partial automorphism between closed substructures? (A sub-hypergraph H of a Steiner triple system S is closed if whenever $\{x, y, z\}$ is a triple of S and $x, y \in H$, then $z \in H$.)*

Acknowledgment

We would like to thank Andy Zucker, the anonymous referee, and David M. Evans for their helpful comments which significantly improved the quality of this paper.

5. EPPA for two-graphs and antipodal metric spaces

DAVID M. EVANS, JAN HUBIČKA, MATĚJ KONEČNÝ, JAROSLAV NEŠETŘIL

Abstract We prove that the class of finite two-graphs has the extension property for partial automorphisms (EPPA, or Hrushovski property), thereby answering a question of Macpherson. In other words, we show that the class of graphs has the extension property for switching automorphisms. We present a short, self-contained, purely combinatorial proof which also proves EPPA for the class of integer valued antipodal metric spaces of diameter 3, answering a question of Aranda et al.

The class of two-graphs is an important new example which behaves differently from all the other known classes with EPPA: Two-graphs do not have the amalgamation property with automorphisms (APA), their Ramsey expansion has to add a graph, it is not known if they have coherent EPPA and even EPPA itself cannot be proved using the Herwig–Lascar theorem.

5.1 Introduction

Two-graphs, introduced by G. Higman and studied extensively since the 1970s (see for example [Cam99, Sei73]), are 3-uniform hypergraphs with the property that on every four vertices there is an even number of hyperedges. A class \mathcal{C} of finite structures (such as hypergraphs) has the *extension property for partial automorphisms* (EPPA, sometimes also called *Hrushovski property*) if for every $\mathbf{A} \in \mathcal{C}$ there exists $\mathbf{B} \in \mathcal{C}$ containing \mathbf{A} as an (induced) substructure such that every isomorphism between substructures of \mathbf{A} extends to an automorphism of \mathbf{B} . We call \mathbf{B} an *EPPA-witness* for \mathbf{A} . We prove:

Theorem 5.1.1. *The class \mathcal{T} of all finite two-graphs has EPPA.*

Our result answers a question of Macpherson which is also stated in Siniora’s PhD thesis [Sin17] and can be seen as a contribution to the ongoing effort of identifying new classes of structures with EPPA. This was started in 1992 by Hrushovski’s proof [Hru92] that the class of all finite graphs has EPPA, and followed by a series of papers dealing with other classes, including [ABWH⁺17b, Con19, Her95, Her98, HL00, HO03, HKN19a, HKN18, Kon19, Ott20, Sol05, Ver08].

All proofs of EPPA in this paper are purely combinatorial and self-contained. The second part of the paper requires some model-theoretical notions and discusses in more detail the interplay of the following properties for which there were no known examples before:

1. The usual procedure for building an EPPA-witness is to construct an incomplete object (where some relations are missing) and later complete it to satisfy axioms of the class without affecting any automorphisms (i.e. one needs to have an *automorphism-preserving completion* [ABWH⁺17b]). This

is not possible for two-graphs and thus makes them related to tournaments which pose a well known open problem in the area, see Remark 5.8.1.

2. The class \mathcal{T} does not have APA (*amalgamation property with automorphisms*). Hodges, Hodkinson, Lascar, and Shelah [H HLS93] introduced this notion and showed that APA together with EPPA imply the existence of *ample generics* (see also [Sin17, Chapter 2]). To the authors' best knowledge, \mathcal{T} is the only known class with EPPA but not APA besides pathological examples, see Section 5.6.
3. In all cases known to the authors except for the class of all finite groups, whenever a class of structures \mathcal{C} has EPPA then expanding a variant of \mathcal{C} by linear orders gives a Ramsey expansion. This does not seem to be the case for two-graphs, see Section 5.7.
4. Solecki and Sinióra [SS19, Sol09] introduced the notion of *coherent EPPA* (see Section 5.2.1) as a way to prove that the automorphism group of the respective Fraïssé limit contains a dense locally finite subgroup. Our method does not give coherent EPPA for \mathcal{T} and thus it makes \mathcal{T} the only known example with EPPA for which coherent EPPA is not known. However, our method does give coherent EPPA for the class of all antipodal metric spaces of diameter 3 and using it we are able to obtain a dense locally finite subgroup of the Fraïssé limit of \mathcal{T} , see Section 5.5.

Two-graphs are closely related to the switching classes of graphs and to double covers of complete graphs [Cam99, Sei73], which is in fact key in this paper. Our results can thus be interpreted as a direct strengthening of the theorem of Hrushovski [Hru92] which states that the class of all finite graphs \mathcal{G} has EPPA. Namely we can consider \mathcal{G} with a richer class of mappings — the switching automorphisms.

Given a graph \mathbf{G} with vertex set G and $S \subseteq G$, the (*Seidel*) *switching* \mathbf{G}_S of \mathbf{G} is the graph created from \mathbf{G} by complementing the edges between S and $G \setminus S$. (That is, for $s \in S$ and $t \in G \setminus S$ it holds that $\{s, t\}$ is an edge of \mathbf{G}_S if and only if $\{s, t\}$ is not an edge of \mathbf{G} . Edges and non-edges with both endpoints in S or $G \setminus S$ are preserved.)

Given a graph \mathbf{H} with vertex set H , a function $f: G \rightarrow H$ is a *switching isomorphism* of \mathbf{G} and \mathbf{H} if there exists $S \subseteq G$ such that f is an isomorphism of \mathbf{G}_S and \mathbf{H} . If $\mathbf{G} = \mathbf{H}$ we call such a function a *switching automorphism*.

Definition 5.1.1. We say that a class $\mathcal{C} \subseteq \mathcal{G}$ has the *extension property for switching automorphisms* if for every $\mathbf{G} \in \mathcal{C}$ there exists $\mathbf{H} \in \mathcal{C}$ containing \mathbf{G} as an induced subgraph such that every switching isomorphism of induced subgraphs of \mathbf{G} extends to a switching automorphism of \mathbf{H} and, moreover, every isomorphism of induced subgraphs of \mathbf{G} extends to an automorphism of \mathbf{H} .

In this language we prove:

Theorem 5.1.2. *The class of all finite graphs \mathcal{G} has the extension property for switching automorphisms.*

Because of the ‘moreover’ part of Definition 5.1.1, Theorem 5.1.2 implies the theorem of Hrushovski. It is also a strengthening of Theorem 5.1.1 by the following well-known correspondence between two-graphs and switching classes [Cam99, Sei73].

Given a graph \mathbf{G} , its *associated two-graph* $T(\mathbf{G})$ is a two-graph on the same vertex set as \mathbf{G} such that $\{a, b, c\}$ is a hyperedge if and only if the three-vertex subgraph induced by \mathbf{G} on $\{a, b, c\}$ has an odd number of edges. Then a function $f: G \rightarrow H$ between graphs \mathbf{G} and \mathbf{H} is a switching isomorphism if and only if it is an isomorphism between the associated two-graphs. Thus, the existence of a switching isomorphism is an equivalence relation on the class of graphs and two-graphs correspond to the equivalence classes (called *switching classes of graphs*).

We shall see that the most natural setting for our proof is to work with the class of all finite integer-valued antipodal metric spaces of diameter 3. Following [ACM21] we call a metric space an *integer-valued metric space of diameter 3* if the distance of every two distinct points is 1, 2 or 3. It is *antipodal* if

1. it contains no triangle with distances 2, 2, 3, and
2. the edges with label 3 form a perfect matching (in other words, for every vertex there is precisely one *antipodal vertex* at distance 3).

We require that the domain and image of a partial automorphism of such a structure should be closed under taking antipodal points. However, this can be assumed without loss of generality because there always is a unique way of extending a partial automorphism not satisfying this condition to one which does. In this language, we can then state our main theorem as:

Theorem 5.1.3. *The class of all finite integer-valued antipodal metric spaces of diameter 3 has coherent EPPA.*

This theorem also holds for all other antipodal metric spaces from Cherlin’s catalogue of metrically homogeneous graphs [Che22], see [Kon20]. It answers affirmatively a question of Aranda, Bradley-Williams, Hubička, Karamanlis, Kompatscher, Konečný and Pawliuk [ABWH⁺17b] and completes the analysis of EPPA for all classes from Cherlin’s catalogue. However, in this note, for brevity, we often refer to an antipodal, integer-valued metric space of diameter 3 as an *antipodal metric space*. Other antipodal metric spaces are not considered.

5.2 Notation and preliminaries

It is in the nature of this paper to consider multiple types of structures. We will use bold letter such as $\mathbf{A}, \mathbf{B}, \mathbf{C}, \dots$ to denote structures ((hyper)graphs or metric spaces defined below) and corresponding normal letters, such as A, B, C, \dots , to denote corresponding vertex sets. Our substructures (sub-(hyper)graphs or subspaces) will always be induced.

Formally, we will consider a metric space to be a complete edge-labelled graph (that is, a complete graph where edges are labelled by the respective distances), or, equivalently, a relational structure with multiple binary relations representing the distances. This justifies that we will speak of pairs of vertices at distance d

as of *edges of length d* . We will, however, use both notions (a vertex set with a distance function or a complete edge-labelled graph) interchangeably. We adopt the standard notion of isomorphism, embedding and substructure.

5.2.1 Coherent EPPA

Suppose \mathcal{C} is a class of finite structures which has the hereditary and joint embedding properties. If \mathcal{C} has EPPA, then it has the amalgamation property, so, assuming that there are only countably many isomorphism types of structures in \mathcal{C} , then we can consider the Fraïssé limit \mathbf{M} of \mathcal{C} . Coherence is a natural strengthening of EPPA which guarantees that $\text{Aut}(\mathbf{M})$ has a dense locally finite subgroup [Sol09, SS19]. Note that the existence of a dense locally finite subgroup of the automorphism group of a homogeneous structure implies that its age has EPPA, but it is not known whether coherent EPPA also follows from this: see Section 5.1 of [SS19]. At the moment all previously known EPPA classes are also coherent EPPA classes. Interestingly, we can prove coherent EPPA for the antipodal metric spaces of diameter 3, but not for two-graphs (this is discussed in Section 5.5). We need to introduce two additional definitions.

Definition 5.2.1 (Coherent maps [Sol09, SS19]). Let X be a set and \mathcal{P} be a family of partial bijections between subsets of X . A triple (f, g, h) from \mathcal{P} is called a *coherent triple* if

$$\text{Dom}(f) = \text{Dom}(h), \text{Range}(f) = \text{Dom}(g), \text{Range}(g) = \text{Range}(h)$$

and

$$h = g \circ f.$$

Let X and Y be sets, and \mathcal{P} and \mathcal{Q} be families of partial bijections between subsets of X and between subsets of Y , respectively. A function $\varphi: \mathcal{P} \rightarrow \mathcal{Q}$ is said to be a *coherent map* if for each coherent triple (f, g, h) from \mathcal{P} , its image $(\varphi(f), \varphi(g), \varphi(h))$ in \mathcal{Q} is coherent. In the case of EPPA, where the elements of \mathcal{Q} are automorphisms of Y , we sometimes refer to the image of φ as a coherent family of automorphisms extending \mathcal{P} .

Definition 5.2.2 (Coherent EPPA [Sol09, SS19]). A class \mathcal{C} of finite structures is said to have *coherent EPPA* if \mathcal{C} has EPPA and moreover the extension of partial automorphisms is coherent. More precisely, for every $\mathbf{A} \in \mathcal{C}$, there exists an EPPA-witness $\mathbf{B} \in \mathcal{C}$ for \mathbf{A} and a coherent map $f \mapsto \hat{f}$ from the family of partial automorphisms of \mathbf{A} to the group of automorphisms of \mathbf{B} . In this case we also call \mathbf{B} a *coherent EPPA-witness* of \mathbf{A} .

5.3 EPPA for antipodal metric spaces

Given an antipodal metric space \mathbf{A} we give a direct construction of a coherent EPPA-witness \mathbf{B} . Some ideas are based on a construction of Hodkinson and Otto [HO03] and some of the terminology is loosely based on Hodkinson's exposition of this construction [Hod02]. Note that our techniques also give a very simple and short proof of EPPA for graphs.

Fix a (finite) antipodal metric space \mathbf{A} . Denote by $M = \{e_1, e_2, \dots, e_n\}$ the set of all edges of \mathbf{A} of length 3. For a function $\chi: M \rightarrow \{0, 1\}$ we denote by $1 - \chi$ the function satisfying $(1 - \chi)(e) = 1 - \chi(e)$ for every $e \in M$. We construct \mathbf{B} as follows:

1. The vertices of \mathbf{B} are all pairs (e, χ) where $e \in M$ and χ is a function from M to $\{0, 1\}$ (called a *valuation function*).
2. Distances for $(e, \chi) \neq (f, \chi') \in B$ are given by the following rules:
 - (i) $d_{\mathbf{B}}((e, \chi), (e, 1 - \chi)) = 3$,
 - (ii) $d_{\mathbf{B}}((e, \chi), (f, \chi')) = 1$ if and only if $\chi(f) = \chi'(e)$,
 - (iii) $d_{\mathbf{B}}((e, \chi), (f, \chi')) = 2$ otherwise.

Lemma 5.3.1. *The structure \mathbf{B} is an antipodal metric space.*

Proof. Given $(e, \chi) \in B$, by (i) we know that there is precisely one vertex at distance 3 (namely $(e, 1 - \chi)$) and thus the edges of length 3 form a perfect matching.

It remains to check that every quadruple $(e, \chi), (e, 1 - \chi), (f, \chi'), (f, 1 - \chi')$ of distinct vertices of \mathbf{B} is an antipodal metric space. By (ii) we know that precisely one of $(f, \chi'), (f, 1 - \chi')$ is at distance 1 from (e, χ) and by (iii) that the other is at distance 2, similarly for $(e, 1 - \chi)$. It also follows that $d_{\mathbf{B}}((e, \chi), (f, \chi')) = d_{\mathbf{B}}((e, 1 - \chi), (f, 1 - \chi'))$ and $d_{\mathbf{B}}((e, \chi), (f, 1 - \chi')) = d_{\mathbf{B}}((e, 1 - \chi), (f, \chi'))$. \square

We now define an embedding $\psi: \mathbf{A} \rightarrow \mathbf{B}$ and refer to its image \mathbf{A}' in \mathbf{B} as a *generic copy* of \mathbf{A} in \mathbf{B} .

Fix an arbitrary function $p: A \rightarrow \{0, 1\}$ such that whenever $d_{\mathbf{A}}(x, x') = 3$, then $p(x) = 1 - p(x')$. This function partitions the vertices of \mathbf{A} into two *podes* such that pairs of vertices at distance 3 are in different podes. For every $1 \leq i \leq n$ we denote by x_i and y_i the endpoints of e_i such that $p(x_i) = 0$ and $p(y_i) = 1$. We construct ψ by putting $\psi(x_i) = (e_i, \chi_i)$ and $\psi(y_i) = (e_i, 1 - \chi_i)$, where χ_i is defined as

$$\chi_i(e_j) = \begin{cases} 0 & \text{if } j \geq i \\ 0 & \text{if } j < i \text{ and } d_{\mathbf{A}}(x_i, x_j) = 1 \\ 1 & \text{otherwise.} \end{cases}$$

It follows from the construction that ψ is indeed an embedding $\mathbf{A} \rightarrow \mathbf{B}$. We put $\mathbf{A}' = \psi(\mathbf{A})$. Now we are ready to show the main result of this section:

Proposition 5.3.2. *With the above notation, the antipodal metric space \mathbf{B} is a coherent EPPA-witness for \mathbf{A}' . Moreover, $p \circ \psi^{-1}$ extends to a function $\hat{p}: B \rightarrow \{0, 1\}$ such that whenever a partial automorphism φ of \mathbf{A}' preserves values of $p \circ \psi^{-1}$, then its coherent extension θ preserves values of \hat{p} .*

Proof. Let φ be a partial automorphism of \mathbf{A}' . Let $\pi: B \rightarrow M$ be the projection mapping $(e, \chi) \mapsto e$. By this projection, φ induces a partial permutation of M and we denote by $\hat{\varphi}$ an extension of it to a permutation of M . To obtain coherence we always extend the permutation in an *order-preserving way*, that is, we enumerate $M \setminus \text{Dom}(\varphi) = \{e_{i_1}, \dots, e_{i_k}\}$ and $M \setminus \text{Range}(\varphi) = \{e_{j_1}, \dots, e_{j_k}\}$,

where $i_1 < \dots < i_k$ and $j_1 < \dots < j_k$, and put $\hat{\varphi}(e_{i_\ell}) = e_{j_\ell}$ for every $1 \leq \ell \leq k$ (cf. [SS19, Sol09]).

Let F be the set consisting of unordered pairs $\{e, f\}$, $e, f \in M$ (possibly $e = f$), such that there exists χ with the property that $(e, \chi) \in \text{Dom}(\varphi)$ and $\chi(f) \neq \chi'(\hat{\varphi}(f))$ for $(\hat{\varphi}(e), \chi') = \varphi((e, \chi))$. We say that these pairs are *flipped* by φ . Because of the choice of \mathbf{A}' , there are zero or two choices for χ for every e , and if there are two, then they are χ and $1 - \chi$ for some χ and both of them give the same result. Note that there may be no η such that $(f, \eta) \in \text{Dom}(\varphi)$.

Define a function $\theta: B \rightarrow B$ by putting

$$\theta((e, \chi)) = (\hat{\varphi}(e), \xi)$$

where

$$\xi(\hat{\varphi}(f)) = \begin{cases} \chi(f) & \text{if } \{e, f\} \notin F \\ 1 - \chi(f) & \text{if } \{e, f\} \in F. \end{cases}$$

First we verify that θ extends φ . Suppose $(e, \chi) \in \text{Dom}(\varphi)$. Write $\theta(e, \chi) = (\hat{\varphi}(e), \xi)$ and $\varphi(e, \chi) = (\hat{\varphi}(e), \chi')$. We must check that $\xi = \chi'$, so we let $f \in M$ and show that $\chi'(\hat{\varphi}(f)) = \xi(\hat{\varphi}(f))$. But this follows easily from the definitions of ξ and F , by considering the cases $\{e, f\} \in F$ and $\{e, f\} \notin F$ separately.

Now we check that θ is an automorphism of \mathbf{B} . It is easy to see that θ is one-to-one (one can construct its inverse) and that it preserves antipodal pairs. To check that $d_{\mathbf{B}}((e, \chi), (f, \eta)) = d_{\mathbf{B}}(\theta((e, \chi)), \theta((f, \eta)))$ for non-antipodal pairs, denote $\theta((e, \chi)) = (\hat{\varphi}(e), \chi')$ and $\theta((f, \eta)) = (\hat{\varphi}(f), \eta')$.

By definition of θ , we have $\chi(f) \neq \chi'(\hat{\varphi}(f))$ if and only if $\{e, f\} \in F$, and analogously, $\eta(e) \neq \eta'(\hat{\varphi}(e))$ if and only if $\{e, f\} \in F$. Putting this together (and again considering the cases $\{e, f\} \in F$ and $\{e, f\} \notin F$ separately) we get $\chi(f) = \eta(e)$ if and only if $\chi'(\hat{\varphi}(f)) = \eta'(\hat{\varphi}(e))$. Together with the definition of $d_{\mathbf{B}}$ this implies that indeed $d_{\mathbf{B}}((e, \chi), (f, \eta)) = d_{\mathbf{B}}(\theta((e, \chi)), \theta((f, \eta)))$.

Thus far, we have shown that \mathbf{B} is an EPPA-witness for \mathbf{A}' .

Now we put $\hat{p}((e, \chi)) = \chi(e)$. Note that $\hat{p}(\psi(x_i)) = \chi_i(e_i) = 0 = p(x_i)$ and similarly $\hat{p}(\psi(y_i)) = p(y_i)$. So \hat{p} extends $p \circ \psi^{-1}$.

Suppose φ preserves values of p . Then for $(e, \chi) \in \text{Dom}(\varphi)$ it holds that $\chi(e) = \chi'(\hat{\varphi}(e))$, where $(\hat{\varphi}(e), \chi') = \varphi((e, \chi))$. Thus there is no $e \in M$ such that $\{e\} \in F$ (for such an e , there would have to be some $(e, \chi) \in \text{Dom} \varphi$, of course). By definition of θ we immediately get $\hat{p}((e, \chi)) = \hat{p}(\theta((e, \chi)))$ for every $(e, \chi) \in B$.

Finally we verify coherence for the above construction. Consider partial automorphisms φ_1, φ_2 and φ of \mathbf{A}' such that φ is the composition of φ_1 and φ_2 . Denote by $\hat{\varphi}_1, \hat{\varphi}_2$ and $\hat{\varphi}$ their corresponding permutations of M constructed above. Let F_1, F_2 and F be the corresponding sets of flipped pairs and θ_1, θ_2 and θ the corresponding extensions. Because the permutations $\hat{\varphi}_1, \hat{\varphi}_2$ and $\hat{\varphi}$ were constructed by extending projections of φ_1, φ_2 and φ (which also compose coherently) in an order-preserving way, we know that $\hat{\varphi}$ is the composition of $\hat{\varphi}_1$ and $\hat{\varphi}_2$.

To see that θ is the composition of θ_1 and θ_2 one first checks that for $e \in \pi(\text{Dom}(\varphi)) = \pi(\text{Dom}(\varphi_1))$ and $f \in M$ one has that $\{e, f\} \in F$ if and only if

$$(\{e, f\} \notin F_1 \text{ and } \{\hat{\varphi}_1(e), \hat{\varphi}_1(f)\} \notin F_2) \text{ or } (\{e, f\} \in F_1 \text{ and } \{\hat{\varphi}_1(e), \hat{\varphi}_1(f)\} \in F_2).$$

The definition of θ_1, θ_2 and θ then gives the required result. In more detail, write $\theta_1(e, \chi) = (\hat{\varphi}_1(e), \xi_1)$ and $\theta_2\theta_1((e, \chi)) = \theta_2(\hat{\varphi}_1(e), \xi_1) = (\hat{\varphi}(e), \xi_2)$, where, for $f \in M$,

$$\xi_2(\hat{\varphi}_2(\hat{\varphi}_1(f))) = \begin{cases} \xi_1(\hat{\varphi}_1(f)) & \text{if } \{\hat{\varphi}_1(e), \hat{\varphi}_1(f)\} \notin F_2 \\ 1 - \xi_1(\hat{\varphi}_1(f)) & \text{if } \{\hat{\varphi}_1(e), \hat{\varphi}_1(f)\} \in F_2. \end{cases}$$

Applying the definition of ξ_1 and using the above observation finishes the calculation. □

5.4 Proofs of the main results

Theorem 5.1.3 is a direct consequence of Proposition 5.3.2. Next we use the correspondence between graphs with switching automorphisms and antipodal metric spaces (or, in the language of Seidel, double-covers of complete graphs [Cam99, Sei73]) to prove Theorem 5.1.2:

Proof of Theorem 5.1.2. Given a graph \mathbf{G} , we construct an antipodal metric space \mathbf{A} on vertex set $G \times \{0, 1\}$ with distances defined as follows:

1. $d_{\mathbf{A}}((x, 0), (x, 1)) = 3$ for every $x \in G$,
2. $d_{\mathbf{A}}((x, i), (y, i)) = 1$ for every $x \neq y$ forming an edge of \mathbf{G} and $i \in \{0, 1\}$,
3. $d_{\mathbf{A}}((x, i), (y, 1 - i)) = 1$ for every $x \neq y$ forming a non-edge of \mathbf{G} and $i \in \{0, 1\}$, and
4. $d_{\mathbf{A}}((x, i), (y, j)) = 2$ otherwise.

Let $p: A \rightarrow \{0, 1\}$ be a function defined by $p((x, i)) = i$. We apply Proposition 5.3.2 to construct an antipodal metric space \mathbf{C} and a function $\hat{p}: C \rightarrow \{0, 1\}$. Construct a graph \mathbf{H} with vertex set $\{x \in C : \hat{p}(x) = 0\}$ with x, y forming an edge if and only if $d_{\mathbf{C}}(x, y) = 1$.

Now consider a partial automorphism φ of \mathbf{G} . This automorphism corresponds to a partial automorphism φ' of \mathbf{A} by putting $\varphi'((x, i)) = (\varphi(x), i)$ for every $x \in \text{Dom}(\varphi)$ and $i \in \{0, 1\}$. φ' then extends to θ which preserves values of \hat{p} . Consequently, θ restricted to H is an automorphism of \mathbf{H} .

Finally consider a partial switching automorphism φ (i.e. a switching isomorphism of induced subgraphs). Let S be the set of vertices switched by φ . Now the partial automorphism of \mathbf{A} is defined by putting $\varphi'((x, i)) = (\varphi(x), i)$ if $x \notin S$ and $\varphi'((x, i)) = (\varphi(x), 1 - i)$ otherwise. Again extend φ' to θ and observe that θ gives a switching automorphism of \mathbf{H} . □

EPPA for two-graphs follows easily too:

Proof of Theorem 5.1.1. Let \mathbf{T} be a finite two-graph, pick an arbitrary vertex $x \in T$ and define a graph \mathbf{G} on the vertex set T such that $\{y, z\} \in E_{\mathbf{G}}$ if and only if $x \notin \{y, z\}$ and $\{x, y, z\}$ is a triple of \mathbf{T} . Observe that $\mathbf{T} = T(\mathbf{G})$. By Theorem 5.1.2, there is a graph \mathbf{H} containing \mathbf{G} such that every switching




Figure 5.1: Two possible antipodal quadruples for choice of a, b, c, d , $d_{\mathbf{A}}(a, b) = d_{\mathbf{A}}(c, d) = 3$.

isomorphism of induced subgraphs of \mathbf{G} extends to a switching automorphism of \mathbf{H} . We claim that $T(\mathbf{H})$ is an EPPA-witness for \mathbf{T} .

To prove that, let $\varphi: \mathbf{T} \rightarrow \mathbf{T}$ be a partial automorphism of \mathbf{T} . By the construction of \mathbf{G} and the correspondence between graphs and two-graphs, φ is a switching isomorphism of subgraphs of \mathbf{G} induced on the domain and range of φ respectively. Hence, it extends to a switching automorphism of \mathbf{H} and thus an automorphism of $T(\mathbf{H})$, which is what we wanted. \square

Remark 5.4.1. Observe that the EPPA-witness given in this proof of Theorem 5.1.1 is not necessarily a coherent EPPA-witness. The problem lies in the proof of Theorem 5.1.2, because for every switching isomorphism of subgraphs of \mathbf{G} there are two corresponding partial automorphisms of \mathbf{A} . For example, if the partial automorphism of \mathbf{G} is a partial identity, one partial automorphism of \mathbf{A} is also a partial identity, while the other flips all the involved edges of length 3. While the first is extended to the identity by the construction in Proposition 5.3.2, the other is extended to a non-trivial permutation of the edges of length 3. This issue carries over to Theorem 5.1.1, and in fact, it seems to be a fundamental obstacle for using antipodal metric spaces to prove coherent EPPA for two-graphs.

5.5 Existence of a dense locally finite subgroup

As discussed in Section 5.2.1 here, Solecki and Sinióra [SS19, Sol09] introduced the notion of *coherent EPPA* for a Fraïssé class as a way to prove that the automorphism group of the respective Fraïssé limit contains a dense locally finite subgroup. While we cannot prove coherent EPPA for \mathcal{T} , Theorem 5.1.3 gives coherent EPPA for the class of all antipodal metric spaces of diameter 3 and thus there is a dense locally finite subgroup of the automorphism group of its Fraïssé limit. We now show how this gives a dense, locally finite subgroup of the automorphism group of the Fraïssé limit of \mathcal{T} .

Let \mathbb{T} be the Fraïssé limit of \mathcal{T} and \mathbb{M} be the Fraïssé limit of the class of all finite antipodal metric spaces of diameter 3. The main result of this section is the following theorem.

Theorem 5.5.1. *There is a dense, locally finite subgroup of $\text{Aut}(\mathbb{T})$.*

In order to prove this, we first describe how to get a two-graph from an antipodal metric space, which is again a well-known construction [Cam99, Sei73].

In an antipodal metric space \mathbf{A} , every quadruple of distinct vertices a, b, c, d such that $d_{\mathbf{A}}(a, b) = d_{\mathbf{A}}(c, d) = 3$ (a pair of edges of label 3) induces one of the two (isomorphic) subspaces depicted in Figure 5.1 — we call these *antipodal quadruples*. However, three edges with label 3 can induce two non-isomorphic




Figure 5.2: Two non-isomorphic antipodal metric spaces with 6 vertices.

structures; the edges of length 1 either form two triangles, or one 6-cycle (see Figure 5.2). This motivates the following correspondence:

Definition 5.5.1 (From antipodal spaces to two-graphs). Let \mathbf{A} be an antipodal metric space. Let M be the set of all edges of \mathbf{A} of length 3 (thus, $|M| = \frac{|A|}{2}$ if \mathbf{A} is finite). Define $T(\mathbf{A})$ to be the 3-uniform hypergraph on vertex set M where $\{a, b, c\}$ is a hyperedge if and only if the substructure of \mathbf{A} induced on the edges a, b, c is isomorphic to that depicted in Figure 5.2 (b).

It is straightforward to verify that $T(\mathbf{A})$ is a two-graph. Clearly there is a natural two-to-one map $\mathbf{A} \rightarrow T(\mathbf{A})$ and this induces a group homomorphism $\text{Aut}(\mathbf{A}) \rightarrow \text{Aut}(T(\mathbf{A}))$. It is also well-known that there is a converse to this construction (usually expressed in terms of double covers of complete graphs). Suppose \mathbf{T} is a two-graph. Let \mathbf{G} be a graph in the switching class of \mathbf{T} (so $T(\mathbf{G}) = \mathbf{T}$) and let \mathbf{A} be the antipodal metric space constructed in the proof of Theorem 5.1.2. Then $T(\mathbf{A}) = \mathbf{T}$ and the proof of Theorem 5.1.2 shows that the map $\text{Aut}(\mathbf{A}) \rightarrow \text{Aut}(\mathbf{T})$ is surjective, as any automorphism of \mathbf{T} is a switching automorphism of \mathbf{G} . More generally, a similar argument gives the following well-known fact:

Lemma 5.5.2. *If $\mathbf{A}_1, \mathbf{A}_2$ are antipodal metric spaces and $\beta: T(\mathbf{A}_1) \rightarrow T(\mathbf{A}_2)$ is an isomorphism of two-graphs, then there is an isomorphism $\alpha: \mathbf{A}_1 \rightarrow \mathbf{A}_2$ which induces β .*

We can now give:

Proof of Theorem 5.5.1. Let \mathbf{T} be the two-graph $T(\mathbb{M})$. By Lemma 5.5.2, any isomorphism between finite substructures of \mathbf{T} lifts to an isomorphism between finite substructures of \mathbb{M} and so, as \mathbb{M} is homogeneous, this partial isomorphism is induced by an automorphism of \mathbb{M} . This shows that \mathbf{T} is homogeneous. The construction of an antipodal metric space from a two-graph shows that \mathbf{T} embeds every finite two-graph, and therefore \mathbf{T} is isomorphic to \mathbb{T} . So we have, again using Lemma 5.5.2, a surjective homomorphism $\alpha: \text{Aut}(\mathbb{M}) \rightarrow \text{Aut}(\mathbb{T})$. As already observed, there is a dense, locally finite subgroup H of $\text{Aut}(\mathbb{M})$. Then $\alpha(H)$ is a dense, locally finite subgroup of $\text{Aut}(\mathbb{T})$. □

5.6 Amalgamation property with automorphisms

Let \mathbf{A}, \mathbf{B}_1 and \mathbf{B}_2 be structures, α_1 an embedding of \mathbf{A} into \mathbf{B}_1 and α_2 an embedding of \mathbf{A} into \mathbf{B}_2 . Then every structure \mathbf{C} with embeddings $\beta_1: \mathbf{B}_1 \rightarrow \mathbf{C}$




Figure 5.3: A failure of APA for two-graphs

and $\beta_2: \mathbf{B}_2 \rightarrow \mathbf{C}$ such that $\beta_1 \circ \alpha_1 = \beta_2 \circ \alpha_2$ is called an *amalgamation* of \mathbf{B}_1 and \mathbf{B}_2 over \mathbf{A} with respect to α_1 and α_2 . Amalgamation is *strong* if $\beta_1(x_1) = \beta_2(x_2)$ if and only if $x_1 \in \alpha_1(A)$ and $x_2 \in \alpha_2(A)$.

For simplicity, in the following definition we will assume that all the embeddings in the definition of amalgamation are in fact inclusions.

Definition 5.6.1 (APA). Let \mathcal{C} be a class of finite structures. We say that \mathcal{C} has the *amalgamation property with automorphisms* (APA) if for every $\mathbf{A}, \mathbf{B}_1, \mathbf{B}_2 \in \mathcal{C}$ such that $\mathbf{A} \subseteq \mathbf{B}_1, \mathbf{B}_2$ there exists $\mathbf{C} \in \mathcal{C}$ which is an amalgamation of \mathbf{B}_1 and \mathbf{B}_2 over \mathbf{A} , has $\mathbf{B}_1, \mathbf{B}_2 \subseteq \mathbf{C}$ and furthermore the following holds:

For every pair of automorphisms $f_1 \in \text{Aut}(\mathbf{B}_1)$, $f_2 \in \text{Aut}(\mathbf{B}_2)$ such that $f_i(A) = A$ for $i \in \{1, 2\}$ and $f_1|_A = f_2|_A$, there is an automorphism $g \in \text{Aut}(\mathbf{C})$ such that $g|_{B_i} = f_i$ for $i \in \{1, 2\}$.

In other words, APA is a strengthening of the amalgamation property which requires that every pair of automorphisms of \mathbf{B}_1 and \mathbf{B}_2 which agree on \mathbf{A} extends to an automorphism of \mathbf{C} .

As was mentioned in the introduction, EPPA + APA imply the existence of ample generics [HHLS93] and it turns out that most of the known EPPA classes also have APA. We now show that this is not the case for two-graphs.

Proposition 5.6.1. *Let \mathbf{A} be the two-graph on two vertices, \mathbf{B}_1 be the two-graph on three vertices with no hyper-edge and \mathbf{B}_2 be the two-graph on three vertices which form a hyper-edge. It is possible to amalgamate \mathbf{B}_1 and \mathbf{B}_2 over \mathbf{A} , but it is not possible to amalgamate them with automorphisms.*

Proof. For convenience, we name the vertices as in Figure 5.3: $A = \{u, v\}$, $B_1 = \{u, v, x_1\}$ and $B_2 = \{u, v, x_2\}$. For $i \in \{1, 2\}$ let f_i be the automorphism of \mathbf{B}_i such that $f_i(x_i) = x_i$, $f_i(u) = v$ and $f_i(v) = u$. Clearly f_1 and f_2 agree on \mathbf{A} .

Consider the amalgamation problem for \mathbf{B}_1 and \mathbf{B}_2 over \mathbf{A} (with inclusions and with automorphisms f_1, f_2) and assume for a contradiction that there is $\mathbf{C} \in \mathcal{T}$ and an automorphism g of \mathbf{C} as in Definition 5.6.1. By the definition of \mathcal{T} , we get that there has to be an even number of triples in \mathbf{C} on $\{u, v, x_1, x_2\}$ and since we know that $\{u, v, x_1\}$ is not a triple and $\{u, v, x_2\}$ is a triple, there actually have to be precisely two triples on $\{u, v, x_1, x_2\}$. Therefore, exactly one of $\{u, x_1, x_2\}$ and $\{v, x_1, x_2\}$ has to form a triple in \mathbf{C} . But this means that g is not an automorphism (as g fixes x_1 and x_2 and swaps u and v), a contradiction.

On the other hand, if we only want to amalgamate \mathbf{B}_1 and \mathbf{B}_2 over \mathbf{A} (not with automorphisms), then this is clearly possible. \square

Remark 5.6.1. In the introduction we mentioned that there are also some pathological examples with EPPA but not APA:

1. Let \mathbf{M} be the two-vertex set with no structure. Its age consists of the empty set, the one-element set and \mathbf{M} . Consider the amalgamation problem for \mathbf{A} the empty set and $\mathbf{B}_1 = \mathbf{B}_2 = \mathbf{M}$. The amalgam has to be \mathbf{M} again. But then it is impossible to preserve all four possible pairs of automorphisms of \mathbf{B}_1 and \mathbf{B}_2 .

This phenomenon clearly happens because this is not a strong amalgamation class (and is exhibited by other non-strong amalgamation classes) and, indeed, disappears when we only consider closed structures.

2. Let now \mathbf{M} be the disjoint union of two infinite cliques, that is, an equivalence relation with two equivalence classes and consider its age. Let \mathbf{A} be the empty structure and $\mathbf{B}_1 = \mathbf{B}_2$ consist of two non-equivalent vertices. This amalgamation problem, again, does not have an APA-solution, because one needs to decide which pairs of vertices will be equivalent and this cannot preserve all four pairs of automorphisms.

This generalises to situations where the homogeneous structure has a definable equivalence relation with finitely many equivalence classes (or, more generally, of finite index).

However, the reason here is that the equivalence classes are algebraic in a quotient, i.e. a similar reason as in the previous point. One can either require the amalgamation problem to specify which classes go to which ones or, equivalently, consider an expansion which *weakly eliminates imaginaries* and the problem disappears.

While these two examples point out that, at least from the combinatorial point of view, one needs a more robust definition for APA, two-graphs seem to *innately not have APA*.

5.7 Ramsey expansion of two-graphs

As was pointed out recently, the methods for proving EPPA and the Ramsey property share many similarities, see e.g. [ABWH⁺17b]. The standard strategy is to study the *completion problem* for given classes and their expansions, see [ABWH⁺17b, EHN21, Kon19, HKN18]. EPPA is usually a corollary of one of the steps towards finding a Ramsey expansion.

The situation is very different for two-graphs. As we shall observe in this section, the Ramsey question can be easily solved using standard techniques, while for EPPA this is not the case (see Remark 5.8.1). This makes two-graphs an important example of the limits of the current methods and shows that the novel approach of this paper is in fact necessary.

We now give the very basic definitions of the structural Ramsey theory.

A class \mathcal{K} of structures is *Ramsey* if for every $\mathbf{A}, \mathbf{B} \in \mathcal{K}$ there exists $\mathbf{C} \in \mathcal{K}$ such that for every colouring of copies of \mathbf{A} in \mathbf{C} there exists a copy of \mathbf{B} in \mathbf{C} that is monochromatic. (By a *copy* of \mathbf{A} in \mathbf{C} we mean any substructure of \mathbf{C} isomorphic to \mathbf{A} .) This statement is abbreviated as $\mathbf{C} \longrightarrow (\mathbf{B})_k^{\mathbf{A}}$.

If class \mathcal{K} has *joint embedding property* (which means that for every $\mathbf{A}_1, \mathbf{A}_2 \in \mathcal{K}$ there is $\mathbf{C} \in \mathcal{K}$ which contains a copy of both \mathbf{A}_1 and \mathbf{A}_2) then EPPA of \mathcal{K} implies amalgamation. It is well known that Ramsey property also implies amalgamation property [Neš89, Neš05] under the assumption of joint embedding.

Every Ramsey class consists of rigid structures, i.e. structures with no non-trivial automorphism. The usual way to establish rigidity is to extend the language (in a model-theoretical way) by an additional binary relation \leq which fixes the ordering of vertices. It is thus a natural question whether the class $\vec{\mathcal{T}}$ of all two-graphs with a linear ordering of the vertices is Ramsey. We now show that the answer is negative and to do it, we need the following statement.

Proposition 5.7.1. *For every two-graph \mathbf{A} there exists a two-graph \mathbf{B} such that every graph in the switching class of \mathbf{B} contains a copy of every graph in the switching class of \mathbf{A} .*

Proof. Denote by \mathbf{G} the disjoint union of all graphs in the switching class of \mathbf{A} . Now let \mathbf{H} be a graph such that every colouring of vertices of \mathbf{H} by 2 colours induces a monochromatic copy of \mathbf{G} (that is, \mathbf{H} is vertex-Ramsey for \mathbf{G}) — it exists by a theorem of Folkman [Fol70, NR76a, NR77a]).

Every graph \mathbf{H}' in the switching class of \mathbf{H} induces a colouring of vertices of \mathbf{H} by two colours: \mathbf{H}' being in the switching class of \mathbf{H} means that there is a set $S \subseteq H$ such that $\mathbf{H}' = \mathbf{H}_S$, the colour classes are then S and $H \setminus S$ respectively.

By the construction of \mathbf{H} we find a copy $\tilde{\mathbf{G}} \subseteq \mathbf{H}$ of \mathbf{G} which is monochromatic with respect to this colouring. This however implies that the graphs induced by \mathbf{H} and \mathbf{H}' on the set \tilde{G} are isomorphic and thus \mathbf{H}' indeed contains every graph in the switching class of \mathbf{A} . Therefore, we can put \mathbf{B} to be the two-graph associated to \mathbf{H} . \square

Corollary 5.7.2. *The class $\vec{\mathcal{T}}$ is not Ramsey for colouring pairs of vertices.*

Proof. Let \mathbf{A} be the two-graph associated to an arbitrary graph containing both an edge and a non-edge, and let \mathbf{B} be the two-graph given by Proposition 5.7.1 for \mathbf{A} . Let $\vec{\mathbf{B}}$ be an arbitrary linear ordering of \mathbf{B} .

Assume that there exists an ordered two-graph $\vec{\mathbf{C}}$ such that

$$\vec{\mathbf{C}} \longrightarrow (\vec{\mathbf{B}})_2^{\vec{\mathbf{E}}}$$

where $\vec{\mathbf{E}}$ is the unique ordered two-graph on 2 vertices.

Let $\vec{\mathbf{I}}$ be an arbitrary graph from the switching class of $\vec{\mathbf{C}}$ (with the inherited order) and colour copies of $\vec{\mathbf{E}}$ red if they correspond to an edge of $\vec{\mathbf{I}}$ and blue otherwise. By the construction of \mathbf{B} it follows that there is no monochromatic copy of $\vec{\mathbf{B}}$, a contradiction. \square

Proposition 5.7.1 says that the expansion of two-graphs adding a particular graph from the switching class has the so-called *expansion property*. As a consequence of the Kechris–Pestov–Todorčević correspondence [KPT05], one gets that every Ramsey expansion of \mathcal{T} has to fix a particular representative of the switching class (see e.g. [NVT15] for details, it is also easy to see this directly). On the other hand, expanding any two-graph by a linear order and a particular graph from the given switching class is a Ramsey expansion by the Nešetřil–Rödl theorem [NR77b].

5.8 Remarks

Remark 5.8.1. The by-now-standard strategy for proving EPPA for class \mathcal{C} in, say, a relational language L may be summarized as follows:

1. Assume that, every pair of vertices of every structure in \mathcal{C} is in some relation. If it is not, we can always add a new binary relation to L and put all pairs into the relation.
2. Study the class \mathcal{C}^- which consists of all L -structures \mathbf{A}^- such that there is $\mathbf{A} \in \mathcal{C}$ which is a *completion* of \mathbf{A}^- , that is, \mathbf{A}^- and \mathbf{A} have the same vertex set A and there is $X \subset \mathcal{P}(A)$, a subset of the powerset of A , such that if $Y \in X$ and $Z \subseteq Y$, then $Z \in X$, and for each relation $R \in L$ it holds that $R^{\mathbf{A}^-} = R^{\mathbf{A}} \cap X$.
3. Find a finite family of L -structures \mathcal{F} such that \mathcal{C}^- is precisely $\text{Forb}_{\text{he}}(\mathcal{F})$, that is, the class of all finite L -structures \mathbf{B} such that there is no $\mathbf{F} \in \mathcal{F}$ with a homomorphism to \mathbf{B} .
4. Prove that in fact for every $\mathbf{A}^- \in \mathcal{C}^-$ there is $\mathbf{A} \in \mathcal{C}$ which is its *automorphism-preserving completion*, that is, \mathbf{A}^- can be obtained from \mathbf{A} as in point 2 and furthermore \mathbf{A}^- and \mathbf{A} have the same automorphisms.
5. Use the Herwig–Lascar theorem [HL00] omitting homomorphisms from \mathcal{F} to get EPPA-witnesses in $\text{Forb}_{\text{he}}(\mathcal{F})$.
6. Take the automorphism-preserving completion of the witnesses to obtain EPPA-witnesses in \mathcal{C} and thus prove EPPA for \mathcal{C} .

In various forms, this strategy has been applied many times, for example in [Sol05, Con19, ABWH⁺17b, Kon19, HKN18]. See also [HN19] where the notion of completions was introduced.

As we have seen in Section 5.6, \mathcal{T} does not admit automorphism-preserving completions (because APA is a weaker property). One can also prove (and it will appear elsewhere), using the negative result of Proposition 5.7.1 and Theorem 2.11 of [HN19], that \mathcal{T} cannot be described by finitely many forbidden homomorphisms (hence in particular there is no finite family \mathcal{F} satisfying point 3 above). This we believe is the first time the Ramsey techniques have been used to prove a negative result for EPPA.

Remark 5.8.2. \mathcal{T} is one of the five reducts of the random graph [Tho91]. Besides \mathcal{T} , the random graph itself and the countable set with no structure, the remaining two corresponding automorphism groups can be obtained by adding an isomorphism between the random graph and its complement and an isomorphism between the generic two-graph and its complement respectively.

By a similar argument, one can prove that the “best” Ramsey expansion (that is, with the expansion property) of these structures is still the ordered random graph. On the other hand, EPPA for these two classes is an open problem (we conjecture that neither of these two classes has EPPA).

Remark 5.8.3. Theorem 5.1.2 implies the following. For every graph \mathbf{G} there exists an EPPA-witness \mathbf{H} with the property that the two-graph associated to \mathbf{H}

is an EPPA-witness for the two-graph associated for \mathbf{G} , in other words, it implies that the class of all graphs and two-graphs respectively are a non-trivial positive example for the following question.

Question 5.8.1. *For which pairs of classes \mathcal{C} , \mathcal{C}^- such that \mathcal{C}^- is a reduct of \mathcal{C} does it hold that for every $\mathbf{A} \in \mathcal{C}$ there is $\mathbf{B} \in \mathcal{C}$ such that \mathbf{B} is an EPPA-witness for \mathbf{A} (in \mathcal{C}) and furthermore if \mathbf{A}^- and \mathbf{B}^- are the corresponding reducts in \mathcal{C}^- then \mathbf{B}^- is an EPPA-witness for \mathbf{A}^- (in \mathcal{C}^-)?*

Remark 5.8.4. As was already mentioned, ample generics are usually proved to exist by showing the combination of EPPA and APA. Siniora in his thesis [Sin17] asked if two-graphs have ample generics. This question still remains open, although we conjecture that it is not the case (ample generics are equivalent to having the so-called *weak amalgamation property for partial automorphisms* and the *joint embedding property for partial automorphisms* and it seems that the reasons for two-graphs not having APA are fundamental enough to also hold in the weak amalgamation context).¹

Acknowledgment

We would like to thank the anonymous referee for their helpful comments.

¹Added August 2019: in joint work with P. Simon, the authors have shown that this conjecture holds.

6. Extending partial isometries of antipodal graphs

MATĚJ KONEČNÝ

Abstract We prove EPPA (extension property for partial automorphisms) for all antipodal classes from Cherlin’s list of metrically homogeneous graphs, thereby answering a question of Aranda et al. This paper should be seen as the first application of a new general method for proving EPPA which can bypass the lack of automorphism-preserving completions. It is done by combining the recent strengthening of the Herwig–Lascar theorem by Hubička, Nešetřil and the author with the ideas of the proof of EPPA for two-graphs by Evans et al.

6.1 Introduction

Let $G = (V, E)$ be a (not necessarily finite) graph and let X, Y be subsets of V . We say that a function $f: X \rightarrow Y$ is a *partial automorphism* of G if f is an isomorphism of $G[X]$ and $G[Y]$, the graphs induced by G on X and Y respectively. This notion naturally extends to arbitrary structures (see Section 6.2).

In 1992 Hrushovski [Hru92] proved that for every finite graph G there is a finite graph H such that G is an induced subgraph of H and every partial automorphism of G extends to an automorphism of H . This property is, in general, called the extension property for partial automorphisms (EPPA):

Definition 6.1.1. Let \mathcal{C} be a class of finite structures. We say that \mathcal{C} has the *extension property for partial automorphisms* (or *EPPA*), also called the *Hrushovski property*, if for every $\mathbf{A} \in \mathcal{C}$ there is $\mathbf{B} \in \mathcal{C}$ such that \mathbf{A} is an (induced) substructure of \mathbf{B} and for every isomorphism f of substructures of \mathbf{A} there is an automorphism g of \mathbf{B} such that $f \subseteq g$. We call such \mathbf{B} an EPPA-witness for \mathbf{A} .

Hrushovski’s proof was group-theoretical, Herwig and Lascar [HL00] later gave a simple combinatorial proof by embedding G into the complement of a Kneser graph. After this, the quest of identifying new classes of structures with EPPA continued with a series of papers including [ABWH⁺17b, Con19, EHKN20, Her95, Her98, HL00, HO03, HKN19a, HKN18, HKN22, Kon19, Ott20, Sol05, Ver08].

Let $G = (V, E)$ be a graph. We say that a (partial) map $f: V \rightarrow V$ is *distance-preserving* if whenever u, v are in the domain of f , the distance between u and v is the same as the distance between $f(u)$ and $f(v)$. Clearly, every automorphism is distance-preserving. In 2005, Solecki [Sol05] (and independently also Vershik [Ver08]) proved that the class of all finite graphs has a variant of EPPA for distance-preserving maps. Namely, they proved that for every finite graph G there is a finite graph H satisfying the following:

1. G is an induced subgraph of H ,
2. whenever u, v are vertices of G , then the distance between u and v in G is the same as in H , and

3. every partial distance-preserving map of G extends to an automorphism of H .

It is not very convenient to work with distance-preserving maps, because they are relative to a graph and thus a distance-preserving map on a subgraph need not be distance-preserving with respect to a supergraph and vice versa. Given a graph $G = (V, E)$, it is more natural to consider the metric space $M = (V, d)$ where $d(u, v)$ is the number of edges of the shortest path from u to v in G (we will call this the *path-metric space* of G). And this is in fact what Solecki and Vershik did — they proved EPPA for all (integer-valued) metric spaces, which is equivalent to EPPA for graphs with distance-preserving maps.

Vershik’s proof is unpublished, Solecki’s proof uses a complicated general theorem of Herwig and Lascar [HL00, Theorem 3.2] about EPPA for structures with forbidden homomorphisms. Hubička, Nešetřil and the author [HKN19a] recently gave a simple self-contained proof of Solecki’s result. There is also a group theoretical proof by Sabok [Sab17] using a construction à la Mackey [Mac66].

This paper continues in this direction. Generalising the concept of distance transitivity, we say that a (countable) connected graph G is *metrically homogeneous* if every partial distance-preserving map of G with finite domain extends to an automorphism of G (so it is, in a sense, an EPPA-witness for itself). Cherlin [Che11] gave a list of countable metrically homogeneous graphs (which is conjectured to be complete and is provably complete in some cases [ACM21, Che22]) in terms of classes of finite metric spaces which embed into the path-metric space of the given metrically homogeneous graph. EPPA and other combinatorial properties of classes from Cherlin’s list were studied by Aranda, Bradley-Williams, Hng, Hubička, Karamanlis, Kompatscher, Pawliuk and the author [ABWH⁺17a, ABWH⁺21, ABWH⁺17b] (see also [Kon18]) and in [ABWH⁺17b] almost all the questions were settled, only EPPA for antipodal classes of odd diameter and bipartite antipodal classes of even diameter (see Section 6.2.3) remained open. An important step was later done by Evans, Hubička, Nešetřil and the author [EHKN20] who proved EPPA for antipodal metric spaces of diameter 3.

In this paper we combine the results of [ABWH⁺17b], ideas from [EHKN20], and the new strengthening of the Herwig–Lascar theorem by Hubička, Nešetřil and the author [HKN22] (stated here in a weaker form as Theorem 6.2.1) and prove the following theorem, thereby answering a question of Aranda et al. (Problem 1.3 in [ABWH⁺17b]) and completing the study of EPPA for classes from Cherlin’s list.

Theorem 6.1.1. *Every class of antipodal metric spaces from Cherlin’s list has EPPA.*

6.2 Preliminaries

A (not necessarily finite) structure \mathbf{A} is *homogeneous* if every partial automorphism of \mathbf{A} with finite domain extends to a full automorphism of \mathbf{A} itself (so it is, in a sense, an EPPA-witness for itself). Gardiner proved [Gar76] that the finite homogeneous graphs are precisely disjoint unions of cliques of the same size, their complements, the 5-cycle and the line graph of $K_{3,3}$. Lachlan and Woodrow

later [LW80] classified the countably infinite homogeneous graphs. These are disjoint unions of cliques of the same size (possibly infinite), their complements, the Rado graph, the K_n -free variants of the Rado graph and their complements.

Every homogeneous structure can be associated with the class of all (isomorphism types of) its finite substructures, which is called its *age*. By the Fraïssé theorem [Fra53], one can reconstruct the homogeneous structure back from this class (because it has the so-called *amalgamation property*). For more on homogeneous structures see the survey by Macpherson [Mac11].

A graph G is *vertex transitive* if for every pair of vertices u, v there is an automorphism sending u to v , it is *edge transitive* if every edge can be sent to every other edge by an automorphism and it is *distance transitive* if for every two pairs of vertices u, v and x, y such that the distance between u and v is the same as the distance between x and y there is an automorphism sending u to x and v to y .

Distance transitivity is a very strong condition. For example, there are only finitely many finite 3-regular distance transitive graphs [BS71] and the full catalogue is available in some other particular cases. However, for larger degrees, the classification is unknown, see e.g. the book by Godsil and Royle [GR01], largely devoted to the study of distance transitive graphs.

Recall that a connected graph G is *metrically homogeneous* if every partial distance-preserving map of G with finite domain extends to an automorphism of G . This is equivalent to saying that the path-metric space of G is homogeneous in the sense of the previous paragraphs. All connected homogeneous graphs are also metrically homogeneous, because every pair of vertices is either connected by an edge or by a path of length 2. Finite cycles of size at least 6 are examples of metrically homogeneous graphs which are not homogeneous.

Remark 6.2.1. If one checks the known classes with EPPA, they will find out that they all are ages of homogeneous structures. This is not a coincidence. It is easy to see that if a class of finite structures \mathcal{C} has EPPA and the *joint embedding property* (for every $\mathbf{A}, \mathbf{B} \in \mathcal{C}$ there is $\mathbf{C} \in \mathcal{C}$ which contains a copy of both of them), then \mathcal{C} is the age of a homogeneous structure provided that it contains at most countably many members up to isomorphism. This restricts the candidate classes for EPPA severely and connects finite combinatorics with the study of infinite homogeneous structures and infinite permutation groups.

In the other direction, EPPA has some implications for the automorphism group (with the *pointwise convergence topology*) of the corresponding homogeneous structure, see for example the paper of Hodges, Hodkinson, Lascar, and Shelah [H HLS93].

6.2.1 Γ_L -structures

An important feature of the strengthening of the Herwig–Lascar theorem by Hubička, Nešetřil and the author (Theorem 6.2.1) is that it allows to also permute the language. Namely, we will work with categories whose objects are the standard model-theoretic structures (in a given language), but the arrows are potentially richer, allowing a permutation of the language. The reader is invited to verify that in the following paragraphs, if the group Γ_L consists of the identity, one

obtains the usual notion of model-theoretic L -structures with the corresponding maps.

The following notions are taken from [HKN22], sometimes stated in a more special form which is sufficient for our purposes. Many of them were introduced by Hubička and Nešetřil [HN19] in the context of structural Ramsey theory (e.g. homomorphism-embeddings or completions).

Let $L = L_{\mathcal{R}} \cup L_{\mathcal{F}}$ be a language with relational symbols $R \in L_{\mathcal{R}}$, each having associated *arities* denoted by $a(R)$ and function symbols $F \in L_{\mathcal{F}}$. All functions in this paper are unary and have unary range. Let Γ_L be a permutation group on L such that each $\alpha \in \Gamma_L$ preserves the partition $L = L_{\mathcal{R}} \cup L_{\mathcal{F}}$ (that is, maps relations to relations and functions to functions) and the arities of all symbols. We will say that Γ_L is a *language equipped with a permutation group*.

A Γ_L -*structure* \mathbf{A} is a structure with *vertex set* A , functions $F_{\mathbf{A}}: A \rightarrow A$ for every $F \in L_{\mathcal{F}}$ and relations $R_{\mathbf{A}} \subseteq A^{a(R)}$ for every $R \in L_{\mathcal{R}}$. We will write structures in bold and their corresponding vertex sets in normal font. If Γ_L is trivial, we will often talk about L -structures instead of Γ_L -structures.

If the set A is finite we call \mathbf{A} a *finite structure*. If the language L contains no function symbols, we call L a *relational language* and say that a Γ_L -structure is a *relational Γ_L -structure*.

A *homomorphism* $f: \mathbf{A} \rightarrow \mathbf{B}$ is a pair $f = (f_L, f_A)$ where $f_L \in \Gamma_L$ and f_A is a mapping $A \rightarrow B$ such that for every $R \in L_{\mathcal{R}}$ and $F \in L_{\mathcal{F}}$ we have:

- (a) $(x_1, x_2, \dots, x_{a(R)}) \in R_{\mathbf{A}} \implies (f_A(x_1), f_A(x_2), \dots, f_A(x_{a(R)})) \in f_L(R)_{\mathbf{B}}$,
and
- (b) $f_A(F_{\mathbf{A}}(x)) = f_L(F)_{\mathbf{B}}(f_A(x))$.

For brevity, we will also write $f(x)$ for $f_A(x)$ in the context where $x \in A$ and $f(S)$ for $f_L(S)$ where $S \in L$. For a subset $A' \subseteq A$ we denote by $f(A')$ the set $\{f(x) : x \in A'\}$ and by $f(\mathbf{A})$ the homomorphic image of a structure \mathbf{A} . Note that we write $f: \mathbf{A} \rightarrow \mathbf{B}$ to emphasize that f respects the structure.

If f_A is injective then f is called a *monomorphism*. A monomorphism f is an *embedding* if for every $R \in L_{\mathcal{R}}$ we have the equivalence in the definition, that is,

$$(x_1, x_2, \dots, x_{a(R)}) \in R_{\mathbf{A}} \iff (f(x_1), f(x_2), \dots, f(x_{a(R)})) \in f(R)_{\mathbf{B}}.$$

If the inclusion $A \subseteq B$ together with the identity of Γ_L form an embedding, we say that \mathbf{A} is a *substructure* of \mathbf{B} and often denote it as $\mathbf{A} \subseteq \mathbf{B}$. For an embedding $f: \mathbf{A} \rightarrow \mathbf{B}$ we say that $f(\mathbf{A})$ is a *copy* of \mathbf{A} in \mathbf{B} . If f is an embedding where f_A is onto, then f is an *isomorphism* and an isomorphism $\mathbf{A} \rightarrow \mathbf{A}$ is called an *automorphism*.

Note that from the previous paragraph it follows that when L contains functions, not every subset of vertices induces a substructure. Namely, every substructure needs to be closed on functions. For example, if L consists of one unary function F , Γ_L contains only the identity and \mathbf{B} is a Γ_L -structure with vertex set $B = \{b_1, b_2\}$ such that $F(b_1) = b_2$ and $F(b_2)$ is not defined, then there is a substructure of \mathbf{B} on the set $\{b_2\}$, but the smallest substructure of \mathbf{B} containing b_1 is \mathbf{B} itself. Generalising this example, we say that for a Γ_L -structure \mathbf{B} and a set A which is a subset of B , the *closure of A in \mathbf{B}* , denoted by $\text{Cl}_{\mathbf{B}}(A)$, is the

smallest substructure of \mathbf{B} containing A . For $x \in B$, we will also write $\text{Cl}_{\mathbf{B}}(x)$ for $\text{Cl}_{\mathbf{B}}(\{x\})$.

Generalising the notion of a graph clique, we say that a Γ_L -structure \mathbf{A} is *irreducible* if for every pair of distinct vertices $x, y \in A$ there is a relation $R \in L$ and a tuple $\bar{r} \in A^{a(R)}$ containing both x and y such that $\bar{r} \in R_{\mathbf{A}}$. Note that the definition of irreducibility from [HKN22] is more general than this one (making more structures irreducible in general languages with functions), but stating it would need some more preliminary definitions and moreover they are equivalent for structures which we will consider in this paper.

Example 6.2.1. If the language only contains unary relations, irreflexive symmetric binary relations and unary functions (which will always be true in this paper), a structure is irreducible if and only if the union of the binary relations is a complete graph.

A homomorphism $f: \mathbf{A} \rightarrow \mathbf{B}$ is a *homomorphism-embedding* if the restriction $f \upharpoonright_{\mathbf{C}}$ is an embedding whenever \mathbf{C} is an irreducible substructure of \mathbf{A} .

6.2.2 EPPA for Γ_L -structures

We next state the main result of [HKN22] for which we need the following definitions, which are mostly variants of the definitions needed for the Hubička–Nešetřil theorem [HN19].

A *partial automorphism* of a Γ_L -structure \mathbf{A} is an isomorphism $f: \mathbf{C} \rightarrow \mathbf{C}'$ where \mathbf{C} and \mathbf{C}' are substructures of \mathbf{A} (remember that it also includes a permutation of the language which is not partial). We say that a class \mathcal{C} of finite Γ_L -structures has the *extension property for partial automorphisms (EPPA)* if for every $\mathbf{A} \in \mathcal{C}$ there is $\mathbf{B} \in \mathcal{C}$ such that \mathbf{A} is a substructure of \mathbf{B} and every partial automorphism of \mathbf{A} extends to an automorphism of \mathbf{B} . We call \mathbf{B} with such a property an *EPPA-witness for \mathbf{A}* . If \mathbf{B} is an EPPA-witness for \mathbf{A} , we say that it is *irreducible-structure faithful* if for every irreducible substructure \mathbf{C} of \mathbf{B} there exists an automorphism g of \mathbf{B} such that $g(C) \subseteq A$. We say that a class \mathcal{C} of finite Γ_L -structures *has EPPA* if there is an EPPA-witness $\mathbf{B} \in \mathcal{C}$ for every $\mathbf{A} \in \mathcal{C}$. We say that \mathcal{C} *has irreducible-structure faithful EPPA* if the witness can always be chosen to be irreducible-structure faithful.

Example 6.2.2.

1. Let \mathbf{A} be the graph on vertices u, v, w containing a single edge uv (here, the language consists of one binary relation and the permutation group is trivial). Then a possible (irreducible-structure faithful) EPPA-witness for \mathbf{A} is the graph \mathbf{B} on vertices u, v, w, x with edges uv and wx .
2. To see an example of a non-trivial permutation group, let L be the language consisting of unary relations R^i , where $1 \leq i \leq 10$ and let Γ_L consist of all permutation of L which fix R^{10} . Let \mathbf{A} be the Γ_L structure on one vertex v such that $R_{\mathbf{A}}^1 = \{v\}$ and $R_{\mathbf{A}}^i = \emptyset$ for every $i \geq 2$. Then every EPPA-witness \mathbf{B} for \mathbf{A} must contain vertices v_2, \dots, v_9 such that $v_i \in R_{\mathbf{B}}^i$ for every $2 \leq i \leq 9$, because \mathbf{B} needs to extend all partial automorphism f^i , $2 \leq i \leq 9$, such that f_A^i is the empty function and $f_L^i \in \Gamma_L$ sends R^1 to R^i .




Figure 6.1: Completions

Definition 6.2.1. Let \mathbf{C} be a Γ_L -structure. A Γ_L -structure \mathbf{C}' is a *completion* of \mathbf{C} if there exists an injective homomorphism-embedding $f: \mathbf{C} \rightarrow \mathbf{C}'$ which fixes every symbol of the language. We say that \mathbf{C}' is an *automorphism-preserving completion* of \mathbf{C} , if $C \subseteq C'$, the inclusion together with the identity from Γ_L give a homomorphism-embedding, for every $\alpha \in \text{Aut}(\mathbf{C})$ there is $\beta \in \text{Aut}(\mathbf{C}')$ such that $\alpha \subseteq \beta$ and moreover the map $\alpha \mapsto \beta$ is a group homomorphism $\text{Aut}(\mathbf{C}) \rightarrow \text{Aut}(\mathbf{C}')$.

In this paper, the languages will contain only unary and binary relations and unary functions, and moreover, whenever \mathbf{C}' will be a completion of \mathbf{C} , it will always hold that $C' = C$ and that the identity is a homomorphism-embedding. In such a case, for every relation $R \in L$ we have $R_{\mathbf{C}} \subseteq R_{\mathbf{C}'}$ with equality for unary R . For binary R it holds that if $(u, v) \in R_{\mathbf{C}'} \setminus R_{\mathbf{C}}$, then for every binary $R^0 \in L$ we have $(u, v) \notin R_{\mathbf{C}}^0$. Furthermore, \mathbf{C}' is an automorphism-preserving completion of \mathbf{C} if and only if $\text{Aut}(\mathbf{C}') = \text{Aut}(\mathbf{C})$.

Example 6.2.3. Consider the class $\mathcal{C}_{\mathbb{N}}$ of all finite integer-valued metric spaces understood as structures in a binary symmetric relational language L with a relation for every nonzero distance (the fact that $d(x, x) = 0$ is implicit). In Figure 6.1 we see the following:

- (a) An L -structure which has an automorphism-preserving completion in \mathcal{C} ,
- (b) one such completion, and
- (c) an L -structure which has no completion in \mathcal{C} .

Definition 6.2.2. Let L be a finite language with relations and unary functions equipped with a permutation group Γ_L . Let \mathcal{E} be a class of finite Γ_L -structures and let \mathcal{K} be a subclass of \mathcal{E} consisting of irreducible structures. We say that \mathcal{K} is a *locally finite subclass* of \mathcal{E} if for every $\mathbf{A} \in \mathcal{K}$ and every $\mathbf{B}_0 \in \mathcal{E}$ there is a finite integer $n = n(\mathbf{A}, \mathbf{B}_0)$ such that every Γ_L -structure \mathbf{B} has a completion $\mathbf{B}' \in \mathcal{K}$ provided that it satisfies the following:

1. For every vertex $v \in B$ we have that $\text{Cl}_{\mathbf{B}}(v)$ lies in a copy of \mathbf{A} ,
2. there is a homomorphism-embedding from \mathbf{B} to \mathbf{B}_0 , and
3. every substructure of \mathbf{B} with at most n vertices has a completion in \mathcal{K} .

We say that \mathcal{K} is a *locally finite automorphism-preserving subclass* of \mathcal{E} if in the condition above, the completion of \mathbf{B} can always be chosen to be automorphism-preserving.

Remark 6.2.2. While in Definition 6.2.2 we promise that $\text{Cl}_{\mathbf{B}}(v)$ lies in a copy of \mathbf{A} , the definition of local finiteness for the Hubička–Nešetřil theorem (Definition 2.4 from [HN19], similarly also the definition of local finiteness from [HKN22]) promises that every irreducible substructure of \mathbf{B} comes from \mathcal{K} . The difference here is due to the fact that the definition of irreducibility is simplified in this paper and does not work well for general languages with functions. In the applications, both conditions are used to ensure that closures behave well in \mathbf{B} .

Example 6.2.4. Let us observe that the class $\mathcal{C}_{\mathbb{N}}$ from Example 6.2.3 is a locally finite automorphism-preserving subclass of the class \mathcal{E} consisting of all finite L -structures (for L from Example 6.2.3), where all relations are symmetric and irreflexive and every pair of vertices is in at most one relation. Fix $\mathbf{A} \in \mathcal{K}$ and $\mathbf{B}_0 \in \mathcal{E}$. The assumption on \mathcal{E} justifies defining a symmetric partial function $d_{\mathbf{B}_0}: B_0^2 \rightarrow \mathbb{N}$ where $d(u, u) = 0$ and $d(u, v) = \ell$ if and only if u and v are in the relation corresponding to ℓ in \mathbf{B}_0 . Let S be the set of all integers ℓ for which there are vertices $u, v \in B_0$ such that $d(u, v) = \ell$. Since \mathbf{B}_0 is finite, S is also finite. Put $n = \max_{a, b \in S} \lceil \frac{a}{b} \rceil$.

Let \mathbf{B} be an L -structure satisfying the conditions of Definition 6.2.2 (since L contains no functions, condition 1 is satisfied trivially). The existence of a homomorphism-embedding from \mathbf{B} to \mathbf{B}_0 implies that all the relations in \mathbf{B} are also symmetric and irreflexive and every pair of vertices of \mathbf{B} is in at most one relation, hence we can analogously define a partial function $d_{\mathbf{B}}: B^2 \rightarrow \mathbb{N}$. Moreover, since there is a homomorphism-embedding from \mathbf{B} to \mathbf{B}_0 , we also get that the only non-empty distance relations in \mathbf{B} are those representing distances from S .

Next we define function $d': B^2 \rightarrow \mathbb{N}$ by

$$d'(x, y) = \min_{\mathbf{P} \text{ is a path } x \rightarrow y \text{ in } \mathbf{B}} \|\mathbf{P}\|,$$

where by a path $x \rightarrow y$ we mean a sequence of distinct vertices $x = p_1, \dots, p_k = y$ such that $d_{\mathbf{B}}(p_i, p_{i+1})$ is defined for every i satisfying $1 \leq i < k$, and we define $\|\mathbf{P}\|$ as $\sum_{i=1}^{k-1} d_{\mathbf{B}}(p_i, p_{i+1})$. It is easy to verify that (B, d') is a metric space with distances from \mathbb{N} and that $d_{\mathbf{B}} \subseteq d'$ if and only if \mathbf{B} contains no *non-metric cycles*, that is, sequences of vertices v_1, \dots, v_k such that $d_{\mathbf{B}}(v_i, v_{i+1})$ is defined for every $1 \leq i \leq k$ (we identify $v_{k+1} = v_1$) and $d_{\mathbf{B}}(v_1, v_k) > \sum_{i=1}^{k-1} d_{\mathbf{B}}(v_i, v_{i+1})$.

Since, clearly, non-metric cycles do not have a completion in $\mathcal{C}_{\mathbb{N}}$, it follows from the definition of n that \mathbf{B} contains no non-metric cycles and hence has a completion $\mathbf{B}' = (B, d')$ in $\mathcal{C}_{\mathbb{N}}$ as requested. Moreover, from the canonicity of the definition of d' it follows that this completion is automorphism-preserving and hence we have proved that $\mathcal{C}_{\mathbb{N}}$ is an automorphism-preserving completion of \mathbf{B} .

This construction of \mathbf{B}' is called the *shortest-path completion* in [HN19] and was already used by Solecki [Sol05] to prove EPPA for the class of all finite metric spaces and by Nešetřil [Neš07] to find a Ramsey expansion of the class of all finite metric spaces.

The main theorem of [HKN22] can be stated as follows.

Theorem 6.2.1 ([HKN22]). *Let L be a finite language with relations and unary functions equipped with a permutation group Γ_L , let \mathcal{E} be a class of finite Γ_L -structures which has irreducible-structure faithful EPPA and let \mathcal{K} be a hereditary locally finite automorphism-preserving subclass of \mathcal{E} with the strong amalgamation property, which consists of irreducible structures. Then \mathcal{K} has EPPA.*

Here \mathcal{K} is *hereditary* if whenever $\mathbf{B} \in \mathcal{K}$ and $\mathbf{A} \subseteq \mathbf{B}$, then also $\mathbf{A} \in \mathcal{K}$. We will not define what the *strong amalgamation property* is (see [HKN22]), but all classes for which we will use Theorem 6.2.1 will have this property.

Since Theorem 6.2.1 has a form of implication, we will need the following theorem from [HKN22] to supply us with the base EPPA class \mathcal{E} .

Theorem 6.2.2 ([HKN22]). *Let L be a finite language with relations and unary functions equipped with a permutation group Γ_L . Then the class of all finite Γ_L -structures has irreducible-structure faithful EPPA.*

Note that combining Example 6.2.4 with Theorems 6.2.1 and 6.2.2 gives a proof of Solecki’s result that finite metric spaces have EPPA (to prove that \mathcal{E} has EPPA, one has to use Theorem 6.2.2 for a finite fragment of L to get an EPPA-witness for a given $\mathbf{A} \in \mathcal{E}$). Both Theorems 6.2.1 and 6.2.2 are proved by an application of the method of valuation functions, a variant of which we will also use in this paper. More precisely, Theorem 6.2.2 is proved by giving an explicit construction of EPPA-witnesses. Theorem 6.2.1 iteratively applies the method of valuation functions to produce, given $\mathbf{A} \in \mathcal{K}$ and its irreducible-structure faithful EPPA-witness $\mathbf{B}_0 \in \mathcal{E}$, an EPPA-witness \mathbf{B} satisfying the conditions of Definition 6.2.2, the automorphism-preserving completion \mathbf{B}' of \mathbf{B} is then the desired EPPA-witness for \mathbf{A} in \mathcal{K} .

6.2.3 Metrically homogeneous graphs

Most of the details of Cherlin’s metric spaces are not important for this paper. We only give the necessary definitions and facts and refer the reader to [Che11], [ABWH⁺17b] or [Kon18].

All the metric spaces we will work with have distances from $\{0, 1, \dots, \delta\}$ for some integer δ . Therefore, we will view them interchangeably as pairs (A, d) where d is the metric, as complete graphs with edges labelled by $\{1, \dots, \delta\}$ (we will call these *complete $[\delta]$ -edge-labelled graphs*) where the labels of every triangle satisfy the triangle inequality, and as relational structures with trivial Γ_L and binary symmetric irreflexive relations R^1, \dots, R^δ (distance 0 is not represented) such that every pair of vertices is in exactly one relation and the triangle inequality is satisfied. The middle point of view works best with the notion of completions: Given a (not necessary complete) $[\delta]$ -edge-labelled graph \mathbf{G} , a $[\delta]$ -edge-labelled graph \mathbf{G}' is a completion of \mathbf{G} if \mathbf{G} is a non-induced subgraph of \mathbf{G}' and the labels are preserved.

We will say that two vertices are at distance a and that they are connected by an edge of length a interchangeably. In particular, when we talk about an edge of a $[\delta]$ -edge-labelled graph, we mean a pair of vertices such that their distance is defined, it does not necessarily mean that they are at distance 1.

A major part of Cherlin’s list of the classes of finite metric spaces which embed into the path-metric space of a countably infinite metrically homogeneous graph

consists of certain 5-parameter classes $\mathcal{A}_{K_1, K_2, C_0, C_1}^\delta$. These are classes of metric spaces with distances $\{0, 1, \dots, \delta\}$ (we call δ the *diameter* of such spaces) omitting certain families of triangles (e.g. triangles of short odd perimeter or triangles of long even perimeter).

A special case of these classes are the *antipodal* classes, where the five parameters have only two degrees of freedom. Here we will denote the antipodal classes as \mathcal{A}_K^δ , where $1 \leq K \leq \frac{\delta}{2}$, or $K = \delta$.¹ \mathcal{A}_K^δ is defined as the class of all finite metric spaces with distances from $\{0, 1, \dots, \delta\}$ such that they contain no triangle with distances a, b, c for which at least one of the following holds:

1. $a + b + c > 2\delta$,
2. $a + b + c$ is odd and $a + b + c < 2K$, or
3. $a + b + c$ is odd and $a + b + c > 2(\delta - K) + 2 \min(a, b, c)$.

However, for our purposes, we need only the following fact:

Fact 6.2.3 (Antipodal spaces). *The following holds in every class $\mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$ of antipodal metric spaces from Cherlin's list:*

1. *The edges of length δ form a matching (that is, for every vertex there is at most one vertex at distance δ from it) and for every $\mathbf{A} \in \mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$ there is a unique $\mathbf{B} \in \mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$ such that $\mathbf{A} \subseteq \mathbf{B}$, the edges of length δ form a perfect matching in \mathbf{B} and every edge of length δ in \mathbf{B} has at least one endpoint from \mathbf{A} .*
2. *For every pair of vertices u, v such that $d(u, v) = \delta$ and for every vertex w we have $d(u, w) + d(v, w) = \delta$.*
3. *If one selects exactly one vertex from each edge of length δ , the metric space they induce belongs to a special (non-antipodal) class of diameter $\delta - 1$ which we will call \mathcal{B}_K^δ .² And the other way around, one can get an antipodal metric space from every metric space $\mathbf{M} \in \mathcal{B}_K^\delta$ by taking two disjoint copies of \mathbf{M} , connecting every vertex to its copy by an edge of length δ and using point 2 to fill-in the missing distances.*

There are two kinds of antipodal classes with different combinatorial behaviour — those that come from a countable bipartite metrically homogeneous graph (they correspond to the case $K = \delta$) and those that come from a non-bipartite one. We will call the first the *bipartite* classes (their members have the property that they contain no triangles, or more generally cycles, of odd perimeter) and we will call the others the *non-bipartite* ones. This is slightly misleading, because some of the finite metric subspaces of the path-metric of a non-bipartite metrically homogeneous graph are surely bipartite, but it should not cause any confusion in this paper. The non-bipartite class of antipodal metric spaces of diameter 3 is closely connected to switching classes of graphs and to two-graphs (see [EHKN20]).

The following fact summarizes results from [ABWH⁺17b] about the non-bipartite odd diameter antipodal classes.

¹If $K \neq \delta$, the other parameters are then defined as $K_1 = K$, $K_2 = \delta - K$, $C_0 = 2\delta + 2$ and $C_1 = 2\delta + 1$, if $K = \delta$, then $K_1 = \infty$ and the other parameters are as before.

²It is in fact $\mathcal{A}_{K_1, K_2, C_0, C_1}^{\delta-1}$ for $K_1 = K$, $K_2 = \delta - K$, $C_0 = 2\delta + 2$ and $C_1 = 2\delta + 1$.

Fact 6.2.4. *Let $\mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$ be a non-bipartite class of antipodal metric spaces of odd diameter δ . Let \mathbf{A} be a $[\delta]$ -edge-labelled graph such that the edges of length δ of \mathbf{A} form a perfect matching and furthermore for every $u, v, w \in A$ such that $d_{\mathbf{A}}(u, v) = \delta$ and $w \neq u, v$, either w is not connected by an edge to either of u, v , or $d_{\mathbf{A}}(u, w) + d_{\mathbf{A}}(v, w) = \delta$. Suppose furthermore that \mathbf{A} contains none of the finitely many cycles forbidden in \mathcal{B}_K^δ .*

Let $f: \binom{A}{2} \rightarrow \{0, 1\}$ be a mapping satisfying the following.

1. *Whenever uv is an edge of \mathbf{A} , then $f(uv) \equiv d_{\mathbf{A}}(u, v) \pmod{2}$.*
2. *Let u_1v_1 and u_2v_2 be two different edges of length δ of \mathbf{A} . Then $f(u_1u_2) = f(v_1v_2)$, $f(u_1v_2) = f(u_2v_1)$ and $f(u_1u_2) \neq f(u_1v_2)$.*

Then there is $\bar{\mathbf{A}} \in \mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$ such that the following holds.

1. *$\bar{\mathbf{A}}$ is a completion of \mathbf{A} with the same vertex set,*
2. *for every edge uv of $\bar{\mathbf{A}}$ it holds that $f(uv) \equiv d_{\bar{\mathbf{A}}}(u, v) \pmod{2}$, and*
3. *Every automorphism of \mathbf{A} which preserves values of f is also an automorphism of $\bar{\mathbf{A}}$.*

Such $\bar{\mathbf{A}}$ can be constructed by picking one vertex from each edge of length δ , considering this auxiliary metric space of diameter $\delta - 1$, completing it using Theorem 4.9 from [ABWH⁺17b] (see also Lemma 4.18 from the same paper, or [HKK18]) and then pulling this completion back using f to decide parities of the edges. The proof then uses the observation that the completion procedure for $\mathcal{A}_{K_1, K_2, C_0, C_1}^{\delta-1}$ from [ABWH⁺17b] preserves the equivalence “ $a \sim \delta - a$ ”. That is, we say that two $[\delta - 1]$ -edge-labelled graphs \mathbf{G} and \mathbf{G}' are equivalent if they share the same vertex set and the same edge set and every edge has either the same label in both \mathbf{G} and \mathbf{G}' , or it has label a in \mathbf{G} and $\delta - a$ in \mathbf{G}' . The completion procedure then produces equivalent graphs whenever given equivalent graphs.

6.3 The odd diameter non-bipartite case

EPPA for the even diameter non-bipartite case was proved in [ABWH⁺17b]. In this section we prove the following proposition.

Proposition 6.3.1. *Let $\mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$ be a non-bipartite class of antipodal metric spaces of odd diameter. Then for every $\mathbf{A} \in \mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$ there is $\mathbf{B} \in \mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$ which is an EPPA-witness for \mathbf{A} .*

Proposition 6.3.1 extends the results of [EHKN20] where it was proved for diameter 3.

6.3.1 Motivation

We first give some motivation and intuition behind Proposition 6.3.1, as its proof is a bit technical. Consider the class \mathcal{A}_1^3 . It consists of all finite complete $[3]$ -edge-labelled graphs which omit triangles with distances $(1, 1, 3)$, $(2, 2, 3)$ and $(3, 3, a)$, where $1 \leq a \leq 3$. In other words, the edges of length 3 form a matching (and




Figure 6.2: Two possible (isomorphic) antipodal quadruples

by Fact 6.2.3 we can assume that it is a perfect matching), and if u, v, w, x are pairwise distinct vertices such that $d(u, v) = d(w, x) = 3$, then they form an *antipodal quadruple*, which means that $d(u, w) = d(v, x)$, $d(u, x) = d(v, w)$, and either $d(u, w) = 1$ and $d(u, x) = 2$, or $d(u, w) = 2$ and $d(u, x) = 1$ (see Figure 6.2).

Suppose that we want to find an EPPA-witness for a single edge of length 3 using Theorem 6.2.1. To do it, we in particular need to show that \mathcal{A}_1^3 is a locally finite automorphism-preserving subclass of the class \mathcal{E} of all [3]-edge-labelled graphs, which has irreducible-structure faithful EPPA by Theorem 6.2.2.

However, this does not hold. Take the disjoint union of two edges of length 3. This clearly has a completion in \mathcal{A}_1^3 (the antipodal quadruple), but it has no automorphism-preserving completion, because one has to pick which edges have length 1 and which edges have length 2.

In order to overcome this issue, we need to expand our structures by some information which will help us decide the parities. At the same time, we have to do it so that there is an expansion \mathbf{A}^+ of \mathbf{A} such that every partial automorphism of \mathbf{A} extends to a partial automorphism of \mathbf{A}^+ . This allows us to later forget the extra information and get an EPPA-witness for \mathbf{A} .

Let L^+ consist of the distance relations R^1 , R^2 and R^3 , a unary function M and two unary relations, T and B (for *top* and *bottom*), equipped with the permutation group Γ_{L^+} consisting of the identity and the transposition $(T B)$.

Let \mathcal{E} be the class of all finite [3]-edge-labelled graphs where the edges of length 3 form a perfect matching. Given $\mathbf{E} \in \mathcal{E}$, we say that a Γ_{L^+} -structure \mathbf{E}^+ is a *suitable expansion* of \mathbf{E} if the following hold:

1. \mathbf{E} and \mathbf{E}^+ share the same vertices and $R_{\mathbf{E}}^i = R_{\mathbf{E}^+}^i$ for every $1 \leq i \leq 3$ (that is, \mathbf{E}^+ and \mathbf{E} also share the distance relations),
2. $M_{\mathbf{E}^+}(u) = v$ if and only if $(u, v) \in R_{\mathbf{E}^+}^3$ (we need this for the strong amalgamation property),
3. every vertex of \mathbf{E}^+ is in precisely one of $T_{\mathbf{E}^+}$ and $B_{\mathbf{E}^+}$,
4. if $u, v \in E^+$ are connected by an edge of an odd length, then precisely one of $\{u, v\}$ is in $T_{\mathbf{E}^+}$ and the other is in $B_{\mathbf{E}^+}$, and
5. if $u, v \in E^+$ are connected by an edge of length 2, then either $\{u, v\} \subseteq T_{\mathbf{E}^+}$, or $\{u, v\} \subseteq B_{\mathbf{E}^+}$.

Note that not every $\mathbf{E} \in \mathcal{E}$ has a suitable expansion, however, \mathbf{A} has two of them and both preserve all partial automorphisms of \mathbf{A} (for this, we need the transposition $(T B)$).

Denote by \mathcal{E}^+ the class of all suitable expansions of structures from \mathcal{E} and similarly define \mathcal{A}_1^{3+} . Theorem 6.2.2 implies that \mathcal{E}^+ has irreducible-structure faithful EPPA.

In order to prove that \mathcal{A}_1^{3+} is a locally finite automorphism-preserving subclass of \mathcal{E}^+ , it is enough to observe that the conditions of Definition 6.2.2 imply that every such \mathbf{B} which we are asked to complete in fact comes from \mathcal{E}^+ , and if we pick $n = 6$, we get that it contains no triangles forbidden in \mathcal{A}_1^3 . It then suffices to define the missing distances according to the unary relations T and B : If uv is not an edge of \mathbf{B} , we put $d(u, v) = 2$ if they are in the same unary relation and $d(u, v) = 1$ otherwise.

In order to prove Proposition 6.3.1, we now generalise the construction above for larger diameters and arbitrary $\mathbf{A} \in \mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$. For the rest of this section, fix $\mathbf{A} \in \mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$. Using Fact 6.2.3, we can without loss of generality assume that for every vertex $v \in A$ there is a vertex $w \in A$ such that $d_{\mathbf{A}}(v, w) = \delta$. Enumerate the edges of \mathbf{A} of length δ as e_1, \dots, e_m and let $D = \{1, 2, \dots, m\}$ be their indices, that is, $|D| = \frac{|A|}{2}$ (we will sometimes treat D also as the set $\{e_1, \dots, e_m\}$ itself using the natural bijection). We furthermore denote $e_i = \{x_i, y_i\}$, where x_i and y_i are vertices of \mathbf{A} .

6.3.2 The expanded language

We will say that a function $\chi: D \rightarrow \{0, 1\}$ is a *valuation function*. For a set $F \subseteq D$, we denote by χ^F the *flip* of χ , that is, the function $D \rightarrow \{0, 1\}$ defined as

$$\chi^F(i) = \begin{cases} 1 - \chi(i) & \text{if } i \in F \\ \chi(i) & \text{otherwise,} \end{cases}$$

and for a permutation ψ of D we denote by χ_ψ the function satisfying $\chi_\psi(i) = \chi(\psi^{-1}(i))$. If χ is a valuation function, ψ is a permutation of D and $F \subseteq D$, then by χ_ψ^F we will mean $(\chi^F)_\psi$, that is, we first apply the flip and then the permutation.

Let L be the language consisting of binary symmetric irreflexive relations R^1, \dots, R^δ representing the distances, a unary function M , and unary relations U_i^X for every $1 \leq i \leq m$ and for every valuation function χ . If \mathbf{A} is an L -structure and v is a vertex of \mathbf{A} such that $v \in U_i^X$, we will say that v has a *unary mark* U_i^X . As in Section 6.3.1, the function M will ensure that the edges of length δ form a matching, the relations U_i^X are generalisations of the relations T and B .

Let $F \subseteq D^2$ be such that if $(i, j) \in F$, then also $(j, i) \in F$ (F is *symmetric*). For every $1 \leq i \leq m$ we let $F_i \subseteq D$ be the set $\{j \in D : (i, j) \in F\}$. We denote by α^F the permutation of L sending $U_i^X \mapsto U_i^{X^{F_i}}$, which fixes M and R^1, \dots, R^δ pointwise. In other words, α^F “flips” the mutual valuations of pairs from F .

For a permutation ψ of D , we denote by α_ψ the permutation of L sending $U_i^X \mapsto U_{\psi(i)}^{X_\psi}$, which fixes M and R^1, \dots, R^δ pointwise. Now we can define Γ_L as the group generated by

$$\{\alpha^F : F \subseteq D^2 \text{ and } F \text{ is symmetric}\} \cup \{\alpha_\psi : \psi \text{ is a permutation of } D\}.$$

Lemma 6.3.2. *For every member $g \in \Gamma_L$ there is a permutation ψ of D and a symmetric subset $F \subseteq D^2$ such that $g = \alpha_\psi \alpha^F$.*

Proof. Put

$$S = \{\alpha^F : F \subseteq D^2 \text{ and } F \text{ is symmetric}\} \cup \{\alpha_\psi : \psi \text{ is a permutation of } D\}.$$

We first show three claims:

Claim 6.3.3. *For every $\alpha^F, \alpha^{F'} \in S$ it holds that $\alpha^F \alpha^{F'} = \alpha^{F''}$, where F'' is the symmetric difference of F and F' (that is, $(i, j) \in F''$ if and only if it is in exactly one of F and F'). Consequently, $\alpha^F \alpha^F = 1$.*

Follows directly from the definitions of α^F and χ^F .

Claim 6.3.4. *For every $\alpha_\psi, \alpha_{\psi'} \in S$ it holds that $\alpha_\psi \alpha_{\psi'} = \alpha_{\psi''}$, where $\psi'' = \psi\psi'$. Consequently, $\alpha_\psi \alpha_{\psi^{-1}} = 1$.*

Again follows directly from the definitions of α_ψ and χ_ψ .

Claim 6.3.5. *For every $\alpha_\psi, \alpha^F \in S$ there is $\alpha^{F'} \in S$ such that $\alpha^F \alpha_\psi = \alpha_\psi \alpha^{F'}$.*

Put $F' = \psi^{-1}(F)$, that is, $F' = \{(\psi^{-1}(i), \psi^{-1}(j)) : (i, j) \in F\}$. The rest is straightforward verification.

We are now ready to prove the statement of this lemma. By definition, every member $g \in \Gamma_L$ can be written as a word consisting of members of S and their inverses. Using Claims 6.3.3 and 6.3.4, we can replace the inverses by members of S , using Claim 6.3.5 we can ensure that the word can be split into two subwords, first consisting only of α_ψ 's and the second consisting of α^F 's. From Claims 6.3.3 and 6.3.4 it follows that there are $\alpha_\psi, \alpha^F \in S$ such that indeed $g = \alpha_\psi \alpha^F$. \square

From now on we will thus denote members of Γ_L by α_ψ^F , where

$$\alpha_\psi^F(U_i^X) = \alpha_\psi(\alpha^F(U_i^X)) = U_{\psi(i)}^{X_{\psi(i)}^{F_i}},$$

and α_ψ^F is the identity on $\{M, R^1, \dots, R^\delta\}$. In other words, α_ψ^F first ‘‘flips’’ the mutual valuations of pairs from F and then permutes the set D .

For notational convenience, whenever \mathbf{C} is a Γ_L -structure, $U_i^\xi \in L$ and $u \in C$ is a vertex such that $u \in U_i^\xi$ and u has no other unary mark, we will denote by $\pi(u) = i$ its *projection* and by $\chi(u) = \xi$ its *valuation*. If u does not have precisely one unary mark, we leave $\pi(u)$ and $\chi(u)$ undefined.

The following (easy) observation says that the unary marks U_i^X indeed generalise the construction from Section 6.3.1.

Observation 6.3.6. *Let \mathbf{C} be a Γ_L -structure such that every vertex of \mathbf{C} has precisely one unary mark, let g be an automorphism of \mathbf{C} and let $u, v \in C$ be arbitrary vertices of \mathbf{C} . Then we have*

$$\chi(u)(\pi(v)) = \chi(v)(\pi(u))$$

if and only if

$$\chi(g(u))(\pi(g(v))) = \chi(g(v))(\pi(g(u))).$$

This implies that the function $f: \binom{C}{2} \rightarrow \{0, 1\}$, defined by $f(uv) = 0$ if $\chi(u)(\pi(v)) = \chi(v)(\pi(u))$ and $f(uv) = 1$ otherwise, is invariant under g and consequently under all automorphisms of \mathbf{C} .

Proof. Assume that $g = (\alpha_\psi^F, g_C)$ and put $F_u = \{j \in D : (\pi(u), j) \in F\}$ and $F_v = \{j \in D : (\pi(v), j) \in F\}$. Since g is an automorphism, we have

$$g(u) \in \alpha_\psi^F(U_{\pi(u)}^{\chi(u)}),$$

hence $\chi(g(u)) = \chi(u)_{\psi}^{F_u}$ and $\pi(g(u)) = \psi(\pi(u))$ and similarly $\chi(g(v)) = \chi(v)_{\psi}^{F_v}$ and $\pi(g(v)) = \psi(\pi(v))$.

It follows that

$$\chi(g(u))(\pi(g(v))) = \chi(u)_{\psi}^{F_u}(\psi(\pi(v))) = \begin{cases} 1 - \chi(u)(\pi(v)) & \text{if } \pi(v) \in F_u \\ \chi(u)(\pi(v)) & \text{otherwise,} \end{cases}$$

and similarly for $\chi(g(v))(\pi(g(u)))$. Since F is symmetric, we have that $\pi(v) \in F_u$ if and only if $\pi(u) \in F_v$ and thus the claim follows. \square

6.3.3 The class \mathcal{K} and completion to it

Let $\mathbf{C} \in \mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$. We say that a Γ_L -structure \mathbf{C}^+ is a *suitable expansion* of \mathbf{C} if the following hold:

1. \mathbf{C} and \mathbf{C}^+ share the same vertex set,
2. for every $1 \leq i \leq \delta$ we have that $R_{\mathbf{C}}^i = R_{\mathbf{C}^+}^i$,
3. $M_{\mathbf{C}^+}(u) = v$ if and only if $d_{\mathbf{C}^+}(u, v) = \delta$,
4. every vertex of \mathbf{C}^+ has precisely one unary mark,
5. if $d_{\mathbf{C}^+}(u, v) = \delta$ and $u \in U_i^\chi$ in \mathbf{C}^+ , then $v \in U_i^{1-\chi}$, where $(1 - \chi)(j) = 1 - \chi(j)$, and
6. in \mathbf{C}^+ it holds that $\chi(u)(\pi(v)) \neq \chi(v)(\pi(u))$ if and only if $d_{\mathbf{C}^+}(u, v)$ is odd.

Denote by \mathcal{K} the class of all suitable expansions of all $\mathbf{C} \in \mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$ where the edges of length δ form a perfect matching (Fact 6.2.3 says that this is without loss of generality; one can always uniquely and canonically add vertices so that this condition is satisfied). Note that it is possible that there is no suitable expansion of a given $\mathbf{C} \in \mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$.

Proposition 6.3.7. *\mathcal{K} is a locally finite automorphism-preserving subclass of \mathcal{E} , the class of all finite Γ_L -structures.*

Proof. Let n be a large enough integer (say, at least 4 and at least twice the number of vertices of the largest forbidden cycle in \mathcal{B}_K^δ) and let $\mathbf{A} \in \mathcal{K}$ and \mathbf{B} be as in Definition 6.2.2. Note that there is an unfortunate notational clash, this \mathbf{A} is different from the structure \mathbf{A} which we fixed at the beginning of this section.

The fact that for every $v \in B$ one has that $\text{Cl}_{\mathbf{B}}(v)$ lies in a copy of \mathbf{A} implies that $M_{\mathbf{B}}(u) = v$ if and only if $d_{\mathbf{B}}(u, v) = \delta$ and furthermore the edges of length δ form a perfect matching in \mathbf{B} (because this holds in \mathbf{A}).

The fact that every substructure of \mathbf{B} on at most n vertices has a completion in \mathcal{K} (which is promised by Definition 6.2.2) implies the following:

1. Every pair of vertices is in at most one distance relation R^i and these relations are symmetric and irreflexive,
2. every vertex of \mathbf{B} is in precisely one unary relation, and
3. if $d_{\mathbf{B}}(u, v) = \delta$ then $v \in U_{\pi(u)}^{1-\chi(u)}$.

We can assume that if $d_{\mathbf{B}}(u, v) = \delta$ and $w \neq u, v$ is a vertex of \mathbf{B} such that at least one of $d_{\mathbf{B}}(u, w)$, $d_{\mathbf{B}}(v, w)$ is defined, then in fact both distances are defined and furthermore $d_{\mathbf{B}}(u, w) + d_{\mathbf{B}}(v, w) = \delta$, because there is a unique way to complete it. It also follows that whenever u, v are vertices such that their distance is defined, then $\chi(u)(\pi(v)) \neq \chi(v)(\pi(u))$ if and only if $d_{\mathbf{B}}(u, v)$ is odd.

Finally, from the definition of n it also follows that \mathbf{B} contains no cycles forbidden in \mathcal{B}_K^δ (we needed n to be twice the number of vertices because Definition 6.2.2 talks about substructures and these need to be closed for functions). Hence if we define the function $f: \binom{B}{2} \rightarrow \{0, 1\}$ as $f(uv) = 0$ if $\chi(u)(\pi(v)) = \chi(v)(\pi(u))$ and $f(uv) = 1$ otherwise, Fact 6.2.4 gives us an automorphism-preserving way to add the remaining non- δ distances, which is exactly what we need for a completion to \mathcal{K} . \square

Let us remark that \mathcal{K} is hereditary: Whenever \mathbf{B} is a substructure of $\mathbf{C} \in \mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$ such that the edges of length δ form a perfect matching in both \mathbf{B} and \mathbf{C} , we have that if \mathbf{C}^+ is a suitable expansion of \mathbf{C} , then the substructure of \mathbf{C}^+ induced on the vertex set B is a suitable expansion of \mathbf{B} .

6.3.4 Constructing the witness

Recall that at the beginning of this section, we fixed $\mathbf{A} \in \mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$ and enumerated its edges of length δ as $e_1 = \{x_1, y_1\}, \dots, e_m = \{x_m, y_m\}$.

For $1 \leq i \leq m$, we define $\chi_i: D \rightarrow \{0, 1\}$ by putting

$$\chi_i(j) = \begin{cases} 1 & \text{if } i > j \text{ and } d_{\mathbf{A}}(x_i, x_j) \text{ is odd,} \\ 0 & \text{otherwise.} \end{cases}$$

We define a suitable expansion $\mathbf{A}^+ \in \mathcal{K}$ of \mathbf{A} by putting, for every $1 \leq i \leq m$, $M_{\mathbf{A}^+}(x_i) = y_i$, $M_{\mathbf{A}^+}(y_i) = x_i$, $x_i \in U_i^{\chi_i}$ and $y_i \in U_i^{1-\chi_i}$. Next we use Theorems 6.2.1 and 6.2.2 with Proposition 6.3.7 to get $\mathbf{B}^+ \in \mathcal{K}$ which is an EPPA-witness for \mathbf{A}^+ (so, in particular, $\mathbf{A}^+ \subseteq \mathbf{B}^+$). Finally, we put \mathbf{B} to be the *reduct* of \mathbf{B}^+ forgetting all the unary marks and the function M . Then indeed, $\mathbf{B} \in \mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$. And since $\mathbf{A}^+ \subseteq \mathbf{B}^+$, we also have $\mathbf{A} \subseteq \mathbf{B}$.

6.3.5 Extending partial automorphisms

We will show that \mathbf{B} extends all partial automorphisms of \mathbf{A} . Fix a partial automorphism φ of \mathbf{A} . Without loss of generality we can assume that whenever $d_{\mathbf{A}}(u, v) = \delta$ and $u \in \text{Dom}(\varphi)$, then also $v \in \text{Dom}(\varphi)$ (because there is a unique way of extending φ to v). Let $\psi: D \rightarrow D$ be an arbitrary permutation of D extending the action of φ on the edges of length δ of \mathbf{A} .

We now define a set $F \subseteq D^2$ of *flipping pairs*. We put (i, j) and (j, i) in F if $x_i \in \text{Dom}(\varphi)$ and $\chi(\varphi(x_i))(\psi(j)) \neq \chi(x_i)(j)$ in \mathbf{A}^+ . Note that if both x_i and x_j

are in the domain of φ then the outcome is the same if we consider x_j instead of x_i , because φ is an automorphism and therefore preserves the parity of $d_{\mathbf{A}}(x_i, x_j)$ and thus also the (non)-equality of the corresponding valuations. Note also that if we considered y_i instead of x_i , the outcome would still be the same.

What remains is to verify that the pair $(\alpha_{\psi}^F, \varphi)$ is a partial (Γ_L) -automorphism of \mathbf{A}^+ . Indeed, assuming that it is the case, we get that it extends to an automorphism (θ_L, θ) of \mathbf{B}^+ , where $\theta_L = \alpha_{\psi}^F$ and $\varphi \subseteq \theta$. But this means that θ is an automorphism of \mathbf{B} extending φ and hence \mathbf{B} is an EPPA-witness for \mathbf{A} . In the rest of this section we verify that $(\alpha_{\psi}^F, \varphi)$ is a partial automorphism of \mathbf{A}^+ . It amounts to (technical) checking that our construction does what it is supposed to do.

From the fact that φ is a partial automorphism of \mathbf{A} we get that $d_{\mathbf{B}^+}(u, v) = d_{\mathbf{B}^+}(\varphi(u), \varphi(v))$ whenever $u, v \in \text{Dom}(\varphi)$. This, together with the assumption that whenever $d_{\mathbf{A}}(u, v) = \delta$ and $u \in \text{Dom}(\varphi)$, then also $v \in \text{Dom}(\varphi)$, implies that if $v \in \text{Dom}(\varphi)$ then $M_{\mathbf{B}^+}(\varphi(v)) = \varphi(M_{\mathbf{B}^+}(v))$, or in other words, φ respects the function M .

It remains to verify that for every $v \in \text{Dom}(\varphi)$ and for every U_i^X we have $v \in U_i^X$ if and only if $\varphi(v) \in \alpha_{\psi}^F(U_i^X)$, or in other words, $\pi(\varphi(v)) = \psi(\pi(v))$ and $\chi(\varphi(v)) = \chi(v)_{\psi}^{F_v}$, where $F_v = \{j \in D : (\pi(v), j) \in F\}$. Since ψ extends the action of φ on the edges of length δ , and since for every $v \in A$ it holds that $\pi(v) = i$ if and only if $v \in e_i$, it follows that for every $v \in \text{Dom}(\varphi)$ we have $\pi(\varphi(v)) = \psi(\pi(v))$.

Analogously, from the definition of F we have that (i, j) and (j, i) are in F if and only if $x_i \in \text{Dom}(\varphi)$ and $\chi(\varphi(x_i))(\psi(j)) \neq \chi(x_i)(j)$. From the construction it follows that this happens if and only if $y_i \in \text{Dom}(\varphi)$ and $\chi(\varphi(y_i))(\psi(j)) \neq \chi(y_i)(j)$. We can summarize these two equivalences as follows: For every $v \in \text{Dom}(\varphi)$ and for every $j \in D$ we have $(\pi(v), j) \in F$ if and only if $\chi(\varphi(v))(\psi(j)) \neq \chi(v)(j)$.

By the definition of α^F , for every $v \in \text{Dom}(\varphi)$ and for every $j \in D$ we have that $\chi(v)_{\psi}^F(j) \neq \chi(v)(j)$ if and only if $(\pi(v), j) \in F$. Consequently, $\chi(v)_{\psi}^F(\psi(j)) \neq \chi(v)(j)$ if and only if $(\pi(v), j) \in F$, which happens if and only if $\chi(\varphi(v))(\psi(j)) \neq \chi(v)(j)$. It follows that $\chi(v)_{\psi}^F = \chi(\varphi(v))$ which concludes the proof of Proposition 6.3.1.

6.3.6 Remarks

1. If we extended the action of φ on the edges of length δ coherently (say, in an order-preserving way), we would get *coherent EPPA* (see [SS19]) as in [EHKN20].
2. The same strategy would also work for proving EPPA for antipodal metric spaces of even diameter, we would only need to pick a subset $O \subset \{0, 1, \dots, \delta\}$ such that $\delta \in O$ and precisely one of $a, \delta - a$ is in O for every $a \in \{0, 1, \dots, \delta\}$ and replace each occurrence of “odd distance” by “distance from O ” and “even distance” by “distance a such that $\delta - a \in O$ ”. (Note that for even δ , we have $\frac{\delta}{2} = \delta - \frac{\delta}{2}$, so $\frac{\delta}{2}$ “is both odd and even” in this sense.)
3. Cherlin also allows to forbid certain sets of $\{1, \delta - 1\}$ -valued metric spaces

(he calls them *Henson constraints*). We chose not to include these classes in order to avoid further technical complications, but using irreducible-structure faithfulness and the fact that the completion from Fact 6.2.4 does not create distances 1 and $\delta - 1$ gives EPPA also in this case.

6.4 The even diameter bipartite case

The odd diameter bipartite case was done in [ABWH⁺17b] (because every edge of length δ has one endpoint in each part of the bipartition and thus there is a unique way of determining parities of the distances, which implies that such classes admit automorphism-preserving completions), so it suffices to deal with the even diameter case. We prove the following proposition.

Proposition 6.4.1. *Let $\mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$ be a bipartite class of antipodal metric spaces of even diameter. Then for every $\mathbf{A} \in \mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$ there is $\mathbf{B} \in \mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$ which is an EPPA-witness of \mathbf{A} .*

The structure of the proof will be very similar to the odd non-bipartite case. We will also introduce some facts from [ABWH⁺17b] about completions, add unary functions and unary marks which will help us decide how to fill-in the missing distances while preserving all necessary automorphisms. We have to be a bit more careful in dealing with the bipartiteness (edges of length δ now lie inside the parts, so we need to make ψ preserve the bipartition, there are also infinitely many forbidden cycles — the odd perimeter ones), but the general structure is identical.

For the rest of the section, fix $\mathbf{A} \in \mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$. We can without loss of generality assume that every vertex $v \in A$ has some vertex $w \in A$ such that $d_{\mathbf{A}}(v, w) = \delta$. Consider the set $\{e_1, \dots, e_m\}$ of edges of \mathbf{A} of length δ and let $D = \{1, 2, \dots, m\}$ be their indices, that is, $|D| = \frac{|A|}{2}$. We denote $e_i = \{x_i, y_i\}$, where x_i and y_i are vertices of \mathbf{A} . Since \mathbf{A} is bipartite, we have that the relation “vertices u and v are at an even distance” is an equivalence relation on \mathbf{A} which has two equivalence classes. Because δ is even, we can assume that $D = D_1 \cup D_2$, where D_1 consists of the indices of edges with both endpoints in one part and D_2 consists of the indices of edges with both endpoints in the other part.

We also assume without loss of generality that $|D_1| = |D_2|$ (otherwise we can add more vertices to \mathbf{A} , and if this larger structure has an EPPA-witness \mathbf{B} , then it is also an EPPA-witness of the original \mathbf{A}).

We will need the following analogue of Fact 6.2.4.

Fact 6.4.2. *Let $\mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$ be a bipartite class of antipodal metric spaces. Let \mathbf{A} be a $[\delta]$ -edge-labelled graph such that the edges of length δ of \mathbf{A} form a perfect matching and furthermore for every $u, v, w \in A$ such that $d_{\mathbf{A}}(u, v) = \delta$ and $w \neq u, v$, either w is not connected by an edge to either of u, v , or $d_{\mathbf{A}}(u, w) + d_{\mathbf{A}}(v, w) = \delta$. Suppose furthermore that \mathbf{A} contains no odd-perimeter cycles and none of the finitely many even-perimeter cycles forbidden in \mathcal{B}_K^δ .*

Let $O \subset \{0, 1, \dots, \delta\}$ be a set such that $\delta \in O$ and exactly one of $a, \delta - a$ is in O for every $a \in \{0, 1, \dots, \delta\}$ and denote by $\delta - O$ the set $\{\delta - a : a \in O\}$.

Let $f: \binom{A}{2} \rightarrow \{0, 1\}$ be a mapping satisfying the following.

1. Whenever uv is an edge of \mathbf{A} , then $f(uv) = 1$ implies that $d_{\mathbf{A}}(u, v) \in O$ and $f(uv) = 0$ implies that $d_{\mathbf{A}}(u, v) \in \delta - O$.³

2. Let u_1v_1 and u_2v_2 be two different edges of length δ of \mathbf{A} . Then $f(u_1u_2) = f(v_1v_2)$, $f(u_1v_2) = f(u_2v_1)$ and $f(u_1u_2) \neq f(u_1v_2)$.

Then there is $\bar{\mathbf{A}} \in \mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$ such that the following holds.

1. $\bar{\mathbf{A}}$ is a completion of \mathbf{A} with the same vertex set,

2. for every edge uv of $\bar{\mathbf{A}}$ it holds that $f(uv) = 1$ implies that $d_{\bar{\mathbf{A}}}(u, v) \in O$ and $f(uv) = 0$ implies that $d_{\bar{\mathbf{A}}}(u, v) \in \delta - O$, and

3. Every automorphism of \mathbf{A} which preserves values of f is also an automorphism of $\bar{\mathbf{A}}$.

6.4.1 The expanded language

As in the odd non-bipartite case, we will call a function $\chi: D \rightarrow \{0, 1\}$ a *valuation function*, adopt the same notions of *flips* χ^F and permutations χ_ψ . We also let L be the same language as in Section 6.3, adding a unary function M and unary relations U_i^X .

In contrast to Section 6.3, we put Γ_L to be the group generated by

$$S = \{\alpha^F : F \subseteq D^2 \text{ and } F \text{ is symmetric}\} \cup \{\alpha_\psi : \psi \text{ is a partition-preserving permutation of } D\},$$

where ψ is *partition-preserving* if either $\psi(D_1) = D_1$ and $\psi(D_2) = D_2$, or $\psi(D_1) = D_2$ and $\psi(D_2) = D_1$. Analogously to Lemma 6.3.2 it follows that every element of Γ_L can be written as the product $\alpha_\psi \alpha^F$, where $\alpha_\psi, \alpha^F \in S$. We will denote $\alpha_\psi^F = \alpha_\psi \alpha^F$.

Again, for a vertex u in a Γ_L -structure which has precisely one unary mark U_i^ξ , we define $\pi(u) = i$ and $\chi(u) = \xi$ and we have the same observation with the same proof as before.

Observation 6.4.3. *Let \mathbf{C} be a Γ_L -structure such that every vertex of \mathbf{C} has precisely one unary mark, let g be an automorphism of \mathbf{C} and let $u, v \in C$ be arbitrary vertices of \mathbf{C} . Then we have*

$$\chi(u)(\pi(v)) = \chi(v)(\pi(u))$$

if and only if

$$\chi(g(u))(\pi(g(v))) = \chi(g(v))(\pi(g(u))).$$

This implies that the function $f: \binom{C}{2} \rightarrow \{0, 1\}$, defined by $f(uv) = 0$ if $\chi(u)(\pi(v)) = \chi(v)(\pi(u))$ and $f(uv) = 1$ otherwise, is invariant under g and consequently under all automorphisms of \mathbf{C} .

Note that the edge-labelled graph formed by the distance relations in \mathbf{C} may contain odd cycles.

³This seemingly sloppy statement is necessary in order to deal with $\frac{\delta}{2}$ being in both O and $\delta - O$ for even δ .

6.4.2 The class \mathcal{K} and completion to it

Now we also have to ensure that structures from \mathcal{K} are bipartite. Let $\mathbf{C} \in \mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$. Since \mathbf{C} is bipartite, we can denote by Q_1, Q_2 its parts (that is, $Q_1 \cup Q_2 = C$ and each of Q_1, Q_2 is an equivalence class of the relation “vertices u and v are at an even distance from each other”). We say that a Γ_L -structure \mathbf{C}^+ is a *suitable expansion of \mathbf{C}* if the following hold:

1. \mathbf{C} and \mathbf{C}^+ share the same vertex set,
2. for every $1 \leq i \leq \delta$ we have that $R_{\mathbf{C}}^i = R_{\mathbf{C}^+}^i$,
3. $M_{\mathbf{C}^+}(u) = v$ if and only if $d_{\mathbf{C}^+}(u, v) = \delta$,
4. every vertex of \mathbf{C}^+ has precisely one unary mark,
5. if $d_{\mathbf{C}^+}(u, v) = \delta$ and $u \in U_i^\chi$ in \mathbf{C}^+ , then $v \in U_i^{1-\chi}$,
6. in \mathbf{C}^+ it holds that if $\chi(u)(\pi(v)) \neq \chi(v)(\pi(u))$ then $d_{\mathbf{C}^+}(u, v) \in O$ and if $\chi(u)(\pi(v)) = \chi(v)(\pi(u))$ then $d_{\mathbf{C}^+}(u, v) \in \delta - O$, and
7. let $P_1 = \{v \in C : \pi(v) \in D_1\}$ and $P_2 = \{v \in C : \pi(v) \in D_2\}$ (where π is taken with respect to \mathbf{C}^+). Then either $P_1 = Q_1$ and $P_2 = Q_2$, or $P_1 = Q_2$ and $P_2 = Q_1$.

Denote by \mathcal{K} the class of all suitable expansions of all $\mathbf{C} \in \mathcal{A}_{K_1, K_2, C_1, C_2}^\delta$ where the edges of length δ form a perfect matching.

Proposition 6.4.4. *\mathcal{K} is a locally finite automorphism-preserving subclass of \mathcal{E} , the class of all finite Γ_L -structures.*

Proof. Let n be a large enough integer (say, at least 4 and at least twice the number of vertices of the largest even-perimeter forbidden cycle in \mathcal{B}_K^δ) and let $\mathbf{A} \in \mathcal{K}$ and \mathbf{B} be as in Definition 6.2.2. (Again, this is not the \mathbf{A} which we fixed at the beginning of this section.)

As for the odd non-bipartite case we get the following:

1. Every vertex of \mathbf{B} is in precisely one unary relation,
2. every pair of vertices is in at most one distance relation R^i (and these relations are symmetric),
3. $M_{\mathbf{B}}(u) = v \iff d_{\mathbf{B}}(u, v) = \delta$,
4. the edges of length δ form a perfect matching in \mathbf{B} ,
5. if $d_{\mathbf{B}}(u, v) = \delta$ then $v \in U_{\pi(u)}^{1-\chi(u)}$,
6. without loss of generality, we can assume that if $d_{\mathbf{B}}(u, v) = \delta$ and $w \neq u, v$ is a vertex of \mathbf{B} such that at least one of $d_{\mathbf{B}}(u, w), d_{\mathbf{B}}(v, w)$ is defined, then both distances are defined and furthermore $d_{\mathbf{B}}(u, w) + d_{\mathbf{B}}(v, w) = \delta$.
7. let u, v be vertices such that their distance is defined. Then $\chi(u)(\pi(v)) \neq \chi(v)(\pi(u))$ implies $d_{\mathbf{A}^+}(u, v) \in O$ and $\chi(u)(\pi(v)) = \chi(v)(\pi(u))$ implies $d_{\mathbf{A}^+}(u, v) \in \delta - O$.

Furthermore, from the last condition for a suitable expansion we also get that two vertices u, v of \mathbf{B} are at an even distance, if and only if there is $i \in \{1, 2\}$ such that $\pi(u), \pi(v) \in D_i$. Note that this implies that \mathbf{B} contains no cycles of odd perimeter (each cycle has to contain an even number of odd edges).

Finally, from the definition of n it also follows that \mathbf{B} contains no even-perimeter cycles forbidden in \mathcal{B}_K^δ . Hence if we define the function $f: \binom{B}{2} \rightarrow \{0, 1\}$ as $f(uv) = 0$ if $\chi(u)(\pi(v)) = \chi(v)(\pi(u))$ and $f(uv) = 1$ otherwise, Fact 6.4.2 gives us an automorphism-preserving way to add the remaining distances, which is exactly what we need for a completion to \mathcal{K} . \square

Let us again remark that \mathcal{K} is hereditary.

6.4.3 Constructing the witness

This is completely the same as for the odd diameter non-bipartite case. We define a Γ_L -structure \mathbf{A}^+ which is a suitable expansion of \mathbf{A} , and use Theorems 6.2.1 and 6.2.2 with Proposition 6.4.4 to get $\mathbf{B}^+ \in \mathcal{K}$ which is an EPPA-witness for \mathbf{A}^+ . Finally, we put \mathbf{B} to be the reduct of \mathbf{B}^+ forgetting all unary marks and all functions M .

6.4.4 Extending partial automorphisms

Again, this is completely the same as before with the exception that the permutation ψ of D has to preserve the bipartition $D = D_1 \cup D_2$ (it can exchange D_1 and D_2). Every partial automorphism φ of \mathbf{A} respects the bipartition, and since we assumed that $|D_1| = |D_2|$, it is always possible to extend φ to a full permutation ψ as needed.

Let us remark that if one is a bit more careful, the same strategy again gives coherent EPPA.

6.5 Conclusion

We are now ready to prove Theorem 6.1.1.

Proof of Theorem 6.1.1. In [ABWH⁺17b], EPPA is proved for the non-bipartite classes of even diameter and bipartite classes of odd diameter. Proposition 6.3.1 proves EPPA for non-bipartite classes of odd diameter and Proposition 6.4.1 proves EPPA for bipartite classes of even diameter, hence Theorem 6.1.1 is proved. \square

We think of this paper as the first example of a more general method for bypassing the lack of an automorphism-preserving completion, namely using the method of valuation functions to add more information to the structures (and thus restrict automorphisms) while preserving all partial automorphisms of one given structure \mathbf{A} , and then plugging this expanded class into the existing machinery. A similar trick can be done also for structures with higher arities, using higher-arity valuation functions (cf. [HKN22]). However, there are still classes where this method does not work, for example the class of tournaments which poses a long-standing important problem in this area.

6.6 Acknowledgements

I would like to thank Jan Hubička, Jaroslav Nešetřil and Gregory Cherlin for valuable advice and comments which significantly improved this paper. I would also like to thank all three anonymous referees for their incredibly helpful comments.

7. Ramsey expansions of 3-hypertournaments

GREGORY CHERLIN, JAN HUBIČKA, MATĚJ KONEČNÝ, JAROSLAV NEŠETŘIL

Abstract We study Ramsey expansions of certain homogeneous 3-hypertournaments. We show that they exhibit an interesting behaviour and, in one case, they seem not to submit to current gold-standard methods for obtaining Ramsey expansions. This makes these examples very interesting from the point of view of structural Ramsey theory as there is a large demand for novel examples.

7.1 Introduction

Structural Ramsey theory studies which homogeneous structures have the so-called Ramsey property, or at least are not far from it (can be *expanded* by some relations to obtain a structure with the Ramsey property). Recently, the area has stabilised with general methods and conditions from which almost all known Ramsey structures follow. In particular, the homogeneous structures offered by the classification programme are well-understood in most cases. Hence, there is a demand for new structures with interesting properties.

In this abstract we investigate Ramsey expansions of four homogeneous 4-constrained 3-hypertournaments identified by the first author [Che] and show that they exhibit an interesting range of behaviours. In particular, for one of them the current techniques and methods cannot be directly applied. There is a big demand for such examples in the area, in part because they show the limitations of present techniques, in part because they might lead to a negative answer to the question whether every structure homogeneous in a finite relational language has a Ramsey expansion in a finite relational language, one of the central questions of the area asked in 2011 by Bodirsky, Pinsker and Tsankov [BPT11].

7.2 Preliminaries

We adopt the standard notions of languages (in this abstract they will be relational only), structures and embeddings. A structure is *homogeneous* if every isomorphism between finite substructures extends to an automorphism. There is a correspondence between homogeneous structures and so-called (*strong*) *amalgamation classes* of finite structures, see e.g. [Hod93]. A structure \mathbf{A} is *irreducible* if every pair of vertices is part of a tuple in some relation of \mathbf{A} .

In this abstract, an *n-hypertournament* is a structure \mathbf{A} in a language with a single n -ary relation R such that for every set $S \subseteq A$ with $|S| = n$ it holds that the automorphism group of the substructure induced on S by \mathbf{A} is precisely $\text{Alt}(S)$, the alternating group on S . This in particular means that exactly half of n -tuples of elements of S with no repeated occurrences are in $R^{\mathbf{A}}$. For $n = 2$ we get standard tournaments, for $n = 3$ this correspond to picking one of the two possible cyclic orientations on every triple of vertices. It should be noted

however, that another widespread usage, going back at least to Assous [Ass86], requires a unique instance of the relation to hold on each n -set. A *holey n -hypertournament* is a structure \mathbf{A} with a single n -ary relation R such that all irreducible substructures of \mathbf{A} are n -hypertournaments. A *hole* in \mathbf{A} is a set of 3 vertices on which there are no relations at all.

Let $\mathbf{A}, \mathbf{B}, \mathbf{C}$ be structures. We write $\mathbf{C} \longrightarrow (\mathbf{B})_2^{\mathbf{A}}$ to denote the statement that for every 2-colouring of embeddings of \mathbf{A} to \mathbf{C} , there is an embedding of \mathbf{B} to \mathbf{C} on which all embeddings of \mathbf{A} have the same colour. A class \mathcal{C} of finite structures has the *Ramsey property (is Ramsey)* if for every $\mathbf{A}, \mathbf{B} \in \mathcal{C}$ there is $\mathbf{C} \in \mathcal{C}$ with $\mathbf{C} \longrightarrow (\mathbf{B})_2^{\mathbf{A}}$ and \mathcal{C}^+ is a *Ramsey expansion* of \mathcal{C} if it is Ramsey and can be obtained from \mathcal{C} by adding some relations. By an observation of Nešetřil [Neš05], every Ramsey class is an amalgamation class under some mild assumptions.

7.2.1 Homogeneous 4-constrained 3-hypertournaments

Suppose that $\mathbf{T} = (T, R)$ is a 3-hypertournament and pick an arbitrary linear order \leq on T . One can define a 3-uniform hypergraph $\hat{\mathbf{T}}$ on the set T such that $\{a, b, c\}$ with $a \leq b \leq c$ is a hyperedge of $\hat{\mathbf{T}}$ if and only if $(a, b, c) \in R$. (Note that by the definition of a 3-hypertournament, it always holds that exactly one of (a, b, c) and (a, c, b) is in R .) This operation has an inverse and hence, after fixing a linear order, we can work with 3-uniform hypergraphs instead of 3-hypertournaments. There are three isomorphism types of 3-hypertournaments on 4 vertices:

\mathbf{H}_4 The homogeneous 3-hypertournament on 4 vertices. For an arbitrary linear order \leq on H_4 , $\hat{\mathbf{H}}_4$ contains exactly two hyperedges. Moreover, they intersect in vertices $a < b$ such that there is exactly one $c \in H_4$ with $a < c < b$.

\mathbf{O}_4 The odd 3-hypertournament on 4 vertices. For an arbitrary linear order \leq , $\hat{\mathbf{O}}_4$ will contain an odd number of hyperedges. Conversely, any ordered 3-uniform hypergraph on 4 vertices with an odd number of hyperedges will give rise to \mathbf{O}_4 .

\mathbf{C}_4 The cyclic 3-hypertournament on 4 vertices. There is a linear order \leq on C_4 such that $\hat{\mathbf{C}}_4$ has all four hyperedges. In other linear orders, $\hat{\mathbf{C}}_4$ might have no hyperedges or exactly two which do not intersect as in \mathbf{H}_4 .

We say that a class \mathcal{C} of finite 3-hypertournaments is *4-constrained* if there is a non-empty subset $S \subseteq \{\mathbf{H}_4, \mathbf{O}_4, \mathbf{C}_4\}$ such that \mathcal{C} contains precisely those finite 3-hypertournaments whose every substructure on four distinct vertices is isomorphic to a member of S . There are four 4-constrained classes of finite 3-hypertournaments which form a strong amalgamation class [Che]. They correspond to the following sets S :

$S = \{\mathbf{C}_4\}$ The *cyclic* ones. These can be obtained by taking a finite cyclic order and orienting all triples according to it. Equivalently, they admit a linear order such that the corresponding hypergraph is complete.

$S = \{\mathbf{C}_4, \mathbf{H}_4\}$ The *even* ones. The corresponding hypergraphs satisfy the property that on every four vertices there are an even number of hyperedges.

$S = \{\mathbf{C}_4, \mathbf{O}_4\}$ The \mathbf{H}_4 -free ones. Note that in some sense, this generalizes the class of finite linear orders: As $\text{Aut}(\mathbf{H}_4) = \text{Alt}(4)$, one can define \mathbf{H}_n to be the $(n - 1)$ -hypertournament on n points such that $\text{Aut}(\mathbf{H}_n) = \text{Alt}(n)$. For $n = 3$, we get that \mathbf{H}_3 is the oriented cycle on 3 vertices and the class of all finite linear orders contains precisely those tournaments which omit \mathbf{H}_3 .

$S = \{\mathbf{C}_4, \mathbf{O}_4, \mathbf{H}_4\}$ The class of all finite 3-hypertournaments.

7.3 Positive Ramsey results

In this section we give Ramsey expansions for all above classes with the exception of the \mathbf{H}_4 -free ones. Let \mathcal{C}_c be the class of all finite cyclic 3-hypertournaments. Let $\vec{\mathcal{C}}_c$ be a class of finite linearly ordered 3-hypertournaments such that $(A, R, \leq) \in \mathcal{C}_c$ if and only if for every $x < y < z \in A$ we have $(x, y, z) \in R$. Notice that for every $(A, R) \in \mathcal{C}$ there are precisely $|A|$ orders \leq such that $(A, R, \leq) \in \vec{\mathcal{C}}_c$ (after fixing a smallest point, the rest of the order is determined by R), and conversely, for every $(A, R, \leq) \in \vec{\mathcal{C}}_c$ we have that $(A, R) \in \mathcal{C}_c$.

It is a well-known fact that every Ramsey class consists of linearly ordered structures [KPT05]. We have seen that after adding linear orders freely, the class of all finite ordered even 3-hypertournaments corresponds to the class of all finite ordered 3-uniform hypergraphs which induce an even number of hyperedges on every quadruple of vertices. These structures are called *two-graphs* and they are one of the reducts of the random graph (one can obtain a two-graph from a graph by putting hyperedges on triples of vertices which induce an even number of edges). Ramsey expansions of two-graphs have been discussed in [EHKN20] and the same ideas can be applied here.

Let $\vec{\mathcal{C}}_e$ consist of all finite structures (A, \leq, E, R) such that (A, \leq) is a linear order, (A, E) is a graph, (A, R) is a 3-hypertournament and for every $a, b, c \in A$ with $a < b < c$ we have that $(a, b, c) \in R$ if and only if there are an even number of edges (relation E) on $\{a, b, c\}$. Otherwise $(a, c, b) \in R$.

Theorem 7.3.1. *The 4-constrained classes of finite 3-hypertournaments with $S \in \{\{\mathbf{C}_4\}, \{\mathbf{C}_4, \mathbf{H}_4\}, \{\mathbf{C}_4, \mathbf{O}_4, \mathbf{H}_4\}\}$ all have a Ramsey expansion in a finite language. More concretely:*

1. $\vec{\mathcal{C}}_c$ is Ramsey.
2. $\vec{\mathcal{C}}_e$ is Ramsey.
3. The class of all finite linearly ordered 3-hypertournaments is Ramsey.

We remark that these expansions can be shown to have the so-called *expansion property* with respect to their base classes, which means that they are the optimal Ramsey expansions (see e.g. Definition 3.4 of [HN19]).

Proof. In $\vec{\mathcal{C}}_c$, R is definable from \leq and we can simply use Ramsey's theorem. Similarly, in $\vec{\mathcal{C}}_e$, R is definable from \leq and E , hence part 2 follows from the Ramsey property of the class of all ordered graphs [NR77b].

To prove part 3, fix a pair of finite ordered 3-hypertournaments \mathbf{A} and \mathbf{B} and use the Nešetřil–Rödl theorem [NR77b] to obtain a finite ordered holey 3-hypertournament \mathbf{C}' such that $\mathbf{C}' \rightarrow (\mathbf{B})_2^{\mathbf{A}}$. The holes in \mathbf{C}' can then be filled in arbitrarily to obtain a linearly ordered 3-hypertournament \mathbf{C} such that $\mathbf{C} \rightarrow (\mathbf{B})_2^{\mathbf{A}}$. \square

7.4 The \mathbf{H}_4 -free case

Let $\mathbf{A} = (A, R)$ be a holey 3-hypertournament. We say that $\bar{\mathbf{A}} = (A, R')$ is a *completion* of \mathbf{A} if $R \subseteq R'$ and $\bar{\mathbf{A}}$ is an \mathbf{H}_4 -free 3-hypertournament. Most of the known Ramsey classes can be proved to be Ramsey by a result of Hubička and Nešetřil [HN19]. In order to apply the result for \mathbf{H}_4 -free 3-hypertournaments, one needs a finite bound c such that whenever a holey 3-hypertournament has no completion, then it contains a substructure on at most c vertices with no completion. (Completions defined in [HN19] do not directly correspond to completions defined here. However, the definitions are equivalent for structures considered in this paper.) We prove the following.

Theorem 7.4.1. *There are arbitrarily large holey 3-hypertournaments \mathbf{B} such that \mathbf{B} has no completion but every proper substructure of \mathbf{B} has a completion.*

This theorem implies that one cannot use [HN19] directly for \mathbf{H}_4 -free hypertournaments. However, a situation like in Theorem 7.4.1 is not that uncommon. There are two common culprits for this, either the class contains orders (for example, failures of transitivity can be arbitrarily large in a holey version of posets) or it contains equivalences (again, failures of transitivity can be arbitrarily large). In the first case, there is a condition in [HN19] which promises the existence of a linear extension, and thus resolves the issue. For equivalences, one has to introduce explicit representatives for equivalence classes (this is called *elimination of imaginaries*) and unbounded obstacles to completion again disappear.

For \mathbf{H}_4 -free hypertournaments neither of the two solutions seems to work. This means that something else is happening which needs to be understood in order to obtain a Ramsey expansion of \mathbf{H}_4 -free tournaments. Hopefully, this would lead to new, even stronger, general techniques.

In the rest of the abstract we sketch a proof of Theorem 7.4.1.

Lemma 7.4.2.

1. Let $\mathbf{G} = (G, R)$ be a holey 3-hypertournament with $G = \{1, 2, 3, 4\}$ such that $(1, 3, 4) \in R$, $(1, 4, 2) \in R$ and $\{1, 2, 3\}$ and $\{2, 3, 4\}$ are holes. Let (G, R') be a completion of \mathbf{G} . If $(1, 2, 3) \in R'$, then $(2, 3, 4) \in R'$.
2. Let $\mathbf{G}^\neg = (G, R)$ be a holey 3-hypertournament with $G = \{1, 2, 3, 4\}$ such that $(2, 4, 3) \in R$, $(1, 4, 2) \in R$ and $\{1, 2, 3\}$ and $\{1, 3, 4\}$ are holes. Let (G, R') be a completion of \mathbf{G}^\neg . If $(1, 2, 3) \in R'$, then $(1, 3, 4) \notin R'$.

Proof. In the first case, suppose that $(1, 2, 3) \in R'$. If $(2, 4, 3) \in R'$, then (G, R') is isomorphic to \mathbf{H}_4 . Hence $(2, 3, 4) \in R'$. The second case is proved similarly. \square

Suppose that $\mathbf{A} = (A, R)$ is a holey 3-hypertournament. For $x, y, z, w \in A$, we will write $xyz \Rightarrow yzw$ if the map $(1, 2, 3, 4) \mapsto (x, y, z, w)$ is an embedding $\mathbf{G} \rightarrow \mathbf{A}$ and we will write $xyz \Rightarrow \neg xzw$ if the map $(1, 2, 3, 4) \mapsto (x, y, z, w)$ is an embedding $\mathbf{G}^\neg \rightarrow \mathbf{A}$. Using the complement of \mathbf{G} , we can define $\neg xyz \Rightarrow \neg yzw$, and using the complement of \mathbf{G}^\neg we can define $\neg xyz \Rightarrow xzw$. This notation can be chained as well, e.g. $xyz \Rightarrow yzw \Rightarrow zwu \Rightarrow \neg zuv$ means that all of $xyz \Rightarrow yzw$, $yzw \Rightarrow zwu$, $zwu \Rightarrow \neg zuv$ are satisfied.

Let $n \geq 6$. We denote by $\mathbf{O}_n = (O_n, R)$ the holey 3-hypertournament with vertex set $O_n = \{1, \dots, n\}$ such that

$$123 \Rightarrow 234 \Rightarrow 345 \Rightarrow \dots \Rightarrow (n-2)(n-1)n \Rightarrow \neg(n-2)n1 \Rightarrow \neg n12 \Rightarrow \neg 123.$$

All triples not covered by these conditions are holes.

Lemma 7.4.3.

1. There is a completion (O_n, R') of \mathbf{O}_n .
2. If (O_n, R') is a completion of \mathbf{O}_n , then $(1, 2, 3) \notin R'$.
3. For every $v \in O_n \setminus \{1, 2, 3\}$ there is a completion $(O_n \setminus \{v\}, R')$ of the structure induced by \mathbf{O}_n on $O_n \setminus \{v\}$ such that $(1, 2, 3) \in R'$.

Proof. For part 1, observe that every set of four vertices of \mathbf{O}_n with at least two different subsets of three vertices covered by a relation is isomorphic to \mathbf{G} , \mathbf{G}^\neg or the complement of \mathbf{G} . It follows that whenever $x, y, z \in O_n$ is a hole such that $x < y < z$, we can put (x, z, y) and its cyclic rotations in R' to get a completion. Part 2 follows by induction on the conditions.

For part 3, we put $(1, 2, 3), (2, 3, 4), \dots, (v-3, v-2, v-1) \in R'$, $(v+1, v+3, v+2), \dots, (n-2, n, n-1) \in R'$ and $(n-2, n, 1), (n, 1, 2) \in R'$. It can be verified that this does not create any copies of \mathbf{H}_4 . A completion of (O_n, R') exists as the class of all finite \mathbf{H}_4 -free tournaments has strong amalgamation. \square

Similarly, for $n \geq 6$, we define $\mathbf{O}_n^\neg = (O_n^\neg, R)$ the holey 3-hypertournament with vertex set $O_n^\neg = \{1, \dots, n\}$ such that

$$\neg 123 \Rightarrow \neg 234 \Rightarrow \neg 345 \Rightarrow \dots \Rightarrow \neg(n-2)(n-1)n \Rightarrow (n-2)n1 \Rightarrow n12 \Rightarrow 123$$

and there are no other relations in R . In any completion (O_n^\neg, R') of \mathbf{O}_n^\neg it holds that $(1, 2, 3) \in R'$, in fact, an analogue of Lemma 7.4.3 holds for \mathbf{O}_n^\neg .

Let \mathbf{B}_n be the holey 3-hypertournament obtained by gluing a copy of \mathbf{O}_n with a copy of \mathbf{O}_n^\neg , identifying vertices 1, 2 and 3. (This means that \mathbf{B}_n has $2n - 3$ vertices.) We now use $\{\mathbf{B}_n : n \geq 6\}$ to prove Theorem 7.4.1.

Proof of Theorem 7.4.1. Assume that (B_n, R') is a completion of \mathbf{B}_n . So in particular, it is a completion of the copies of \mathbf{O}_n and \mathbf{O}_n^\neg . By Lemma 7.4.3 and its analogue for \mathbf{O}_n^\neg , we have that $(1, 2, 3) \notin R'$ and $(1, 2, 3) \in R'$, a contradiction.

Pick $v \in B_n$ and consider the structure \mathbf{B}_n^v induced by \mathbf{B}_n on $B_n \setminus \{v\}$. We prove that \mathbf{B}_n^v has a completion. If $v \notin \{1, 2, 3\}$, one can use part 3 of Lemma 7.4.3 and its analogue for \mathbf{O}_n^\neg to complete the copy of \mathbf{O}_n and \mathbf{O}_n^\neg (one of them missing a vertex) so that they agree on $\{1, 2, 3\}$. Using strong amalgamation, we get a completion of \mathbf{B}_n^v . If $v \in \{1, 2, 3\}$, we pick an arbitrary completion of \mathbf{O}_n and \mathbf{O}_n^\neg , remove v from both of them, and let the completion of \mathbf{B}_n to be the strong amalgamation of the completions over $\{1, 2, 3\} \setminus v$. \square

The following question remains open.

Question 7.4.4. *What is the optimal Ramsey expansion for the class of all finite \mathbf{H}_4 -free hypertournaments? Does it have a Ramsey expansion in a finite language?*

8. Big Ramsey degrees of 3-uniform hypergraphs are finite

MARTIN BALKO, DAVID CHODOUNSKÝ, JAN HUBIČKA, MATĚJ KONEČNÝ,
LLUIS VENA

Abstract We prove that the universal homogeneous 3-uniform hypergraph has finite big Ramsey degrees. This is the first case where big Ramsey degrees are known to be finite for structures in a non-binary language.

Our proof is based on the vector (or product) form of Milliken’s Tree Theorem and demonstrates a general method to carry existing results on structures in binary relational languages to higher arities.

8.1 Introduction

Given 3-uniform hypergraphs \mathbf{A} and \mathbf{B} , we denote by $\binom{\mathbf{B}}{\mathbf{A}}$ the set of all embeddings from \mathbf{A} to \mathbf{B} . We write $\mathbf{C} \longrightarrow (\mathbf{B})_{k,\ell}^{\mathbf{A}}$ to denote the following statement:

For every colouring χ of $\binom{\mathbf{C}}{\mathbf{A}}$ with k colours, there exists an embedding $f: \mathbf{B} \rightarrow \mathbf{C}$ such that χ does not take more than ℓ values on $\binom{f[\mathbf{B}]}{\mathbf{A}}$.

For a countably infinite structure \mathbf{B} and its finite induced sub-structure \mathbf{A} , the *big Ramsey degree* of \mathbf{A} in \mathbf{B} is the least number $\ell \in \omega + 1$ such that $\mathbf{B} \longrightarrow (\mathbf{B})_{k,\ell}^{\mathbf{A}}$ for every $k \in \omega$; see [KPT05]. A countably infinite structure \mathbf{B} has *finite big Ramsey degrees* if the big Ramsey degree of \mathbf{A} in \mathbf{B} is finite for every finite substructure \mathbf{A} of \mathbf{B} .

A countable hypergraph \mathbf{A} is *(ultra)homogeneous* if every isomorphism between finite induced sub-hypergraphs extends to an automorphism of \mathbf{A} . It is well known that there is (up to isomorphism) a unique countable homogeneous 3-uniform hypergraph \mathbf{H} with the property that every countable 3-uniform hypergraph can be embedded into \mathbf{H} , see e.g. [Mac11].

Solving a question of Sauer¹ we prove the following result, which was announced in [BCH⁺19].

Theorem 8.1.1. *The universal homogeneous 3-uniform hypergraph \mathbf{H} has finite big Ramsey degrees.*

Our result is a contribution to the ongoing project of characterising big Ramsey degrees of homogeneous structures [KPT05], [Tod10, Chapter 6]. The origin of this project is in the work of Galvin [Gal68, Gal69] who proved that the big Ramsey degree of pairs in the order of the rationals, denoted by (\mathbb{Q}, \leq) , is equal to 2. Subsequently, Laver in late 1969 proved that in fact (\mathbb{Q}, \leq) has finite big Ramsey degrees, see [Dev79, Page 73],[EH74, Lav84] and Devlin determined the exact values of ℓ [Dev79], [Tod10, Theorems 6.22 and 6.23]. Using Milliken’s Tree

¹Personal communication, 2014.

Theorem (Theorem 8.2.2) for a single binary tree his argument is particularly intuitive: The vertices of a binary tree can be seen as finite $\{0, 1\}$ -words and those as rationals in the range $(0, 1)$ written as binary numbers (with additional digit 1 added to the end of each word to avoid ambiguities). Since this is a dense linear order it follows that (\mathbb{Q}, \leq) can be embedded to it. A colouring of finite subsets of \mathbb{Q} corresponds then to a colouring of finite subtrees and thus leads to an application of Milliken’s Tree Theorem, see [Tod10, Section 6.3] for further details.

Similar ideas can be applied to graphs; here a graph is coded using a binary tree, where 1 is used to code an edge. In the *passing number representation*, a pair of words w, w' with $|w| \leq |w'|$ is adjacent if $w'_{|w|} = 1$ [Tod10, Theorem 6.25]. Here, $|w|$ denotes the length of the word w and w_i is the letter of w on index i , where the indices start from 0. This representation was used by Sauer [Sau06] and Laflamme, Sauer, Vuksanovic [LSV06] who, in 2006, characterised big Ramsey degrees of the Rado graph. This was refined to unconstrained structures in binary languages [LSV06] and additional special classes [DLS16, NVT09, LNVTS10, Maš20]. Milliken’s Tree Theorem remained the key in all results in the area (here we consider Ramsey’s Theorem as a special case of Milliken’s Tree Theorem for the unary tree). See also [Dob20c] for a recent survey.

A generalization of these results to structures of higher arities (such as hypergraphs) and to structures forbidding non-trivial substructures remained open for over a decade. Recent connections to topological dynamics [Zuc19] renewed the interest in the area and both these problems were solved recently. Using set-theoretic techniques, Dobrinen [Dob20a] proved that the universal homogeneous triangle free graph has finite big Ramsey degrees and subsequently generalized this result to graphs omitting a clique K_k for any $k \geq 3$ [Dob23]. This was further generalised by Zucker [Zuc22] to free amalgamation classes in binary languages. Hubička applied parameter spaces and Carlson–Simpson’s Theorem to show the finiteness of big Ramsey degrees for partial orders and metric spaces [Hub20a] giving also a straightforward proof of [Dob20a].

The main goal of this note is to demonstrate a proof technique, which allows to reprove some of the aforementioned results in the context of relational structures with higher arities. The proof, for the first time in this area, makes use of the vector (also called product) form of Milliken’s Tree Theorem. To our knowledge this may be also the first combinatorial application of Milliken’s Tree Theorem for trees of unbounded branching [Dod15].

A special case of Theorem 8.1.1 (for colouring vertices) also follows from recent results of Sauer [Sau20] and Coulson, Dobrinen, Patel [CDP20].

8.2 Preliminaries

Our argument will make use of the vector (or product) form of Milliken’s Tree Theorem. All definitions and results in this section are taken from [DK16]. Given an integer ℓ , we use both the combinatorial notion $[\ell] = \{1, \dots, \ell\}$ and the set-theoretical convention $\ell = \{0, 1, \dots, \ell - 1\}$.

A *tree* is a (possibly empty) partially ordered set $(T, <_T)$ such that, for every $t \in T$, the set $\{s \in T : s <_T t\}$ is finite and linearly ordered by $<_T$. All trees considered are finite or countable. All nonempty trees we consider are *rooted*, that

is, they have a unique minimal element called the *root* of the tree. An element $t \in T$ of a tree T is called a *node* of T and its *level*, denoted by $|t|_T$, is the size of the set $\{s \in T : s <_T t\}$. Note that the root has level 0. For $D \subseteq T$, we write $L_T(D) = \{|t|_T : t \in D\}$ for the *level set* of D in T . We use $T(n)$ to denote the set of all nodes of T at level n , and by $T(<n)$ the set $\{t \in T : |t|_T < n\}$. The *height* of T is the minimal natural number h such that $T(h) = \emptyset$. If there is no such number h , then we say that the height of T is ω . We denote the height of T by $h(T)$.

Given a tree T and nodes $s, t \in T$ we say that s is a *successor* of t in T if $t \leq_T s$. The node s is an *immediate successor* of t in T if $t <_T s$ and there is no $s' \in T$ such that $t <_T s' <_T s$. We denote the set of all successors of t in T by $\text{Succ}_T(t)$ and the set of immediate successors of t in T by $\text{ImmSucc}_T(t)$. We say that the tree T is *finitely branching* if $\text{ImmSucc}_T(t)$ is finite for every $t \in T$.

For $s, t \in T$, the *meet* $s \wedge_T t$ of s and t is the largest $s' \in T$ such that $s' \leq_T s$ and $s' \leq_T t$. A node $t \in T$ is *maximal* in T if it has no successors in T . The tree T is *balanced* if it either has infinite height and no maximal nodes, or all its maximal nodes are in $T(h-1)$, where h is the height of T .

A *subtree* of a tree T is a subset T' of T viewed as a tree equipped with the induced partial ordering such that $s \wedge_{T'} t = s \wedge_T t$ for each $s, t \in T'$. Note that our notion of a subtree differs from the standard terminology, since we require the additional condition about preserving meets.

Definition 8.2.1. A subtree S of a tree T is a *strong subtree* of T if either S is empty, or S is nonempty and satisfies the following three conditions.

1. The tree S is rooted and balanced.
2. Every level of S is a subset of some level of T , that is, for every $n < h(S)$ there exists $m \in \omega$ such that $S(n) \subseteq T(m)$.
3. For every non-maximal node $s \in S$ and every $t \in \text{ImmSucc}_T(s)$ the set $\text{ImmSucc}_S(s) \cap \text{Succ}_T(t)$ is a singleton.

Observation 8.2.1. If E is a subtree of a balanced tree T , then there exists a strong subtree $S \supseteq E$ of T such that $L_T(E) = L_T(S)$. \square

A *vector tree* (sometimes also called a *product tree*) is a finite sequence $\mathbf{T} = (T_1, \dots, T_d)$ of trees having the same height $h(T_i)$ for all $i \in [d]$. This common height is the *height* of \mathbf{T} and is denoted by $h(\mathbf{T})$. A vector tree $\mathbf{T} = (T_1, \dots, T_d)$ is *balanced* if the tree T_i is balanced for every $i \in [d]$.

If $\mathbf{T} = (T_1, \dots, T_d)$ is a vector tree, then a *vector subset* of \mathbf{T} is a sequence $\mathbf{D} = (D_1, \dots, D_d)$ such that $D_i \subseteq T_i$ for every $i \in [d]$. We say that \mathbf{D} is *level compatible* if there exists $L \subseteq \omega$ such that $L_{T_i}(D_i) = L$ for every $i \in [d]$. This (unique) set L is denoted by $L_{\mathbf{T}}(\mathbf{D})$ and is called the *level set* of \mathbf{D} in \mathbf{T} .

Definition 8.2.2. Let $\mathbf{T} = (T_1, \dots, T_d)$ be a vector tree. A *vector strong subtree* of \mathbf{T} is a level compatible vector subset $\mathbf{S} = (S_1, \dots, S_d)$ of \mathbf{T} such that S_i is a strong subtree of T_i for every $i \in [d]$.

For every $k \in \omega + 1$ with $k \leq h(\mathbf{T})$, we use $\text{Str}_k(\mathbf{T})$ to denote the set of all vector strong subtrees of \mathbf{T} of height k . We also use $\text{Str}_{\leq k}(\mathbf{T})$ to denote the set of all strong subtrees of \mathbf{T} of height at most k .

Theorem 8.2.2 (Milliken [Mil79]). *For every rooted, balanced and finitely branching vector tree \mathbf{T} of infinite height, every nonnegative integer k and every finite colouring of $\text{Str}_k(\mathbf{T})$ there is $\mathbf{S} \in \text{Str}_\omega(\mathbf{T})$ such that the set $\text{Str}_k(\mathbf{S})$ is monochromatic.*

8.3 Proof of Theorem 8.1.1

Given an integer $n \geq 0$, a $\{0, 1\}$ -vector \vec{v} of length n is a function $\vec{v}: n \rightarrow 2$. We write $|\vec{v}| = n$ to denote the length of \vec{v} and, for $i < |\vec{v}|$, we use v_i to denote the i^{th} coordinate $\vec{v}(i)$ of \vec{v} . In particular, we permit the empty vector and the first coordinate has index 0. We use the standard vocabulary for matrices. An $n \times n$ $\{0, 1\}$ -matrix A is a function $A: n \times n \rightarrow 2$. We use $|A| = n$ to denote the number of rows (and columns) of A . For $k \leq |A|$, we write $A \upharpoonright k$ for the sub-matrix of A with domain $k \times k$. The i^{th} row of the matrix A is the vector $\vec{v}: j \mapsto A_{i,j}$. The value $A(i, j)$, the entry in the i^{th} row and the j^{th} column of A is denoted by $A_{i,j}$. Note that we start the indexing of entries of A from 0. The matrix A is *strictly lower triangular* if $A_{i,j} = 0$ for all i and j with $i \leq j$.

The main idea of the proof of Theorem 8.1.1 is to extend the passing number representation of graphs to 3-uniform hypergraphs. This can be done naturally when one understands the passing number representation in the context of adjacency matrix of a graph as outlined below.

Consider the universal countable homogeneous graph \mathbf{R} (the Rado graph) and enumerate it by fixing its vertex set ω . This yields the *asymmetric adjacency matrix* A of \mathbf{R} . (Recall that this is an infinite $\{0, 1\}$ -matrix with $A_{j,i} = 1$ if and only if $i < j$ and i is adjacent to j in \mathbf{R} .) Assign to every vertex $i \in \omega$ a $\{0, 1\}$ -word $w(i)$ which corresponds to the strictly sub-diagonal part of the i^{th} row of A . It follows that, for all $i, j \in \omega$ with $i < j$, we have $|w(i)| = i$ and $w(j)_i = w(j)_{|w(i)|} = 1$ if and only if i is adjacent to j in \mathbf{R} . This exactly corresponds to the passing number representation used to show that big Ramsey degrees of \mathbf{R} are finite [Tod10, Theorem 6.25], [Sau06, LSV06].

Now consider the countable homogeneous 3-uniform hypergraph \mathbf{H} and put $H = \omega$. Proceeding analogously as before, one can consider the *asymmetric adjacency tensor* of \mathbf{H} which is a function $A': n \times n \times n \rightarrow 2$ defined by $A'(k, j, i) = A'_{k,j,i} = 1$ if and only if $i < j < k$ and i, j, k forms an hyper-edge of \mathbf{H} . Now assign to every vertex $i \in \omega$ an $i \times i$ matrix $M(i)$ such that $i < j < k$ forms a hyper-edge of \mathbf{H} if and only if $M(k)_{j,i} = 1$. This matrix can again be seen as the sub-diagonal part of a “slice” of the adjacency tensor A' since $M(k)_{j,i} = A'_{j,k,i}$ for every $i, j < k$.

Our proof of Theorem 8.1.1 is based on a refinement of this matrix representation. However, in contrast to binary structures, we need to solve one additional difficulty. The tree of matrices (see T_2 in Definition 8.3.1 and Figure 8.1) is no longer uniformly branching and there is no bound on the number of nodes of a strong subtree of a given height. This is the main motivation for using a vector tree we define now.

Definition 8.3.1. Let $\mathbf{T} = (T_1, T_2)$ be the vector tree, where:

1. The binary tree $(T_1, <_{T_1})$ consists of all finite $\{0, 1\}$ -vectors ordered by the end-extension. More precisely, we have $\vec{u} \leq_{T_1} \vec{v}$ if $|\vec{u}| \leq |\vec{v}|$ and $u_i = v_i$ for




Figure 8.1: First 4 levels of the tree T_2 .

every $i \in |u|$. The root of T_1 is the empty vector.

- Nodes of the tree T_2 are all finite strictly lower triangular (square) $\{0, 1\}$ -matrices ordered by extension. That is, we have $A \leq_{T_2} B$ if and only if $|A| \leq |B|$ and $A_{i,j} = B_{i,j}$ for every $i, j \in |A|$. The root of T_2 is the empty matrix; see Figure 8.1.

Remark 8.3.1. The tree T_2 corresponds to the tree of 1-types of \mathbf{H} (see for example [CDP20]). The tree T_1 in our construction has a natural meaning too: while the tree T_2 represents vertices and hyper-edges, the tree T_1 represents the union of all graphs that are created from a 3-uniform hypergraph by fixing a vertex v and considering the graph on the same vertex set with edges induced by hyper-edges containing v .

A key element of the proof is a correspondence between strong vector subtrees of the vector tree \mathbf{T} and special subtrees of \mathbf{T}_1 with shape isomorphic to the initial segments of \mathbf{T}_1 .

For a matrix A and a vector \vec{v} with $|A| = |\vec{v}| = n$, the *extension* $A \hat{\ } \vec{v}$ of A is the $(n + 1) \times (n + 1)$ matrix

$$\left(\begin{array}{c|c} A & \begin{array}{c} 0 \\ \vdots \\ 0 \end{array} \\ \hline \vec{v} & 0 \end{array} \right).$$

More precisely, the matrix $A \hat{\ } \vec{v}$ is given by setting

- $(A \hat{\ } \vec{v})_{i,j} = A_{i,j}$ for all $i, j \in n$,
- $(A \hat{\ } \vec{v})_{n,i} = v_i$ for every $i \in n$, and
- $(A \hat{\ } \vec{v})_{j,n} = 0$ for every $j \in n + 1$.

Note that if A is a strictly lower triangular matrix, then $A \hat{\ } \vec{v}$ is strictly lower triangular as well.

Definition 8.3.2. Let (S_1, S_2) be a strong vector subtree of \mathbf{T} of height $k \in \omega + 1$. In other words, $(S_1, S_2) \in \text{Str}_k(\mathbf{T})$. The *valuation tree* $\text{val}(S_1, S_2)$ corresponding to (S_1, S_2) is a subset of S_2 defined by the following recursive rules:




Figure 8.2: A valuation tree (right) constructed from a vector strong subtree (left).

1. The root of $\text{val}(S_1, S_2)$ is the root of S_2 .
2. If $A \in \text{val}(S_1, S_2)$, $\vec{v} \in S_1(|A|_{S_2})$, and $\{C\} = \text{ImmSucc}_{S_2}(A) \cap \text{Succ}_{T_2}(A \wedge \vec{v})$, then $C \in \text{val}(S_1, S_2)$.
3. There are no other nodes in $\text{val}(S_1, S_2)$.

Note that $\text{val}(S_1, S_2)$ is a subtree of S_2 and hence also a subtree of T_2 . Also, the height of $\text{val}(S_1, S_2)$ equals k and the number of nodes of $\text{val}(S_1, S_2)$ depends only on k , see Lemma 8.3.1.

A tree $T \subseteq T_2$ is a *valuation tree* if $T = \text{val}(S_1, S_2)$ for some $(S_1, S_2) \in \text{Str}_{\leq \omega}(\mathbf{T})$.

Example 8.3.1. See Figure 8.2 for an example of a valuation tree constructed from a product strong subtree.

For two subtrees T and T' of T_2 , a function $f: T \rightarrow T'$ is a *structural isomorphism* if it is an isomorphism of trees (preserving relative heights of nodes), and for every $A, B, C \in T$ with $|A| \leq |B| < |C|$ it also holds that $f(C)|_{|f(B)|, |f(A)|} = C|_{|B|, |A|}$.

Lemma 8.3.1. *For every $k \in \omega + 1$ and every valuation subtree T of T_2 of height k , there exists a unique structural isomorphism $f: T_2(<k) \rightarrow T$.*

Proof. For $k \in \omega$ we use induction on k . There is nothing to prove for $k = 0$. If $k = 1$, then the function f mapping the empty matrix to the root of T is the unique structural isomorphism. Assume the induction hypothesis does hold for $k > 0$. Let $T = \text{val}(S_1, S_2)$ be of height $k + 1$. By the induction hypothesis for the valuation tree $T(<k)$ there exists a unique structural isomorphism $f: T_2(<k) \rightarrow T(<k)$. Denote by $e: k + 1 \rightarrow \omega$ the increasing enumeration of $L_{T_2}(T)$. Fix a node $\vec{u} \in T_1(k - 1)$. Then there is a unique $\vec{v} \in S_1(k - 1)$ such that $u_i = v_{e(i)}$ for every $i \in k$. For every $A \in T_2(k - 1)$, there is a unique $C \in \text{ImmSucc}_{S_2}(f(A)) \cap \text{Succ}_{T_2}(f(A) \wedge \vec{v})$. We extend the map f by declaring $f: A \wedge \vec{u} \mapsto C$, and we do this for each choice of \vec{u} and A . It is easy to check that the extended map is a structural isomorphism of $T_2(<k + 1)$ and T , and that the extension was in fact defined in the unique possible way.

If $k = \omega$, then by the induction hypothesis there are structural isomorphisms $f_i: T_2(<i) \rightarrow T(<i)$ for each $i \in \omega$. Since these isomorphisms are unique, we get $f_i \subset f_j$ for $i < j$, and $f = \bigcup \{f_i : i \in \omega\}$ is the desired structural isomorphism. On the other hand, if $g: T_2 \rightarrow T$ is a structural isomorphism, then for each $i \in \omega$ the restriction $g \upharpoonright T_2(<i) \rightarrow T(<i)$ is a structural isomorphism and due to the induction hypotheses has to be equal to f_i , consequently $g = f$. \square

Definition 8.3.3. Let \mathbf{G} be a 3-uniform hypergraph defined by the following two rules:

1. The vertex set of \mathbf{G} consists of all nodes of T_2 . In particular, the vertices of \mathbf{G} are square $\{0, 1\}$ -matrices.
2. There is a hyperedge $\{A, B, C\}$ in \mathbf{G} if and only if $A, B, C \in T_2$ are matrices satisfying $|A| < |B| < |C|$ and $C_{|B|,|A|} = 1$.

Recall that \mathbf{H} denotes the universal countable homogeneous 3-uniform hypergraph. Without loss of generality, we assume that the vertex set of \mathbf{H} is ω . Let $\varphi: \mathbf{H} \rightarrow \mathbf{G}$ be an embedding defined by setting $\varphi(i) = A^i$. Here, A^i is a $(2i+1) \times (2i+1)$ matrix such that if $\{j, k, i\}$ is a hyper-edge of \mathbf{H} with $j < k < i$, then $A^i_{2k+1, 2j} = A^i_{2k+1, 2j+1} = 1$ and there are no other non-zero values in A^i ; see Example 8.3.2.

Example 8.3.2. Assume that \mathbf{H} starts with vertices $\{0, 1, 2, 3\}$ with hyper-edges $\{0, 1, 2\}$, $\{0, 1, 3\}$ and $\{1, 2, 3\}$. Then the corresponding images in \mathbf{G} are:

$$\varphi(0) = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \varphi(1) = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \varphi(2) = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \varphi(3) = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

It is easy to check that φ is indeed a hypergraph embedding. We also have the following simple observation.

Observation 8.3.2. For any two matrices $A, B \in \varphi[\mathbf{H}]$,

1. all even rows are constant 0-vectors and so $|A \wedge_{T_2} B|$ is odd, and
2. if $\vec{v} \neq \vec{u}$ are two rows of A and B , respectively, then $|\vec{v} \wedge_{T_1} \vec{u}|$ is even. \square

Now, we prove the last auxiliary result that we use in the proof of Theorem 8.1.1. This is a standard step of constructing the envelope of a set as used by Laver and Milliken [Tod10, Section 6.2]. Here we additionally need to take care of the interactions between the two trees.

Lemma 8.3.3. For every $k \in \omega$, there exists $R(k) \in \omega$ such that, for every set $S \subset \omega$ of size k , there exists a valuation tree of height at most $R(k)$ containing all vertices of $\varphi[S]$.

Proof. Choose an arbitrary natural number k , we will show that there is a number $R(k)$ with the desired property. To do so, let S be a set of k elements from ω . We will construct the required strong vector subtree (S_1, S_2) of \mathbf{T} such that $\text{val}(S_1, S_2)$ contains all matrices from $\varphi[S]$. The construction will take a determined number of steps and the upper bound on the height of the constructed tree (S_1, S_2) will thus be a function of k .

To achieve this, we define envelopes E_1 and E_2 in the trees T_1 and T_2 , respectively, by first collecting all necessary matrices in E_2 , then inserting all necessary

vectors into E_1 , which in turn requires adding more matrices to E_2 . The important upshot of our construction is that we can argue that this process promptly terminates and the resulting envelopes are bounded in size.

We proceed in four steps, first defining auxiliary sets $E_1^0 \subseteq T_1$ and $E_2^0 \subseteq T_2$ that will be further extended to E_1 and E_2 , respectively.

- (i) Let $E_2^0 = \{A \wedge_{T_2} B : A, B \in \varphi[S]\} \subset T_2$. This is necessary to obtain a subtree of T_2 . Observe that $|E_2^0| \leq 2k - 1$.
- (ii) Let $E_1^0 \subset T_1$ consist of all $|B|^{\text{th}}$ rows of A for all $A, B \in E_2^0$, $|B| < |A|$ and a constant 0-vector of length $\max L_{T_2}(E_2^0)$. This is necessary to obtain a valuation tree that contains all of E_2^0 . The additional zero vector is added to make the level sets of E_1^0 and E_2^0 equal. Observe that $|E_1^0| \leq |E_2^0|^2 + 1$.
- (iii) Let $E_1 = \{\vec{u} \wedge_{T_1} \vec{v} : \vec{u}, \vec{v} \in E_1^0\}$. This is necessary to obtain a subtree of T_1 . Observe that $|E_1| \leq 2|E_1^0| - 1$.
- (iv) Let E_2 extend E_2^0 by all matrices $A \upharpoonright |\vec{v}|$ where $A \in E_2^0$ and $\vec{v} \in E_1^0$. This is necessary in order to synchronize levels between both subtrees. Observe that $|E_2| \leq |E_2^0|(|E_1| + 1)$.

It follows from step (iii) that E_1 is meet closed in T_1 and thus it is a subtree of T_1 . Similarly, step (i) implies that E_2^0 is a subtree of T_2 . Thus also E_2 is a subtree of T_2 , as we did not introduce any new meets in step (iv). It follows from the upper bounds on $|E_1|$ and $|E_2|$ that the height of E_2 is bounded from above by a function of k .

By step (iv), the level sets of E_1 and E_2 are the same, that is, $L = L_{T_2}(E_2) = L_{T_1}(E_1)$. Now, let S_1 be some strong subtree of T_1 containing E_1 such that $L_{T_1}(S_1) = L$ and let S_2 be some strong subtree of T_2 containing E_2 such that $L_{T_2}(S_2) = L$. Such trees S_1 and S_2 exist by Observation 8.2.1.

We claim that $\text{val}(S_1, S_2)$ contains $\varphi[S]$. Choose any matrix $A \in \varphi[S] \subseteq E_2 \subseteq S_2$. We prove by induction on the level $\ell \in L$, where $\ell \leq |A|$, that $A \upharpoonright \ell \in \text{val}(S_1, S_2)$. For the base of the induction, if ℓ is the minimal element of L , then $A \upharpoonright \ell$ is the root of S_2 and hence the root of $\text{val}(S_1, S_2)$.

To prove the induction step, we need to check that if $A \upharpoonright \ell \in S_2$ for $\ell < |A|$, then the ℓ^{th} row \vec{v} of A , is a node of S_1 . Observe that level sets of the constructed sets “ E ” are extended only at steps (i) and (iii) of the construction. Moreover, all new levels introduced during step (i) are odd by part (1) of Observation 8.3.2 while levels introduced at step (iii) are even by part (2) of Observation 8.3.2.

We distinguish two cases based on the parity of the level ℓ . If ℓ is odd, then $\ell \in L_{T_2}(E_2^0)$, since all odd levels are introduced only in step (i). By step (ii), we then have $\vec{v} \in E_1^0$. Since $E_1^0 \subseteq E_1 \subseteq S_1$, we have $\vec{v} \in S_1$.

Otherwise ℓ is even. Then \vec{v} is a constant 0-vector by part (1) of Observation 8.3.2. We have $\vec{v} \in S_1$, since one of the maximal nodes of S_1 is a constant 0-vector by step (ii).

Altogether, $\text{val}(S_1, S_2)$ contains all matrices from $\varphi[S]$, which finishes the proof. \square

We can now proceed with the proof of Theorem 8.1.1.

Proof of Theorem 8.1.1. Fix a finite 3-uniform hypergraph \mathbf{A} . Recall that we want to prove that there exists a number $\ell = \ell(\mathbf{A})$ such that for every finite k

$$\mathbf{H} \longrightarrow (\mathbf{H})_{k,\ell}^{\mathbf{A}}.$$

That is, for every colouring $\chi^0: \binom{\mathbf{H}}{\mathbf{A}} \rightarrow k$ there is an embedding $g: \mathbf{H} \rightarrow \mathbf{H}$ such that χ^0 does not take more than ℓ values on $\binom{g[\mathbf{H}]}{\mathbf{A}}$.

Consider the 3-uniform hypergraph \mathbf{G} introduced in Definition 8.3.3. Since \mathbf{G} is a countable 3-uniform hypergraph, it follows from the properties of \mathbf{H} that there is an embedding $\theta: \mathbf{G} \rightarrow \mathbf{H}$. Consider the colouring $\chi: \binom{\mathbf{G}}{\mathbf{A}} \rightarrow k$ obtained by setting $\chi(\tilde{\mathbf{A}}) = \chi^0(\theta(\tilde{\mathbf{A}}))$ for every $\tilde{\mathbf{A}} \in \binom{\mathbf{G}}{\mathbf{A}}$.

Consider the vector tree $\mathbf{T} = (T_1, T_2)$ given by Definition 8.2.2. Let $h = R(|\mathbf{A}|)$ be given by Lemma 8.3.3. Let \mathbf{G}_h be the induced sub-hypergraph of \mathbf{G} on $T_2(<h)$. We enumerate the copies of \mathbf{A} in $\binom{\mathbf{G}_h}{\mathbf{A}}$ as $\{\tilde{\mathbf{A}}_i : i \in \ell\}$ for some $\ell \in \omega$ which will give the upper bound on the big Ramsey degree of \mathbf{A} .

By Lemma 8.3.1, for every valuation tree T of height h , there is a structural isomorphism $f_T: \mathbf{G}_h \rightarrow T$ that is also an isomorphism of the corresponding sub-hypergraphs of \mathbf{G} . Let $\mathbf{S} = (S_1, S_2)$ be a strong subtree of \mathbf{T} of height h and consider the structural isomorphism $f = f_{\text{val}(S_1, S_2)}: \mathbf{G}_h \rightarrow \text{val}(S_1, S_2)$. Put

$$\bar{\chi}(\mathbf{S}) = \langle \chi(f(\tilde{\mathbf{A}}_i)) : i \in n \rangle,$$

which is a finite colouring of $\text{Str}_h(\mathbf{T})$. By Theorem 8.2.2, there is an infinite strong subtree of \mathbf{T} monochromatic with respect to $\bar{\chi}$. Let U be its corresponding valuation subtree. The structural isomorphism $\psi: T_2 \rightarrow U$ given by Lemma 8.3.1 is a hypergraph embedding $\psi: \mathbf{G} \rightarrow \mathbf{G}$. Since, by Lemma 8.3.3, every $\tilde{\mathbf{A}} \in \binom{\mathbf{G}}{\mathbf{A}}$ is contained in at least one valuation subtree of height h , we know that χ takes at most ℓ different values on $\binom{\psi[\mathbf{G}]}{\mathbf{A}}$. Considering the embedding $\varphi: \mathbf{H} \rightarrow \mathbf{G}$ defined earlier, the image $\theta[\psi[\varphi[\mathbf{H}]]]$ is the desired copy $g[\mathbf{H}]$ of \mathbf{H} , in which copies of \mathbf{A} have at most n different colours in χ^0 . \square

8.4 Concluding remarks

1. The construction naturally generalises to d -uniform hypergraphs for any $d \geq 2$. Identifying the underlying set of the d -uniform hypergraph with ω we get a d -dimensional adjacency $\{0, 1\}$ -tensor. We can now consider the hypergraph consisting of $(d-1)$ -dimensional ‘sub-diagonal’ $\{0, 1\}$ -tensors with the edge relation being defined analogously as in the 3-uniform case. The $(d-1)$ -dimensional tensors ordered by extension now form a tree T_{d-1} . Sub-hypergraphs isomorphic to T_{d-1} will be again constructed using Milliken’s Tree Theorem used for the vector tree $\mathbf{T}_d = (T_1, \dots, T_{d-1})$ and by defining valuation subtrees of T_{d-1} . Nodes of a tree T_i with $1 \leq i \leq d-1$ are $\{0, 1\}$ -tensors of order i ordered analogously as in the tree T_2 used in Section 8.3. The definition of the valuation tree from Section 8.3 also naturally generalises; given a vector strong subtree $\mathbf{S} = (S_1, \dots, S_{d-1})$ one first obtains the valuation tree $\text{val}(S_1, S_2)$. Based on $\text{val}(S_1, S_2)$ and S_3 the valuation subtree $\text{val}(S_1, S_2, S_3)$ of T_3 can be constructed in analogy to Definition 8.3.2. The construction then proceeds similarly for

higher orders, defining $\text{val}(S_1, S_2, \dots, S_i)$ for all $i < d$. The final valuation tree $\text{val}(S_1, S_2, \dots, S_{d-1})$ is the desired subtree of T_{d-1} . A detailed description of these constructions is going to appear in full generality in [BCH⁺20].

2. More generally, structures in a finite relational language with symbols of maximum arity d can be represented by vector trees $\mathbf{T}_d = (T_1, \dots, T_{d-1})$. In this case, the nodes of a tree T_i with $1 \leq i \leq d-1$ are sequences of tensors of order $i - d + a$ for every relational symbol of arity $a > i - d$.

3. We aimed for simplicity in our proof of Theorem 8.1.1. The bounds obtained in the proof are not optimal. Structures defined in [BCH⁺19] can be used to produce a more careful embedding of hypergraphs to \mathbf{H} . They describe the order in which the branchings of the trees T_1 and T_2 and of the actual vertices appear. This is also going to appear in [BCH⁺20].

8.5 Acknowledgment

We would like to thank to Stevo Todorćević for his kind remarks and for helpful discussion regarding the history and context of this area. We are also grateful to three anonymous referees whose suggestions improved the presentation of this paper.

9. Big Ramsey degrees and infinite languages

SAMUEL BRAUNFELD, DAVID CHODOUNSKÝ, NOÉ DE RANCOURT, JAN HUBIČKA, JAMAL KAWACH, MATĚJ KONEČNÝ

Abstract This paper investigates big Ramsey degrees of unrestricted relational structures in (possibly) infinite languages. While significant progress has been made in studying big Ramsey degrees, many classes of structures with finite small Ramsey degrees still lack an understanding of their big Ramsey degrees. We show that if there are only finitely many relations of every arity greater than one, then unrestricted relational structures have finite big Ramsey degrees, and give some evidence that this is tight. This is the first time that finiteness of big Ramsey degrees has been established for an infinite-language random structure. Our results represent an important step towards a better understanding of big Ramsey degrees for structures with relations of arity greater than two.

9.1 Introduction

Given L -structures \mathbf{A} , \mathbf{B} and \mathbf{C} , we write $\mathbf{C} \longrightarrow (\mathbf{B})_{k,\ell}^{\mathbf{A}}$ to denote the following statement:

For every colouring $\chi: \binom{\mathbf{C}}{\mathbf{A}} \rightarrow k$ there exists an embedding $f: \mathbf{B} \rightarrow \mathbf{C}$ such that χ takes at most ℓ values on $\binom{f[\mathbf{B}]}{\mathbf{A}}$.

For a countably infinite structure \mathbf{B} and its finite induced substructure \mathbf{A} , the *big Ramsey degree* of \mathbf{A} in \mathbf{B} is the least number $\ell \in \omega + 1$ such that $\mathbf{B} \longrightarrow (\mathbf{B})_{k,\ell}^{\mathbf{A}}$ for every $k \in \omega$; see [KPT05]. A countably infinite structure \mathbf{B} has *finite big Ramsey degrees* if the big Ramsey degree of \mathbf{A} in \mathbf{B} is finite for every finite substructure \mathbf{A} of \mathbf{B} .

The study of big Ramsey degrees dates back to Ramsey's theorem itself which can be stated as $(\omega) \longrightarrow (\omega)_{k,1}^n$ for every $n, k \in \omega$, where we understand the ordinals ω and n as structures with their natural linear orders (using the standard set-theoretic convention that $n = \{0, 1, \dots, n-1\}$ and $\omega = \{0, 1, \dots\}$). However, the real origin of this project lies in the work of Galvin [Gal68, Gal69] who proved that the big Ramsey degree of pairs in the order of the rationals, denoted by (\mathbb{Q}, \leq) , is equal to 2. Subsequently, Laver in late 1969 proved that (\mathbb{Q}, \leq) has in fact finite big Ramsey degrees, see [EH74, Lav84], and Devlin determined the exact values of ℓ [Dev79, Page 73]. In 2006 Sauer [Sau06] proved that the Rado graph has finite big Ramsey degrees and Laflamme, Sauer, Vuksanovic [LSV06] obtained their exact values. Behind both this result and the result for (\mathbb{Q}, \leq) was Milliken's tree theorem for a single binary tree. This was refined to unconstrained structures in binary languages [LSV06] and additional special classes [DLS16, NVT09, LNVTS10]. Milliken's tree theorem remained the key in all results in the area (here we consider Ramsey's Theorem as a special case of Milliken's tree theorem for the unary tree). See also [Dob20c] for a recent survey.

Balko, Chodounský, Hubička, Konečný, Nešetřil, and Vena recently applied the product version of Milliken’s tree theorem to prove that the generic countable 3-uniform hypergraph has finite big Ramsey degrees. [BCH⁺22, BCH⁺19] The method could be extended to prove big Ramsey degrees of the generic countable k -uniform hypergraph for an arbitrary k , and in this paper we further extend these results and prove the following theorem (the definition of an unrestricted structure is given later, see Definition 9.2.1):

Theorem 9.1.1. *Let L be a relational language with finitely many relations of every arity greater than one and with finitely or countably many unary relations and let \mathbf{H} be an unrestricted L -structure where all relations are injective. Then \mathbf{H} has finite big Ramsey degrees.*

We believe that this result is the limit of how far Milliken’s tree theorem can be pushed in this area (at least by using the *passing number representation* and its generalisations). In Section 9.6 we give evidence for this and discuss infinite lower bounds. In fact, we conjecture that this theorem is tight when there are only finitely many unary relations (see Section 9.7.1 and Conjecture 9.7.8).

Besides Milliken’s tree theorem, other partition theorems have been used in the area, such as the Carlson–Simpson theorem [Hub20a, BCD⁺21a, BCH⁺21a] or various custom theorem proved using forcing [CDP22b, Dob20a, Dob23, Zuc22]. While we only use Milliken’s theorem here, the product tree structure we develop is, to a large degree, inherent to the problem, not to the method. For this reason we believe that our development of the concept of valuation trees (extending [BCH⁺22]) and k -enveloping embeddings will serve as an important basis for future big Ramsey theorems for structures with relations of arity greater than two.

9.2 Preliminaries

A *relational language* L is a collection of symbols, each having an associated *arity*, denoted by $a(R) \in \omega$. An L -*structure* \mathbf{A} consists of a vertex set A and an *interpretation* of every $R \in L$, which is $R_{\mathbf{A}} \subseteq A^{a(R)}$. We say that a relation R is *injective* in \mathbf{A} if every tuple $\bar{x} \in R_{\mathbf{A}}$ is injective (i.e. contains no repeated occurrences of vertices) and it is *symmetric* if whenever $\bar{x} \in R_{\mathbf{A}}$ and \bar{y} is a permutation of \bar{x} then $\bar{y} \in R_{\mathbf{A}}$. Equivalently, we can consider a symmetric relation to be a subset of $\binom{A}{a(R)}$. An L -*hypergraph* is an L -structure where all relations are injective and symmetric and every tuple is in at most one relation (so it can be seen as an edge-colored hypergraph with number of colours for each arity given by L).

We adopt the standard model-theoretic notions of embeddings etc. with one exception: Unless explicitly stated otherwise, every structure in this paper will be implicitly equipped with an enumeration (i.e. a linear order for finite structures and an ω -type order for countably infinite structures) and all embeddings will be monotone with respect to the enumerations. When we do not explicitly describe the enumeration, one can pick an arbitrary one.

Since all structures will be at most countable, we can assume without loss of generality that the vertex set of every structure is either some natural number

n or ω , the symbol \leq will always denote the standard order of natural numbers and all embeddings will be monotone with respect to \leq . We will always assume that $\leq \notin L$. If \mathbf{A} and \mathbf{B} are L -structures, the symbol $\binom{\mathbf{B}}{\mathbf{A}}$ denotes the set of all embeddings $\mathbf{A} \rightarrow \mathbf{B}$. (Remember that these are monotone with respect to \leq .)

Given a class of finite and countably infinite structures \mathcal{C} , a structure $\mathbf{A} \in \mathcal{C}$ is *universal* if for every $\mathbf{B} \in \mathcal{C}$ there exists an embedding $\mathbf{B} \rightarrow \mathbf{A}$. Examples of universal structures come, for example, from Fraïssé theory which produces special (homogeneous) unenumerated structures which retain their universality (for all enumerations of members of the respective class) even when enumerated. Often we will say that \mathbf{A} is a universal structure of some kind (e.g. \mathbf{A} is a universal \mathcal{F} -free structure). This will mean that \mathbf{A} is a structure of the given kind and is universal for all countable structures of that kind.

Definition 9.2.1. An L -structure \mathbf{F} is *covered by a relation* if there is some relation $R \in L$ and a tuple \bar{x} containing all vertices of \mathbf{F} such that $\bar{x} \in R_{\mathbf{F}}$. If \mathcal{F} is a collection of L -structures and \mathbf{A} is an L -structure such that there is no $\mathbf{F} \in \mathcal{F}$ with an embedding $\mathbf{F} \rightarrow \mathbf{A}$, we say that \mathbf{A} is *\mathcal{F} -free*. An L -structure is *unrestricted* if there is a family \mathcal{F} containing only finite L -structures which are covered by a relation such that \mathbf{A} is a universal \mathcal{F} -free structure.

Note that, in particular, every unrestricted structure is countable (by our definition of universality).

9.2.1 Milliken's tree theorem

Our argument will make use of the vector (or product) form of Milliken's tree theorem. All definitions and results in this section are taken from [DK16]. Given an integer ℓ , we use both the combinatorial notion $[\ell] = \{1, \dots, \ell\}$ and the set-theoretical convention $\ell = \{0, 1, \dots, \ell - 1\}$.

A *tree* is a (possibly empty) partially ordered set $(T, <_T)$ such that, for every $t \in T$, the set $\{s \in T : s <_T t\}$ is finite and linearly ordered by $<_T$. All trees considered in this paper will be finite or countable. All nonempty trees we consider are *rooted*, that is, they have a unique minimal element called the *root* of the tree. An element $t \in T$ of a tree T is called a *node* of T and its *level*, denoted by $|t|_T$, is the size of the set $\{s \in T : s <_T t\}$. Note that the root has level 0. For $D \subseteq T$, we write $L_T(D) = \{|t|_T : t \in D\}$ for the *level set* of D in T . We use $T(n)$ to denote the set of all nodes of T at level n , and by $T(<n)$ the set $\{t \in T : |t|_T < n\}$. The *height* of T is the smallest natural number h such that $T(h) = \emptyset$. If there is no such number h , then we say that the height of T is ω . We denote the height of T by $h(T)$.

Given a tree T and nodes $s, t \in T$ we say that s is a *successor* of t in T if $t \leq_T s$. The node s is an *immediate successor* of t in T if $t <_T s$ and there is no $s' \in T$ such that $t <_T s' <_T s$. We denote the set of all successors of t in T by $\text{Succ}_T(t)$ and the set of immediate successors of t in T by $\text{ImmSucc}_T(t)$. We say that the tree T is *finitely branching* if $\text{ImmSucc}_T(t)$ is finite for every $t \in T$.

For $s, t \in T$, the *meet* of s and t , denoted by $s \wedge_T t$, is the largest $s' \in T$ such that $s' \leq_T s$ and $s' \leq_T t$. A node $t \in T$ is *maximal* in T if it has no successors in T . The tree T is *balanced* if it either has infinite height and no maximal nodes, or all its maximal nodes are in $T(h - 1)$, where h is the height of T .

A *subtree* of a tree T is a subset T' of T viewed as a tree equipped with the induced partial ordering such that $s \wedge_{T'} t = s \wedge_T t$ for each $s, t \in T'$. Note that our notion of a subtree differs from the standard terminology, since we require the additional condition about preserving meets.

Definition 9.2.2. A subtree S of a tree T is a *strong subtree* of T if either S is empty, or S is nonempty and satisfies the following three conditions.

1. The tree S is rooted and balanced.
2. Every level of S is a subset of some level of T , that is, for every $n < h(S)$ there exists $m \in \omega$ such that $S(n) \subseteq T(m)$.
3. For every non-maximal node $s \in S$ and every $t \in \text{ImmSucc}_T(s)$ the set $\text{ImmSucc}_S(s) \cap \text{Succ}_T(t)$ is a singleton.

Observation 9.2.1. *If E is a subtree of a balanced tree T , then there exists a strong subtree $S \supseteq E$ of T such that $L_T(E) = L_T(S)$. \square*

A *vector tree* of dimension $d \in \omega + 1$ (often also called *product tree*) is a sequence $\mathbf{T} = (T_i : i \in d)$ of trees having the same height $h(T_i)$ for all $i \in d$. This common height is the *height* of \mathbf{T} and is denoted by $h(\mathbf{T})$. A vector tree $\mathbf{T} = (T_i : i \in d)$ is *balanced* if the tree T_i is balanced for every $i \in d$.

If $\mathbf{T} = (T_i : i \in d)$ is a vector tree, then a *vector subset* of \mathbf{T} is a sequence $\mathbf{D} = (D_i : i \in d)$ such that $D_i \subseteq T_i$ for every $i \in d$. We say that \mathbf{D} is *level compatible* if there exists $L \subseteq \omega$ such that $L_{T_i}(D_i) = L$ for every $i \in d$. This (unique) set L is denoted by $L_{\mathbf{T}}(\mathbf{D})$ and is called the *level set* of \mathbf{D} in \mathbf{T} . If $k \in d$, we denote $\mathbf{T} \upharpoonright_k = (T_i : i \in k)$ and $\mathbf{T}^{(k)} = (T_i : k \leq i < d)$.

Definition 9.2.3. Let $\mathbf{T} = (T_i : i \in d)$ be a vector tree. A *strong vector subtree* (or *strong product subtree*) of \mathbf{T} is a level compatible vector subset $\mathbf{S} = (S_i : i \in d)$ of \mathbf{T} such that S_i is a strong subtree of T_i for every $i \in d$.

For every $k \in \omega + 1$ with $k \leq h(\mathbf{T})$, we use $\text{Str}_k(\mathbf{T})$ to denote the set of all strong vector subtrees of \mathbf{T} of height k . We also use $\text{Str}_{\leq k}(\mathbf{T})$ to denote the set of all strong subtrees of \mathbf{T} of height at most k .

Theorem 9.2.2 (Milliken [Mil79]). *For every rooted, balanced and finitely branching vector tree \mathbf{T} of infinite height and finite dimension, every non-negative integer k and every finite colouring of $\text{Str}_k(\mathbf{T})$ there is $\mathbf{S} \in \text{Str}_\omega(\mathbf{T})$ such that the set $\text{Str}_k(\mathbf{S})$ is monochromatic.*

9.3 Valuation trees

Given $n \in \omega + 1$ and $0 < \ell < \omega$, we denote $I_\ell^n = \{(i_0, \dots, i_{\ell-1}) : n > i_0 > \dots > i_{\ell-1} \geq 0\}$. If $n, \ell \in \omega + 1$, we put $I_{<\ell}^n = \bigcup_{1 \leq k < \ell} I_k^n$. Let $\sigma = (\sigma_1, \sigma_2, \dots)$ be an infinite sequence of positive natural numbers, we will call it a *signature*. Let $f : I_{<\omega}^n \rightarrow \omega$ be a function such that if $\bar{x} \in I_\ell^n$ then $f(\bar{x}) < \sigma_\ell$. We call such f a *valuation function of level n and signature σ* (or a σ -*valuation function of level n* , or just a *valuation function of level n* if σ is clear from the context), and write $|f| = n$. Given $i \in \omega$ we denote by $\sigma^{(i)}$ the *i -shift* of σ defined by $\sigma_j^{(i)} = \sigma_{j+i}$.

Given a signature σ , let $\mathbf{T}^\sigma = (T_0, T_1, \dots)$ be the infinite-dimensional vector tree where T_i consists of all $\sigma^{(i)}$ -valuation functions of a finite level ordered by inclusion. Often σ will be implicit from the context and we will write only \mathbf{T} for \mathbf{T}^σ . Note that if σ is constant 1 from some point on, \mathbf{T}^σ is essentially finite-dimensional (in the sense that from some point on, each T_i is just the chain of constant zero valuation functions).

A key element of our construction is a correspondence between strong vector subtrees of the vector tree \mathbf{T} and special subtrees of T_0 with shape isomorphic to initial segments of T_0 . If f is a $\sigma^{(i)}$ -valuation function of level n , g is a $\sigma^{(i+1)}$ -valuation function of level n and h is a $\sigma^{(i)}$ -valuation function of level $n+1$, we say that h is an *extension* of f by g and write $h \in f \hat{\wedge} g$ if

$$h(x_0, \dots, x_d) = \begin{cases} f(x_0, \dots, x_d) & \text{if } x_0 < n, \\ g(x_1, \dots, x_d) & \text{if } x_0 = n \text{ and } d > 0. \end{cases}$$

Note that there are always $\sigma_1^{(i)}$ -many possible extensions of f by g which differ on their value at the singleton n , and we will use $f \hat{\wedge} g$ to denote this set.

Definition 9.3.1. Let σ be an arbitrary signature, consider $\mathbf{T} = \mathbf{T}^\sigma$, let $k \in \omega+1$ and let $(S_i : i \in k)$ be a strong vector subtree of $\mathbf{T} \upharpoonright_k$ of height at least k . We now define by induction on k a subset of S_0 which we call the *valuation tree* $\text{val}(S_i : i \in k)$:

If $k = 0$ then $\text{val}() = \emptyset$. For $0 < k < \omega$, we put $S' = \text{val}(S_1, \dots, S_{k-1})$ using the induction hypothesis and define $\text{val}(S_i : i \in k)$ by the following recursive rules:

1. The root of $\text{val}(S_i : i \in k)$ is the root of S_0 .
2. If $f \in \text{val}(S_i : i \in k)$, $g \in S'(|f|_{S_0})$, then $\text{ImmSucc}_{S_0}(f) \cap \text{Succ}_{T_0}(f \hat{\wedge} g) \subseteq \text{val}(S_i : i \in k)$ (note that this set has size $\sigma_1^{(i)}$).
3. There are no other nodes in $\text{val}(S_i : i \in k)$.

If $k = \omega$ then we put $\text{val}(S_i : i \in k) = \bigcup_{j \in \omega} \text{val}(S_i : i \in j)$.

A tree $T \subseteq T_0$ is a *valuation tree* if $T = \text{val}(S_i : i \in k)$ for some $(S_i : i \in k) \in \text{Str}_{\leq \omega}(\mathbf{T})$.

Note that $\text{val}(S_i : i \in k)$ is a subtree of S_0 and hence also a subtree of T_0 . Also, the height of $\text{val}(S_i : i \in k)$ equals k and the number of nodes of $\text{val}(S_i : i \in k)$ depends only on k and σ , see Lemma 9.3.1.

Note that if $(S_i : i \in k)$ is a strong vector subtree of $\mathbf{T} \upharpoonright_k$ then $\text{val}(S_i : i \in k-1)$ is a valuation tree and $\text{val}(S_i : i \in k-1) \subseteq \text{val}(S_i : i \in k)$. Also note that $T_0(< k) = \text{val}(T_i : i \in k)$ for every $k \in \omega+1$.

Example 9.3.1. Figure 9.1 depicts an example of a valuation tree. In this example, $\sigma = (1, 2, 1, 1, \dots)$, hence T_0 consists of all functions which assign 0 or 1 to every decreasing sequence of length 2, T_1 is a binary tree and T_k is just a chain for $k \geq 2$. Given a strong vector subtree (S_0, S_1, S_2) of (T_0, T_1, T_2) of height 3 (depicted by thick nodes and thick successor relations) we construct the corresponding valuation tree $\text{val}(S_0, S_1, S_2) \subseteq T_0$. Note that the topmost level of S_1 and the two topmost levels of S_2 are actually not used in the construction of $\text{val}(S_0, S_1, S_2) \subseteq T_0$, to depict this, they have grey colour.




Figure 9.1: A valuation tree (right) constructed from a strong vector subtree (left).

Let σ be an arbitrary signature, consider $\mathbf{T} = \mathbf{T}^\sigma$, and let T and T' be subtrees of T_0 . A function $\psi: T \rightarrow T'$ is a *structural embedding* if it is an embedding of trees (preserving meets and relative heights of nodes), and for every $x_0, \dots, x_{d-1} \in T$ with $|x_0| > \dots > |x_{d-1}|$ it holds that

$$\psi(x_0)(|\psi(x_1)|, \dots, |\psi(x_{d-1})|) = x_0(|x_1|, \dots, |x_{d-1}|).$$

Lemma 9.3.1. *Let σ be an arbitrary signature and consider $\mathbf{T} = \mathbf{T}^\sigma$. For every $k \in \omega + 1$ and every valuation subtree T of T_0 of height k , there exists a unique structural embedding $f: T_0(<k) \rightarrow T$.*

Proof. For $k \in \omega$ we use induction on k . Cases $k = 0$ and $k = 1$ are simple. Assume now that the induction hypothesis holds for some $k > 0$ and let $T = \text{val}(S_0, \dots, S_k)$ have height $k + 1$. Put $T' = \text{val}(S_1, \dots, S_k)$ (it has height k). By the induction hypothesis there exists a unique structural embedding $f: T_0(<k) \rightarrow T(<k)$ and a unique structural embedding $f': T_1(<k) \rightarrow T'$.

We will now extend f to $T_0(k)$. Fix $u \in T_0(k)$ and denote by $e: k + 1 \rightarrow \omega$ the increasing enumeration of $L_{T_0}(T)$. Let v be the predecessor of u on level $k - 1$, let $w \in T_1(k - 1)$ be the unique node such that $u \in v \wedge w$ and denote $x = u(k - 1)$. Then there is a unique $u' \in \text{ImmSucc}_{S_0}(f(v)) \cap \text{Succ}_{T_0}(f(v) \wedge f'(w))$ with $u'(e(k - 1)) = x$. Put $f(u) = u'$. It is easy to check that the extended map is a structural embedding $T_0(<k + 1) \rightarrow T$ and that the extension was defined in the unique possible way.

If $k = \omega$, we have proved that there are structural embeddings $f_i: T_0(<i) \rightarrow T(<i)$ for each $i \in \omega$. Since these isomorphisms are unique, we get $f_i \subset f_j$ for $i < j$, and $f = \bigcup \{f_i : i \in \omega\}$ is the desired structural embedding. On the other hand, if $g: T_0 \rightarrow T$ is a structural embedding then for each $i \in \omega$ the restriction $g \upharpoonright T_0(<i): T_0(<i) \rightarrow T(<i)$ is a structural embedding and thus has to be equal to f_i , consequently $g = f$. \square

9.4 L -structures on valuation trees

Most of this section will be spent proving the following proposition which will be the key ingredient in the proof of Theorem 9.1.1.

Proposition 9.4.1. *Let L be a relational language which consists of no unary relations and finitely many relations of every arity, and let \mathbf{H} be a universal countable L -hypergraph. Then \mathbf{H} has finite big Ramsey degrees.*

For this section, fix such L and \mathbf{H} . Assume that $H = \omega$. Let n_i be the number of relations of arity i and assume that $L = \{R^{i,j} : i \in \omega, 1 \leq j \leq n_i\}$ such that $R^{i,j}$ has arity i . Let σ be the signature defined by $\sigma_i = n_{i-1} + 2$ and let $\mathbf{T} = \mathbf{T}^\sigma$.

We define an L -hypergraph \mathbf{G} on T_0 by putting $\{x_1, \dots, x_i\} \in R_{\mathbf{G}}^{i,j}$ if and only if $1 \leq j \leq n_i$, $|x_1| > \dots > |x_i|$ and $x_1(|x_2|, \dots, |x_i|) = j$.

Both choices $j = 0$ and $j = \sigma_i + 1$ represent a non-relation. We will later need both of them for technical reasons. Equip \mathbf{G} with the enumeration defined by $f \leq g$ if and only if either $|f| < |g|$, or $|f| = |g|$ and $f(\bar{x}) < g(\bar{x})$, where \bar{x} is the lexicographically smallest tuple where f and g differ. Observe that a structural embedding $T_0 \rightarrow T_0$ induces an (increasing) embedding $\mathbf{G} \rightarrow \mathbf{G}$ with this enumeration.

Our aim is to embed \mathbf{H} into \mathbf{G} , transfer colourings of substructures of \mathbf{H} into colourings of vector subtrees of \mathbf{T} and use Milliken's theorem to obtain the desired Ramsey result. However, not all embeddings $\mathbf{H} \rightarrow \mathbf{G}$ are created equal. In order to prove a more robust result which can later be used in characterising the exact values of big Ramsey degrees, we first need to introduce some terminology.

Definition 9.4.1. Let f be a valuation function of dimension ω with $|f| = n$ and let $\bar{x} = (x_0, \dots, x_{m-1})$ with $n > x_0 > \dots > x_{m-1} \geq 0$. We denote by $f^{\bar{x}}$ the valuation function of level x_{m-1} defined by $f^{\bar{x}}(\bar{y}) = f(\bar{x} \hat{\ } \bar{y})$ and call it the \bar{x} -slice of f . An m -slice of f is an arbitrary \bar{x} -slice of f where $|\bar{x}| = m$.

In particular, a 0-slice of f is just f . Note that the elements of T_i are precisely the i -slices of elements of T_0 .

Definition 9.4.2. Let $k \in \omega$, let $\varphi: \omega \rightarrow T_0$ be an embedding $\mathbf{H} \rightarrow \mathbf{G}$ and denote by $R = \varphi[\omega]$ its range. We say that φ is a k -enveloping embedding if there are two disjoint sets $O, B \subseteq \omega$ such that the following holds for every $f \in R$, for every $m < k$ and for every $\bar{x} = (x_1, \dots, x_m)$:

1. If $f^{\bar{x}}$ is not constant zero then $x_i \in O$ for every i , and
2. if $g \in R$ or g is constant zero, g' is an m -slice of g , and $f^{\bar{x}}$ and g' are incomparable then $|f^{\bar{x}} \wedge g'| \in B$.

In this case we call the members of O the *original* levels and the members of B the *branching* levels. The slice $f^{\bar{x}}$ is called *original* if $x_i \in O$ (that is, all the sliced levels are original).

The intuition is that, when embedding $\mathbf{H} \rightarrow \mathbf{G}$, we “allocate” several levels before each image of a vertex of \mathbf{H} for branching, so that we can have strong control over what happens on the branching levels. Often (e.g. when there is an upper bound on the number of relations of every arity) it is possible to get an embedding which is k -enveloping for every k . However, when the number of relations of higher arities is increasing, we were so far unable to get such a uniform embedding (and we conjecture that it does not exist).

Observation 9.4.2. A k -enveloping embedding is also k' -enveloping for every $k' < k$. □

Lemma 9.4.3. For every $k \in \omega$, there is a k -enveloping embedding $\varphi: \mathbf{H} \rightarrow \mathbf{G}$.

Proof. Recall that the vertex set of H is ω . Define

$$J = \omega \cup \{\text{branch}^s(x_0, \dots, x_{m-1}) : 1 \leq s < \max_{i \in k} \{\sigma_{i+m}\}, x_0 > \dots > x_{m-1} \geq 0\},$$

where we treat $\text{branch}^s(\bar{x})$ as a formal expression. Define a linear order \triangleleft on J , putting $i \triangleleft j$ if and only if $i < j$, putting $i \triangleleft \text{branch}^s(x_0, \dots, x_{m-1})$ if and only if $i < x_0$ and putting $\text{branch}^s(\bar{x}) \triangleleft \text{branch}^t(\bar{y})$ if and only if either \bar{x} is lexicographically smaller than \bar{y} , or $\bar{x} = \bar{y}$ and $s < t$.

Let $\phi: \omega \rightarrow \omega$ be the function such that $\phi(i) = |\{w \in J : w \triangleleft i\}|$ (in particular, the $\phi(i)$ -th element of J is i). For every formal expression $\text{branch}^s(\bar{x}) \in J$ we now define its value

$$\text{branch}^s(\bar{x}) = |\{w \in J : w \triangleleft \text{branch}^s(\bar{x})\}|.$$

The range of ϕ will be the original levels O and the other levels (i.e. the range of branch) will be the branching levels B . Note that ϕ is an increasing function.

Let $\varphi: \mathbf{H} \rightarrow \mathbf{G}$ be the map where $\varphi(i) \in T_0$ is a valuation function with $|\varphi(i)| = \phi(i)$ and all entries equal to 0 except for the following two cases:

1. If $i > x_0 > \dots > x_{m-1}$ and $\{i, x_0, \dots, x_{m-1}\} \in R_{\mathbf{H}}^{m+1, r}$ then we have $\varphi(i)(\psi(x_0), \dots, \psi(x_{m-1})) = r$.
2. If $n < k$, $i > x_0 > \dots > x_{n-1} > y_0 > \dots > y_{m-1}$, and

$$\{i, x_0, \dots, x_{n-1}, y_0, \dots, y_{m-1}\} \in R_{\mathbf{H}}^{n+m+1, r}$$

then we have

$$\varphi(i)(\psi(x_0), \dots, \psi(x_{n-1}), \text{branch}^r(y_0, \dots, y_{m-1})) = \sigma_{n+m} - 1.$$

Part (1) ensures that φ is an embedding. We will prove that it is in fact k -enveloping, that is, we will verify points 1 and 2 of Definition 9.4.2.

Point 1 follows straightforwardly from our construction of φ : The only tuples not valued by 0 either consist of original levels only, or the very last level is branching.

To see point 2, suppose that $f' = f^{\bar{x}}$, $g' = g^{\bar{y}}$, $|\bar{x}| = |\bar{y}| < k$, and f' and g' are incomparable. This means that there is a level n such that $f' \upharpoonright_{I_{\bar{w}}^n} = g' \upharpoonright_{I_{\bar{w}}^n}$, but $f' \upharpoonright_{I_{\bar{w}}^{n+1}} \neq g' \upharpoonright_{I_{\bar{w}}^{n+1}}$. That is equivalent to n being the least integer for which there exists a decreasing sequence \bar{z} such that

$$f(\bar{x} \wedge n \wedge \bar{z}) \neq g(\bar{y} \wedge n \wedge \bar{z}).$$

If n is branching then we are done. So n is original, that is, $n = \phi(n')$ for some $n' \in \omega$. Let $\bar{z} = (z_1, \dots, z_p)$ and denote $a = f(\bar{x} \wedge n \wedge \bar{z})$ and $b = g(\bar{y} \wedge n \wedge \bar{z})$. We know that $a \neq b$. We also know that z_1, \dots, z_{p-1} are original (because only the last level of a tuple valued by a non-zero integer can be non-original), so $\phi^{-1}(z_i)$ is defined for every $1 \leq i \leq p-1$.

First suppose that $z_p = \text{branch}^c(w_1, \dots, w_q)$ for some c . Then we know that exactly one of a, b is equal to zero. In this case put

$$\bar{z}' = (\phi^{-1}(z_1), \dots, \phi^{-1}(z_{p-1}), w_1, \dots, w_q).$$

By the construction we know that $f(\bar{x} \wedge \text{branch}^c(k' \wedge \bar{z}')) \neq g(\bar{x} \wedge \text{branch}^c(k' \wedge \bar{z}'))$, because again exactly one of them is equal to 0. And as $\text{branch}^c(k' \wedge \bar{z}') < \phi(k') = k$, we get a contradiction with minimality of k .

So z_p is original. Put $\bar{z}' = (\phi^{-1}(z_1), \dots, \phi^{-1}(z_p))$ and assume without loss of generality that $a \neq 0$ (at least one of a, b is non-zero). Since z_p is original, from the construction it follows that $a \neq \sigma_{|\bar{x} \wedge k \wedge \bar{z}|} - 1$, which implies that $\text{branch}^a(k' \wedge \bar{z}')$ is defined. Note that $f(\bar{x} \wedge \text{branch}^a(k' \wedge \bar{z}')) \neq 0$ while $g(\bar{x} \wedge \text{branch}^a(k' \wedge \bar{z}')) = 0$. Again, as $\text{branch}^a(k' \wedge \bar{z}') < k$, we get a contradiction with minimality of k , which verifies that point 2 is also satisfied, hence φ is indeed k -enveloping. \square

Remark 9.4.1. Note that if there is an absolute bound N on the number of relations of any arity then there exists an embedding which is k -enveloping for every $k \in \omega$. Indeed, one can pretend in the above proof that $k = \omega$ and use the fact that we always have $\max_{i \in \omega} \{\sigma_{i+m}\} \leq N$ for any $m \in \omega$.

When there is no such bound, we have been unable to produce such an embedding and we conjecture that actually there is no embedding where all envelopes would be bounded:

Conjecture 9.4.4. *If for every $n \in \omega$ there is a $a \in \omega$ such that the number of relations of arity a is at least n then there is no embedding $\varphi: \mathbf{H} \rightarrow \mathbf{G}$ with the property that for every $k \in \omega$ there exists $R(k) \in \omega$ such that every set $S \subseteq \varphi[\mathbf{H}]$ with $|S| = k$ has an envelope of height at most $R(k)$.*

9.4.1 Envelopes

As was noted earlier, our goal is to transfer colourings of substructures of a nice copy of \mathbf{H} in \mathbf{G} to colourings of valuation subtrees of T_0 .

Definition 9.4.3. Given $S \subseteq T_0$ and a valuation subtree $T \subseteq T_0$ of height $k \in \omega + 1$, we say that T is an *envelope* of S if $S \subseteq T$.

Having an enveloping embedding allows us to envelope finite subsets of its range in bounded-height valuation trees:

Lemma 9.4.5. *For every $k \in \omega$ there exists $R(k) \in \omega$ such that for every k -enveloping embedding $\varphi: \mathbf{H} \rightarrow \mathbf{G}$ and every set $S \subset H$ of size k it holds that $\varphi[S]$ has an envelope of height at most $R(k)$.*

Proof. Fix a set S of k elements from H . Put $E_0^1 = \varphi[S]$ and $E_0^2 = \{f \wedge_{T_0} g : f, g \in \varphi[S]\} \subset T_0$. Define by induction sets E_i^0, E_i^1 and E_i^2 for $1 \leq i < R(k)$, where $R(k)$ will be defined later, as follows:

1. $E_i^0 = \{f^{(|g|)} : f, g \in E_{i-1}^2, |g| < |f|\} \subseteq T_i$,
2. $E_i^1 = E_i^0 \cup \{z\}$, where z is the constant zero valuation function of level $\max\{|f| : f \in E_i^0\}$,
3. $E_i^2 = \{f \wedge_{T_i} g : f, g \in E_i^1\}$.

Claim 9.4.6. *The following properties hold for every $0 \leq i < R(k)$:*

1. $E_i^2 \supseteq E_i^1 \supseteq E_i^0$ and E_i^2 is a subtree of T_i ,

2. each element of E_i^1 is either constant zero or an original i -slice of a member of $\varphi[S]$ and each element of E_i^2 is a restriction of an element of E_i^1 ,
3. there are at most $\max(0, k - i)$ levels with non-zero members of E_i^1 ,
4. for $i > 0$, every level in $L_{T_i}(E_i^2) \setminus L_{T_{i-1}}(E_{i-1}^2)$ is branching,
5. if $\ell \in L_{T_i}(E_i^2)$ then either ℓ is branching or there is $f \in \varphi[S]$ with $|f| = \ell$.
6. for $i > 0$ we have that $L_{T_{i-1}}(E_{i-1}^2) \setminus L_{T_i}(E_i^2) = \{\max(L_{T_{i-1}}(E_{i-1}^2))\}$,
7. for $i > 0$, every level in $L_{T_i}(E_i^2)$ is strictly smaller than $\max(L_{T_{i-1}}(E_{i-1}^2))$.

We will proceed by induction on i . For $i = 0$ this is immediate. So $i > 0$ and we know that all properties hold for $i - 1$. Property (1) is straightforward from the definition. We know that (2) holds for $i - 1$, so E_{i-1}^0 consists of slices of members of $\varphi[S]$ and constant zero functions, so this is true also for E_i^1 , and in constructing E_i^2 we are only adding restrictions of members of E_i^1 .

Note that if $i > k$ then, by (3), we know that we only had constant zero functions in E_{i-1}^1 and consequently also in E_{i-1}^2 . So it suffices to prove (3) for $i \leq k$. By (5) for $i - 1$ we have that the only original levels we can slice with correspond to members of $\varphi[S]$ and by (2) we know that the only non-zero members of E_i^1 are original i -slices of members of $\varphi[S]$. To get an i -slice, we need $i + 1$ members of $\varphi[S]$ (one to slice and the other i for a decreasing sequence of original levels to slice with), hence there are only $k - i$ possible last elements of the decreasing sequence. This verifies (3).

To see (4), note that we only add new levels compared to E_{i-1}^2 when taking meets in the construction of E_i^2 . This means that we can assume that $i < k$ as otherwise all elements of E_i^1 are comparable. By (2) we know that each member of E_i^1 is either constant zero or an original i -slice of a member of $\varphi[S]$, and since φ is k -enveloping and $i < k$ it follows that their meets happen on branching levels.

Property (5) is immediate from (4) and the induction hypothesis. Property (6) is also easy, in the construction of E_i^0 we only lose the highest level. Having (6), (7) is again easy, because we only add new levels as meets of functions from existing levels. This finishes the proof of Claim 9.4.6. \square

Property (3) implies that each member of E_k^1 is constant zero, hence $E_k^2 = E_k^1$. It follows that for every $i > k$ we have $E_i^1 = E_i^2$ and all its elements are constant zeros. Consequently $L_{T_i}(E_i^2) \subsetneq L_{T_{i-1}}(E_{i-1}^2)$ and these two level sets differ precisely by $\max(L_{T_{i-1}}(E_{i-1}^2))$ by (6).

Put $\mathfrak{L} = \bigcup_{i=0}^k L_{T_i}(E_i^2)$. Clearly, for every $i \geq 1$ it holds that $|L_{T_i}(E_i^1)| \leq |L_{T_{i-1}}(E_{i-1}^2)|$ and for every $i \geq 0$ we have $|L_{T_i}(E_i^2)| \leq 2|L_{T_i}(E_i^1)| - 1$. Together with $|L_{T_0}(E_0^1)| \leq k$ this implies that $|\mathfrak{L}|$ is bounded from above by some $R(k)$ which is a function of k . In particular, $|L_{T_j}(E_j^2)| \leq R(k) - j$, and so $E_{R(k)-1}^2$ is a singleton set containing a constant zero function.

For every $i \in R(k)$ define

$$E_i^3 = \{f \upharpoonright_\ell : f \in E_i^2, \ell \in \mathfrak{L}, |f| \geq \ell\}.$$

Note that $L_{T_i}(E_i^3) = \mathfrak{L}$ and that E_i^3 is a subtree of T_i for every $i \in R(k)$. For every $i \in R(k)$ let S_i be some strong subtree of T_i containing E_i^3 such that $L_{T_i}(S_i) = \mathfrak{L}$. Such trees exist by Observation 9.2.1.

It remains to prove that $\varphi[S] \subseteq \text{val}(S_i : i \in R(k))$. We will prove the following stronger result:

Claim 9.4.7. *For every $i \in R(k)$ and $\ell \in \mathfrak{L}$ it holds that if $f \in E_{R(k)-i-1}^3$ and $\ell \leq |f|$ then $f \upharpoonright_{\ell} \in \text{val}(S_{R(k)-i-1}, \dots, S_{R(k)-1})$.*

We will prove this statement by double induction on i (outer induction) and $\ell \in \mathfrak{L}$ (inner induction). For $i = 0$ this is easy because $E_{R(k)-1}^3$ consists only of the root of $S_{R(k)-1}$ which is also the constant zero function of level $\min(\mathfrak{L})$ and is the root of $\text{val}(S_{R(k)-1})$.

Assume now that the statement is true for every $j < i$, every $\ell \in \mathfrak{L}$ and also for i and every $\ell' \in \mathfrak{L}$ such that $\ell' < \ell$, we will prove it for i and ℓ . Pick an arbitrary $f \in E_{R(k)-i-1}^3$. If $\ell = \min(\mathfrak{L})$ then $f \upharpoonright_{\ell}$ is the root of $S_{R(k)-i-1}$ and hence the root of $\text{val}(S_{R(k)-i-1}, \dots, S_{R(k)-1})$.

So $\ell > \min(\mathfrak{L})$. Let ℓ' be the largest member of \mathfrak{L} smaller than ℓ . By the construction, $f \upharpoonright_{\ell'} \in E_{R(k)-i-1}^3$, and so $f \upharpoonright_{\ell'} \in \text{val}(S_{R(k)-i-1}, \dots, S_{R(k)-1})$ by the induction hypothesis for i and ℓ' .

If ℓ' is branching then $f^{(\ell')}$ is the constant zero valuation function of level ℓ' and as such is in $E_{R(k)-i}^3$ and consequently in $\text{val}(S_{R(k)-i}, \dots, S_{R(k)-1})$ by the induction hypothesis. Otherwise ℓ' is original, but then $f^{(\ell')} \in E_{R(k)-i}^0 \subseteq E_{R(k)-i}^3$ and hence also in $\text{val}(S_{R(k)-i}, \dots, S_{R(k)-1})$ by the induction hypothesis.

So we know that

$$f \upharpoonright_{\ell'} \in \text{val}(S_{R(k)-i-1}, \dots, S_{R(k)-1})$$

and

$$f^{(\ell')} \in \text{val}(S_{R(k)-i}, \dots, S_{R(k)-1}).$$

The definition of valuation tree then gives that $f \upharpoonright_{\ell} \in \text{val}(S_{R(k)-i-1}, \dots, S_{R(k)-1})$ which concludes the proof of Claim 9.4.7. \square

As a special case of Claim 9.4.7 we get that $\varphi[S] \subseteq E_0^1 \subseteq E_0^3 \subseteq \text{val}(S_i : i \in R(k))$. This means that there indeed is a valuation tree of height at most $R(k)$ which contains $\varphi[S]$. \square

We can now proceed with the proof of Proposition 9.4.1.

Proof of Proposition 9.4.1. Fix a finite L -hypergraph \mathbf{A} with $|A| = k$, a k -enveloping embedding $\varphi : \mathbf{H} \rightarrow \mathbf{G}$ (for example one given by Lemma 9.4.3), and a colouring $\chi^0 : \binom{\mathbf{H}}{\mathbf{A}} \rightarrow p$. Since \mathbf{H} is universal, it follows that there is an embedding $\theta : \mathbf{G} \rightarrow \mathbf{H}$. Consider the colouring $\chi : \binom{\mathbf{G}}{\mathbf{A}} \rightarrow p$ obtained by setting $\chi(\tilde{\mathbf{A}}) = \chi^0(\theta(\tilde{\mathbf{A}}))$ for every $\tilde{\mathbf{A}} \in \binom{\mathbf{G}}{\mathbf{A}}$.

Let $h = R(k)$ be given by Lemma 9.4.5 and let \mathbf{G}_h be the induced sub- L -hypergraph of \mathbf{G} on $T_0(<h)$. We enumerate the copies of \mathbf{A} in $\binom{\mathbf{G}_h}{\mathbf{A}}$ as $\{\tilde{\mathbf{A}}_i : i \in \ell\}$ for some $\ell \in \omega$.

By Lemma 9.3.1, for every valuation tree T of height h , there is a structural embedding $f_T : \mathbf{G}_h \rightarrow T$ that is also an isomorphism of the corresponding substructures of \mathbf{G} . Let $\mathbf{S} = (S_i : i \in h)$ be a strong subtree of $\mathbf{T} \upharpoonright_h$ of height h and consider the structural embedding $f = f_{\text{val}(S_i : i \in h)} : \mathbf{G}_h \rightarrow \text{val}(S_i : i \in h)$. Put

$$\bar{\chi}(\mathbf{S}) = \langle \chi(f(\tilde{\mathbf{A}}_i)) : i \in \ell \rangle,$$

which is a finite colouring of $\text{Str}_h(\mathbf{T} \upharpoonright_h)$. By Theorem 9.2.2, there is an infinite strong subtree of $\mathbf{T} \upharpoonright_h$ monochromatic with respect to $\bar{\chi}$. Extend it arbitrarily to an infinite strong subtree of \mathbf{T} with the same level set and let U be its corresponding valuation subtree. Note that the extension does not influence in any way the valuation subtrees of height h . The structural embedding $\psi: T_0 \rightarrow U$ given by Lemma 9.3.1 is a hypergraph embedding $\psi: \mathbf{G} \rightarrow \mathbf{G}$.

We claim that $\theta[\psi[\varphi[\mathbf{H}]]]$ is the desired copy $g[\mathbf{H}]$ of \mathbf{H} , in which copies of \mathbf{A} have at most ℓ different colours in χ^0 . This is true, because by Lemma 9.4.5 every copy of \mathbf{A} in $\varphi[\mathbf{H}]$ is contained in a valuation subtree of height h . All of these subtrees have the same colour (with respect to $\bar{\chi}$) in $\psi[\varphi[\mathbf{H}]]$, and so we know that χ takes at most ℓ different values on $\binom{\psi[\varphi[\mathbf{H}]]}{\mathbf{A}}$. Consequently, χ^0 attains at most ℓ different values on $\binom{\theta[\psi[\varphi[\mathbf{H}]]]}{\mathbf{A}}$. \square

9.5 The main results

In this section we do three simple constructions on top of Proposition 9.4.1 in order to prove Theorem 9.1.1.

9.5.1 Unary relations

First, we introduce a general construction for adding unary relations in order to prove the following theorem.

Theorem 9.5.1. *Let L be a relational language with finitely many relations of every arity greater than one and with finitely or countably many unary relations. Let \mathbf{H} be a countable universal L -hypergraph. Then \mathbf{H} has finite big Ramsey degrees.*

Proof. Assume that the unary relations of L are $\{U^{i+1} : i \in u\}$ for some $u \in \omega + 1$. Let L^- be the language which one gets from L by removing all unary relations and let \mathbf{M} be a universal L^- -hypergraph (it exists for example by the Fraïssé theorem [Fra53]). Without loss of generality we assume that the vertex set of both \mathbf{M} and \mathbf{H} is ω . We now define an L -structure \mathbf{G} as follows:

1. The vertex set of \mathbf{G} is $G = \{(v, i) : v \in \omega, i \in \min(v + 1, u)\}$ with enumeration given by the lexicographic order,
2. vertex (v, i) is in unary relation U^j if and only if $i = j$ (in particular, $(v, 0)$ is in no unary relations), and
3. $((x_0, i_0), \dots, (x_n, i_n)) \in R_{\mathbf{G}}$ if and only if $(x_0, \dots, x_n) \in R_{\mathbf{M}}$.

Fix some finite L -hypergraph \mathbf{A} and let \mathbf{A}^- be its L^- -reduct. By Proposition 9.4.1 there is $\ell \in \omega$ such that $\mathbf{M} \rightarrow (\mathbf{M})_{k, \ell}^{\mathbf{A}^-}$ for every $k \in \omega$. We will now prove that $\mathbf{G} \rightarrow (\mathbf{H})_{k, \ell}^{\mathbf{A}}$ for an arbitrary $k \in \omega$.

Let $\pi: G \rightarrow M$ be the map sending $(v, i) \mapsto v$. We say that a map $f: X \rightarrow G$ is *transversal* if $\pi \circ f$ is injective. Note that if $f: \mathbf{A} \rightarrow \mathbf{G}$ is a transversal embedding then $\pi \circ f$ is an embedding $\mathbf{A}^- \rightarrow \mathbf{M}$. A colouring $\chi_0: \binom{\mathbf{G}}{\mathbf{A}} \rightarrow k$ then induces a partial colouring $\chi: \binom{\mathbf{M}}{\mathbf{A}^-} \rightarrow k$ by ignoring non-transversal copies and

composing with π . There may be copies of \mathbf{A}^- which do not get any colour, we assign them a colour arbitrarily.

Since $\mathbf{M} \rightarrow (\mathbf{M})_{k,\ell}^{\mathbf{A}^-}$, there is an embedding $\psi: \mathbf{M} \rightarrow \mathbf{M}$ with χ attaining at most ℓ colours on $(\psi[\mathbf{M}]_{\mathbf{A}^-})$. Let $\psi_0: \mathbf{G} \rightarrow \mathbf{G}$ be the embedding mapping $(v, i) \mapsto (\psi(v), i)$. We have that χ_0 attains at most ℓ colours on transversal copies from $(\psi_0[\mathbf{G}]_{\mathbf{A}})$.

Let \mathbf{H}^- be the L^- -reduct of \mathbf{H} forgetting unary relations. Since \mathbf{H}^- is a countable L^- -hypergraph, there is an embedding $\varphi_0: \mathbf{H}^- \rightarrow \mathbf{M}$ with the property that $\varphi_0(x) \geq i_x$, where $i_x = 0$ if x is in no unary relation of \mathbf{H} and $i_x = j$ if $x \in U_{\mathbf{H}}^j$. It is straightforward to check that the map $\varphi: H \rightarrow G$, defined by $\varphi(x) = (\varphi_0(x), i_x)$, is a transversal embedding $\mathbf{H} \rightarrow \mathbf{G}$, hence χ_0 attains at most ℓ colours on $(\psi_0[\varphi[\mathbf{H}]]_{\mathbf{A}})$. Consequently, $\mathbf{G} \rightarrow (\mathbf{H})_{k,\ell}^{\mathbf{A}}$.

Knowing that $\mathbf{G} \rightarrow (\mathbf{H})_{k,\ell}^{\mathbf{A}}$, proving that $\mathbf{H} \rightarrow (\mathbf{H})_{k,\ell}^{\mathbf{A}}$ is straightforward: \mathbf{G} is a countable L -hypergraph and so there is an embedding $\theta: \mathbf{G} \rightarrow \mathbf{H}$, which means that a colouring of $(\mathbf{H})_{k,\ell}^{\mathbf{A}}$ restricts to a colouring of $(\mathbf{G})_{k,\ell}^{\mathbf{A}}$ exactly as in the proof of Proposition 9.4.1. \square

9.5.2 Non- L -hypergraphs

For the constructions above it was convenient to work with L -hypergraphs. In this section we state a folkloristic result showing that one does not lose any generality working with L -hypergraphs only. Its proof is a straightforward verification of the construction.

Lemma 9.5.2. *Let L be a relational language and let \mathcal{C} be the class of all L -structures where every relation is injective and every vertex is in exactly one unary relation. For every $i \in \omega$, put*

$$M_i = \{(R, \pi) : R \in L, i = a(R), \pi \in \text{Sym}(i)\}.$$

Given $\mathbf{A} \in \mathcal{C}$ and vertices $x_1 < \dots < x_n \in \mathbf{A}$, put

$$M_{\mathbf{A}}(x_1, \dots, x_n) = \{(R, \pi) \in M_n : (x_{\pi(1)}, \dots, x_{\pi(n)}) \in R_{\mathbf{A}}\}.$$

Define a language L' containing an i -ary relation R^S for every nonempty finite $S \subseteq M_i$, for every $i \in \omega$, and let \mathcal{C}' be the class of all L' -hypergraphs. If L has only finitely many relations of every arity greater than one and finitely or countably many unary relations then so does L' .

Define T to be the map assigning to every $\mathbf{A} \in \mathcal{C}$ a structure $T(\mathbf{A}) \in \mathcal{C}'$ on the same vertex set such that

$$\{x_1 < \dots < x_n\} \in R_{T(\mathbf{A})}^{M_{\mathbf{A}}(x_1, \dots, x_n)}$$

if $M_{\mathbf{A}}(x_1, \dots, x_n)$ is nonempty. There are no other relations in $T(\mathbf{A})$. Define U to be the map assigning to every $\mathbf{A} \in \mathcal{C}'$ a structure $U(\mathbf{A}) \in \mathcal{C}$ on the same vertex set such that whenever we have $x_1 < \dots < x_n \in \mathbf{A}$ with $\{x_1, \dots, x_n\} \in R_{\mathbf{A}}^S$, we put

$$(x_{\pi(1)}, \dots, x_{\pi(n)}) \in R_{U(\mathbf{A})}$$

for every $(R, \pi) \in S$. There are no other relations in $U(\mathbf{A})$.

Then T and U are mutually inverse and define a bijection between \mathcal{C} and \mathcal{C}' . Moreover, given $\mathbf{A}, \mathbf{B} \in \mathcal{C}$ and a function $f: A \rightarrow B$, it holds that f is an embedding $\mathbf{A} \rightarrow \mathbf{B}$ if and only if it is an embedding $T(\mathbf{A}) \rightarrow T(\mathbf{B})$. \square

Corollary 9.5.3. *Let L be a relational language with finitely many relations of every arity greater than one and with finitely or countably many unary relations and let \mathbf{H} be a countable universal L -structure where all relations are injective and every vertex is in exactly one unary relation. Then \mathbf{H} has finite big Ramsey degrees.*

Note that this statement is the same as the statement of Theorem 9.1.1 for $\mathcal{F} = \emptyset$.

9.5.3 Forbidding structures

Proof of Theorem 9.1.1. Since \mathbf{H} is unrestricted, it is countable and hence there are only countably many types of vertices. Hence, without loss of generality, we can assume that each vertex of \mathbf{H} is in exactly one unary relation (by changing the language similarly as in Lemma 9.5.2) and that \mathcal{F} forbids no single unary relation (otherwise we can remove it from the language).

Let \mathbf{M} be a countable universal L -structure where all relations are injective and every vertex is in one unary relation (it exists for example by the Fraïssé theorem [Fra53]). We say that a subset $S \subseteq M$ is *bad* if \mathbf{M} induces a structure from \mathcal{F} on S and we say that a tuple \bar{x} of elements of \mathbf{M} is *bad* if it contains a bad subset. Let \mathbf{G} be the L -structure on the same vertex set as \mathbf{M} such that $\bar{x} \in R_{\mathbf{G}}$ if and only if $\bar{x} \in R_{\mathbf{M}}$ and \bar{x} is not bad (i.e. we remove bad tuples from all relations). Clearly, \mathbf{G} is \mathcal{F} -free, and note that \mathbf{G} and \mathbf{M} have the same unary relations. We will prove that \mathbf{G} has finite big Ramsey degrees. Since \mathbf{G} embeds into \mathbf{H} and any embedding $\mathbf{H} \rightarrow \mathbf{M}$ is also an embedding $\mathbf{H} \rightarrow \mathbf{G}$, this would imply that \mathbf{H} has finite big Ramsey degrees, thereby proving the theorem.

Fix a finite \mathcal{F} -free L -structure \mathbf{A} where all relations are injective. Let ι be the identity map understood as a function $G \rightarrow M$ and let $\mathbf{A}_0, \dots, \mathbf{A}_m$ be some enumeration of all isomorphism types of structures from $\{\iota \circ f[\mathbf{A}] : f \text{ is an embedding } \mathbf{A} \rightarrow \mathbf{G}\}$, that is, it is an enumeration of all possible isomorphism types which, after removing bad tuples from relations, are isomorphic to \mathbf{A} . There are only finitely many of them because they have the same unary relations as \mathbf{A} and there are only finitely many L -structures on a given number of vertices with given unary relations, up to isomorphism.

For every $0 \leq i \leq m$, let ℓ_i be the big Ramsey degree of \mathbf{A}_i in \mathbf{M} (ℓ_i is finite by Corollary 9.5.3) and put $\ell = \sum_{i=0}^m \ell_i$. We now prove that $\mathbf{G} \rightarrow (\mathbf{G})_{k,\ell}^{\mathbf{A}}$ for every $k \in \omega$.

Fix a colouring $\chi: \binom{\mathbf{G}}{\mathbf{A}} \rightarrow k$ and let $\chi_i: \binom{\mathbf{M}}{\mathbf{A}_i} \rightarrow k$, $0 \leq i \leq m$, be the colourings obtained from χ by composing with ι . By inductive usage of Corollary 9.5.3 we get an embedding $f: \mathbf{M} \rightarrow \mathbf{M}$ such that, for every $0 \leq i \leq m$, χ_i attains at most ℓ_i colours on $\binom{f[\mathbf{M}]}{\mathbf{A}_i}$. By the construction it follows that f is also an embedding $\mathbf{G} \rightarrow \mathbf{G}$, and since every copy in $\binom{f[\mathbf{G}]}{\mathbf{A}}$ corresponds to a copy of \mathbf{A}_i in $f[\mathbf{M}]$ for some i , it follows that χ attains at most ℓ colours on $\binom{f[\mathbf{G}]}{\mathbf{A}}$, hence the big Ramsey degree of \mathbf{A} in \mathbf{G} is indeed finite. \square

9.6 Infinite big Ramsey degrees?

As soon as one has infinite branching, the Milliken theorem stops being true even for colouring vertices (see Proposition 9.6.3). This means that one cannot generalise our methods directly for languages with infinitely many relations of some arity at least 2. In fact, no known methods generalise because all of them find very specific *tree-like* copies and one can construct infinite colourings which are persistent on these copies. Doing this in full generality requires developing the theory of *weak types* and it will appear elsewhere. Here we only show a special case which is technically much simpler but only works for binary relations.

Definition 9.6.1. Let \mathbf{H} be a countable relational structure with vertex set ω . Given $X \subseteq \omega$, we define the *type of X* , denoted by $\text{tp}_{\mathbf{H}}(X)$, to be the isomorphism type of the substructure of \mathbf{H} induced on X . (Note that this corresponds to the quantifier-free type over the empty set from model theory if we consider the enumeration to be part of the language.)

Let $f: \mathbf{H} \rightarrow \mathbf{H}$ be an embedding. We say that f is *1-tree-like* (or, in this paper, simply *tree-like*) if for every finite $X = \{x_0 < \dots < x_m\} \subset \omega$, every $0 \leq i \leq m$ and every $x \in \omega$ such that $x > x_m$ there exists $y > x_m$ such that

$$\text{tp}_{\mathbf{H}}(f[X] \cup \{f(y)\}) = \text{tp}_{\mathbf{H}}(X \cup \{x\})$$

and moreover

$$\text{tp}_{\mathbf{H}}(\{0, \dots, f(x_0) - 1, f(y)\}) = \text{tp}_{\mathbf{H}}(\{0, \dots, f(x_0) - 1, f(x_i)\}).$$

For every embedding f we have that $\text{tp}_{\mathbf{H}}(f[X] \cup \{f(x)\}) = \text{tp}_{\mathbf{H}}(X \cup \{x\})$. The second condition says that for tree-like embeddings we have some control even over types with respect to the ambient structure \mathbf{H} within which our copy lies. Note that every structural embedding $\mathbf{T}^\sigma \rightarrow \mathbf{T}^\sigma$ is a tree-like embedding of the corresponding hypergraphs.

Example 9.6.1. Let L be a language consisting of infinitely many binary relations and let \mathbf{R} be the countable homogeneous L -hypergraph (that is, an infinite-edge-coloured countable random graph where the colour classes are generic). One can repeat our constructions for \mathbf{R} and get the everywhere infinitely branching tree $T = [\omega]^{<\omega}$ as T_0 . If the Milliken theorem was true for infinitely branching trees, it would produce tree-like copies (which would arise simply as strong subtrees isomorphic to $[\omega]^{<\omega}$).

On the other hand, assume that the vertex set of \mathbf{R} is ω and consider any embedding $f: \mathbf{R} \rightarrow \mathbf{R}$ such that $f(i)$ is connected to vertex 0 by the i -th relation (such an embedding exists by the extension property). This embedding is somehow as far as possible from being tree-like, because everything branches at level 0 and one cannot say anything about the behaviour of this copy with respect to the external vertices.

All known big Ramsey methods produce tree-like copies, and whenever the exact big Ramsey degrees are known, the proof can be adapted to show that every copy contains a subcopy which is “weakly tree-like”. Example 9.6.1 shows that such a property fails for the infinite-edge-coloured random graph. We believe that working in the category of tree-like embeddings (or some other variant of

nice embeddings) would be a completely reasonable thing to do as it still captures (a lot of) the combinatorial complexity of the problems, in addition to allowing one to prove negative results.

In this section we will see an instance of this. Given a relational structure \mathbf{A} , its Gaifman graph is the graph on the same vertex set where vertices $x \neq y \in A$ are connected by an edge if and only if there exists a tuple \bar{z} of vertices of \mathbf{A} containing both x and y which belongs to a relation of \mathbf{A} . A relational structure \mathbf{A} is *irreducible* if its Gaifman graph is a complete graph.

Theorem 9.6.1. *Let L be a relational language, let \mathcal{F} be a set of finite irreducible L -structures and let \mathbf{H} be a universal \mathcal{F} -free structure. Assume that there are \mathcal{F} -free structures \mathbf{B} and \mathbf{U} consisting of one vertex, and infinitely many pairwise non-isomorphic 1-vertex \mathcal{F} -free extensions $\mathbf{C}_0, \mathbf{C}_1, \dots$ of \mathbf{B} such that the added vertices come last in the enumeration and each of them is isomorphic to \mathbf{U} . Then there exists an \mathcal{F} -free structure \mathbf{A} on three vertices and a colouring $c: \binom{\mathbf{H}}{\mathbf{A}} \rightarrow \omega$ such that if $f: \mathbf{H} \rightarrow \mathbf{H}$ is a tree-like embedding then c attains all values on $\binom{f[\mathbf{H}]}{\mathbf{A}}$.*

In other words, the big Ramsey degree of \mathbf{A} is infinite for tree-like copies. Note that if \mathcal{F} either contains no enumeration or all enumerations of every finite \mathbf{A} , the class of all finite \mathcal{F} -free structures is a free amalgamation class when one ignores the enumerations.

Corollary 9.6.2. *Let L be a relational language containing infinitely many binary relations and let \mathbf{H} be a countable universal L -hypergraph. Then there exists an L -hypergraph \mathbf{A} on three vertices and a colouring $c: \binom{\mathbf{H}}{\mathbf{A}} \rightarrow \omega$ such that if $f: \mathbf{H} \rightarrow \mathbf{H}$ is a tree-like embedding then c attains all values on $\binom{f[\mathbf{H}]}{\mathbf{A}}$.*

Proof. Note that being an L -hypergraph can be described by a set \mathcal{F} of forbidden irreducible structures (namely those where a tuple is in a relation, but not all possible permutations of the tuple are). Let \mathbf{B} be the empty hypergraph on 1 vertex which is in no unary relations and let $\mathbf{C}_0, \mathbf{C}_1, \dots$ be all possible one-vertex extensions of \mathbf{B} by a vertex in no unary relation. There are infinitely many of them as they correspond to binary relations in L . Therefore the conditions of Theorem 9.6.1 are satisfied and the conclusion follows. \square

A particular example of a structure satisfying this corollary is the infinite-edge-coloured random graph from Example 9.6.1.

Our proof of Theorem 9.6.1 is derived from a proof that the Halpern–Läuchli theorem does not hold for the tree $[\omega]^{<\omega}$, which we now present as a nice warm-up. We believe that this proof is folkloristic, but we were unable to find it in the literature.

Proposition 9.6.3. *Let $T = [\omega]^{<\omega}$ be the tree of all finite sequences of natural numbers and let \sqsubseteq be the usual tree order by end-extension. There is a colouring $c: T \rightarrow \omega$ such that whenever T' is a strong subtree of T of infinite height then $c[T'] = \omega$.*

Proof. Given $t \in T$, we denote by $|t|$ the length (level) of t , we put $w(t) = |t| + \sum_{i < |t|} t(i)$ and we define $\ell(t)$ to be the least ℓ such that $w(t \upharpoonright_\ell) \geq |t|$. Note that $\ell(t)$ always exists as $w(t) \geq |t|$.

We define a colouring $c: T \rightarrow \omega$ putting $c(t) = w(t \upharpoonright_{\ell(t)}) - |t|$. Let T' be an arbitrary strong subtree of T of infinite height, let r be the root of T' and let $n \in \omega$ be such that $n > w(r)$ and $T' \cap T(n) \neq \emptyset$. Put $k = n - w(r) - 1$. Now we will prove that $c[T'] = \omega$. For that, fix a colour $x \in \omega$ and find $t \in T'$ such that $|t| = n$ and $r^\wedge(k+x) \sqsubseteq t$ (such t exists as T' is a strong subtree of T). Now, $w(r) < n$, so $\ell(t) > |r|$. On the other hand, $w(r^\wedge(k+x)) = k+x+w(r)+1 = n+x \geq n$, hence $\ell(t) = |r|+1$ and $c(t) = x$. This means that for every $x \in \omega$ we can find $t \in T'$ such that $c(t) = x$, hence indeed $c[T'] = \omega$. \square

Note that in the construction of the colouring c we essentially only needed to be able to address a particular level $\ell(t)$ on which we knew that the passing numbers attain all possible values. We will now show how this idea can be adapted to prove infinite big Ramsey degrees.

Proof of Theorem 9.6.1. Let $\mathbf{B}_0, \mathbf{B}_1, \dots$ be an enumeration of all copies of \mathbf{B} in \mathbf{H} such that the vertex of \mathbf{B}_i comes in the enumeration of \mathbf{H} before the vertex of \mathbf{B}_j whenever $i < j$. There are infinitely many of them because the infinite disjoint union of copies of \mathbf{B} is \mathcal{F} -free and thus embeds into \mathbf{H} .

Given a vertex $v \in \mathbf{H}$ isomorphic to \mathbf{U} (i.e. it has the same unary relations), we let $b(v)$ be the least integer such that there is $w \in \mathbf{B}_{b(v)}$ with $w \geq v$, and let $s(v): b(v) \rightarrow \omega$ be the sequence satisfying $s(v)_i = x$ if and only if the structure induced by \mathbf{H} on $\mathbf{B}_i \cup \{v\}$ is isomorphic to \mathbf{C}_x .

Given a vertex $v \in \mathbf{H}$ isomorphic to \mathbf{B} , we let $j(v)$ be such that $\{v\} = B_{j(v)}$. Given $s \in [\omega]^{<\omega}$ and $n \leq |s|$, we put $w(s) = |s| + \sum_{i \in |s|} s(i)$ and define $\ell(s, n)$ to be the least ℓ such that $w(s \upharpoonright_\ell) \geq n$. It exists since $n \leq |s|$.

Let \mathbf{A} be an enumerated structure on three vertices such that the second vertex in the enumeration is isomorphic to \mathbf{B} , the third one is isomorphic to \mathbf{U} and there are no non-unary relations in \mathbf{A} . Clearly, \mathbf{A} is \mathcal{F} -free.

For a copy $\widetilde{\mathbf{A}} \subseteq \mathbf{H}$, we denote by $n(\widetilde{\mathbf{A}})$ the level of the first vertex and we put $s(\widetilde{\mathbf{A}}) = s(v) \upharpoonright_{j(w)}$ where w is the second and v is the third vertex of \mathbf{A} . We define a colouring $c: \binom{\mathbf{H}}{\mathbf{A}} \rightarrow \omega$ putting

$$c(\widetilde{\mathbf{A}}) = \begin{cases} w(s \upharpoonright_{\ell(s,n)}) - n & \text{if } n \leq |s|, \text{ where } s = s(\widetilde{\mathbf{A}}) \text{ and } n = n(\widetilde{\mathbf{A}}), \\ 0 & \text{otherwise.} \end{cases}$$

Let $f: \mathbf{H} \rightarrow \mathbf{H}$ be an arbitrary tree-like embedding of \mathbf{H} to \mathbf{H} . We will prove that $c[\binom{f[\mathbf{H}]}{\mathbf{A}}] = \omega$. Given $p \in \omega$, we will construct a copy $\widetilde{\mathbf{A}} \in \binom{f[\mathbf{H}]}{\mathbf{A}}$ with $c(\widetilde{\mathbf{A}}) = p$. Assume that the vertex set of \mathbf{H} is ω and let $r_0 < r_1$ be arbitrary vertices of \mathbf{H} such that r_0 is isomorphic to \mathbf{B} and r_1 is isomorphic to \mathbf{U} . Let $r_2 \in \omega$ be chosen such that:

1. $r_2 > r_1$,
2. r_2 has the same unary relation as the first vertex of \mathbf{A} , and
3. $f(r_2)$ is the n -th vertex of \mathbf{H} for some n with $n > w(s(f(r_1)) \upharpoonright_{j(f(r_0))})$.

Such a vertex exists, because the infinite disjoint union of vertices with the unary is \mathcal{F} -free, and so \mathbf{H} contains infinitely many such vertices. Let $r_3 \in \omega$ be an

arbitrary vertex such that $r_3 > r_2$, $j(f(r_3)) > n$, r_3 is isomorphic to \mathbf{B} and there are no relations on $\{r_2, r_3\}$.

Put $q = n - w(s(f(r_1)) \upharpoonright_{j(f(r_0))}) - 1 + p$, $X = \{r_0, r_1, r_2, r_3\}$ and $i = 1$. Pick any $x \in \omega$ such that $x > r_3$, $\{r_0, x\}$ is isomorphic to \mathbf{C}_q and there are no relations on $\{r_2, x\}$ and $\{r_3, x\}$. Since f is tree-like, we get $y \in \omega$ such that $y > r_3$,

$$\text{tp}_{\mathbf{H}}(\{f(r_0), f(r_2), f(r_3), f(y)\}) = \text{tp}_{\mathbf{H}}(\{r_0, r_2, r_3, x\})$$

and

$$\text{tp}_{\mathbf{H}}(\{0, \dots, f(r_0) - 1, f(y)\}) = \text{tp}_{\mathbf{H}}(\{0, \dots, f(r_0) - 1, f(r_1)\}).$$

Put $v = f(y)$ and observe that v satisfies the following:

1. $\{f(r_0), v\}$ is isomorphic to \mathbf{C}_p , so in particular v is isomorphic to \mathbf{U} ,
2. there are no relations on $\{f(r_2), v\}$ and $\{f(r_3), v\}$,
3. $v > f(r_3)$ and thus $b(v) > j(f(r_3)) > n$,
4. $s(f(r_1)) \upharpoonright_{j(f(r_0))} \sqsubseteq s(v)$,
5. $s(v)_{j(f(r_0))} = q$, and
6. \mathbf{H} induces a copy of \mathbf{A} on $\{f(r_2), f(r_3), v\}$.

Let $\widetilde{\mathbf{A}}$ be the copy of \mathbf{A} induced on $\{f(r_2), f(r_3), v\}$. Put

$$\begin{aligned} s &= s(\widetilde{\mathbf{A}}) \\ &= s(v) \upharpoonright_{j(f(r_3))} \\ &= s(f(r_1)) \upharpoonright_{j(f(r_0))} \widehat{q} \widehat{s'} \end{aligned}$$

for some sequence s' and note that $n(\widetilde{\mathbf{A}}) = n$ (which was defined when we chose r_2). By the choice of r_3 we know that $|s| > n$ and so $\ell(s, n)$ is defined. By the choice of n we know that

$$w(s(f(r_1)) \upharpoonright_{j(f(r_0))}) < n$$

and from the choice of q we have that

$$w(s(f(r_1)) \upharpoonright_{j(f(r_0))}) + q + 1 = n + p,$$

hence $\ell(s, n) = j(f(r_0)) + 1$. Thus indeed

$$c(\widetilde{\mathbf{A}}) = w(s \upharpoonright_{j(f(r_0))+1}) - n = p.$$

□

9.7 Conclusion

Being the first positive big Ramsey result for random structures in infinite languages, this paper helps locate the boundaries of finiteness of big Ramsey degrees. However, there still remain a lot of open problems even regarding unrestricted structures.

While we have proved finiteness of big Ramsey degrees, we do not know their exact values or descriptions. We believe that this is an important problem (which is open even for 3-uniform hypergraphs), but at the same time it is tractable. We think it makes sense to start with L -hypergraphs with no unary relations, where we expect the main combinatorial difficulties will already present themselves, while it will hopefully avoid some technical complications.

We believe that for finite L the problem might be easier, or at least the description of the exact big Ramsey degrees might be simpler. For this reason we state it as a separate problem.

Problem 9.7.1. *Characterise the exact big Ramsey degrees of countable universal L -hypergraphs for finite L with no unary relations.*

Note that big Ramsey degrees are a property of the bi-embeddability type, hence all countable universal L -hypergraphs have the same big Ramsey degrees.

Problem 9.7.2. *Characterise the exact big Ramsey degrees of countable universal L -hypergraphs for L which has no unary relations and has finitely many relations of every arity.*

The next natural step seems to be adding unary relations.

Problem 9.7.3. *Characterise the exact big Ramsey degrees of countable universal L -hypergraphs for L which has at most countably many unary relations and has finitely many relations of every arity greater than one.*

Finally, one can ask for the exact variant of Theorem 9.1.1.

Problem 9.7.4. *Characterise the exact big Ramsey degrees for structures considered in Theorem 9.1.1.*

Of course, our result is just a small step towards the grand goal of characterising big Ramsey degrees of all structures (or perhaps, more realistically, all structures with known small Ramsey degrees). At this point, however, it seems that even free amalgamation classes behave in unexpectedly complex ways.

9.7.1 Infinite big Ramsey degrees

In Section 9.6 we proved that one cannot hope to strengthen our methods to prove big Ramsey results with infinitely many binary relations. We are confident that this argument generalises to higher arities and the following problem has a solution (we believe that we have such a strengthening of tree-likeness, it will appear elsewhere).

Problem 9.7.5. *Define a concept of tree-likeness such that all structural embeddings are tree-like and prove an analogue of Corollary 9.6.2 for the cases when there are infinitely many relations of some higher arity.*

However, even for binary relations the situation is interesting. Corollary 9.6.2 shows that with respect to tree-like embeddings, the infinite-edge-coloured random graph has infinite big Ramsey degree for a particular triple. On the other hand, it is indivisible by the standard argument (try to embed it into one colour class, if it fails then all realisations of some type are in the other colour class which hence contains a monochromatic copy). We do not know what the big Ramsey degree of an edge of some particular colour is:

Question 9.7.6. *Given the infinite-edge-coloured random graph (Example 9.6.1) and its substructure \mathbf{A} on two vertices, is the big Ramsey degree of \mathbf{A} finite? Is it finite with respect to tree-like embeddings?*

Of course, big Ramsey degrees in general for this structure remain open as well:

Question 9.7.7. *Does the infinite-edge-coloured random graph have finite big Ramsey degrees?*

A positive answer to this question would likely require developing new non-tree-like methods for finding oligochromatic copies, which would be a major event for the area. Note that this shows that big Ramsey degrees are much more subtle than small Ramsey degrees, because when we only want to find a finite oligochromatic structure, the whole problem only touches finitely many colours and hence reduces to the finite-language problem.

In fact, it is even more subtle than this: Consider a countable infinite-edge-coloured graph \mathbf{G} such that edges containing the i -th vertex only use the first i colours. This is a countable infinite-edge-coloured graph which is universal for all *finite* infinite-edge-coloured graphs as well as all countable finitely-edge-coloured graphs, but it does not embed the infinite-edge-coloured random graph. At the same time, an application of the Milliken theorem on the tree which branches $n + 1$ times on level n shows that this graph has finite big Ramsey degrees, and the Laflamme–Sauer–Vuksanovic arguments [LSV06] give the exact big Ramsey degrees for this structure (which in turn recovers the exact big Ramsey degrees of all finite-edge-coloured random graphs).

We conjecture that the answers to these questions are negative. In fact, we believe that Theorem 9.1.1 is tight if there are finitely many unary relations (see the following paragraphs why one has to assume this):

Conjecture 9.7.8. *Let L be a relational language with finitely many unary relations and infinitely many relations of some arity $a \geq 2$. Let \mathbf{H} be an unrestricted L -structure realising all relations from L . Then there is a finite L -structure \mathbf{A} whose big Ramsey degree in \mathbf{H} is infinite. Moreover, the number of vertices of \mathbf{A} only depends on a .*

If one allows infinitely many unary relations, there are new related structures with yet different behaviour. Consider, for example, a language with infinitely many unary relations and the universal structure where each vertex is in exactly one unary. One can define binary relations on this structure by the pair of unaries on the respective vertices. This structure has infinitely many binary relations but it has finite big Ramsey degrees by Theorem 9.1.1 (in fact, to prove it it suffices

to repeat the proof of Theorem 9.5.1 on ω which has all big Ramsey degrees equal to one by the Ramsey theorem).

Or consider language L with unary relations U^1, \dots and binary relations R^1, \dots , let every vertex have exactly one unary relation, only allow any binary relations between vertices from the same unary relation, and only allow binary relations R^1, \dots, R^i between vertices from unary U^i . In other words, look at the disjoint union of infinitely many edge-coloured random graphs such that the i -th of them has all vertices in the unary U^i and has i colours. This structure is unrestricted, but Theorem 9.1.1 does not capture it because of the infinitely many binary relations. However, for a fixed finite \mathbf{A} , we can restrict ourselves to the substructure induced on the unaries which appear in \mathbf{A} for which Theorem 9.1.1 can be applied. Since the substructures on different unaries are disjoint, one can then simply add the remaining unaries back.

However, if one allows binary relations even between vertices with different unary relations, but vertices in U^i can only participate in R^1, \dots, R^i then the argument from the previous paragraph no longer applies, because we have no guarantee that the oligochromatic copy in finitely many unaries will be generic with respect to the rest, that is, that one will be able to extend it to a full copy. (For example, what might happen is that U^i is present in \mathbf{A} , U^j is not and the monochromatic copy of the restriction will be such that there are no edges between its vertices from U^i and vertices from U^j .)

We believe that such structures still have finite big Ramsey degrees and that in order to prove it, one only needs to use a stronger statement of Theorem 9.1.1 which promises that the oligochromatic copy comes from a product tree which will make it possible to extend this copy to a copy using all unaries. It seems that a proper formulation of Conjecture 9.7.8 should speak about the tree of types having a dense set of vertices on which it branches uniformly.

9.7.2 Small Ramsey degrees and the partite lemma

Note that the proof of Proposition 9.4.1 actually gives something stronger than just a finite number of colours: It proves that the colour of every copy only depends on how it embeds into its envelope (this is usually called the *embedding type*). Part of the job when characterising the exact big Ramsey degrees is constructing embeddings which realise as few embedding types as possible.

It was discovered by Hubička [Hub20a] that even without knowing the exact big Ramsey degrees, one can often use the big Ramsey upper bound to give exact small Ramsey degrees. For unrestricted structures, the argument is as follows: Given a finite (enumerated) structure \mathbf{A} , pick an arbitrary relation $R \in L$ of arity $a \geq 2$ and extend \mathbf{A} to \mathbf{A}' adding a vertex b_v for every $v \in A$, and putting $b_v < w$ for every $v, w \in A$ and $b_v < b_w$ if and only if $v > w$. Finally, add vertices c_1, \dots, c_{a-2} which come very first in the enumeration. \mathbf{A} will be a substructure of \mathbf{A}' and we will only add relations $(c_1, \dots, c_{a-2}, b_v, v) \in R_{\mathbf{A}'}$ for every $v \in A$.

Note that the lexicographic order on A within \mathbf{A}' is the same as the enumeration of \mathbf{A} and that if \mathbf{B} is a substructure of \mathbf{A} then \mathbf{B}' is a substructure of \mathbf{A}' . Moreover, \mathbf{A}' describes one particular embedding type of \mathbf{A} (namely the one where first all vertices branch and only then they are coded, and they branch so that the lex-order coincides with the enumeration).

If we only colour this particular embedding type of \mathbf{A}' , we get a monochromatic subtree. Given any finite \mathbf{B} which contains \mathbf{A} as a substructure, this subtree will contain a copy of (the embedding type described by) \mathbf{B}' in which all copies of (embedding types described by) \mathbf{A}' will be monochromatic. However, as we noted in the previous paragraphs, every copy of \mathbf{A} inside this \mathbf{B} has embedding type described by \mathbf{A}' , hence in fact all copies of \mathbf{A} inside this \mathbf{B} will be monochromatic. By compactness we did not need to find the whole monochromatic subtree, just a finite initial segment of it, which contains a unique copy of \mathbf{C}' for some \mathbf{C} which hence satisfies $\mathbf{C} \rightarrow (\mathbf{B})_{k,1}^{\mathbf{A}}$. Thus we get, in particular, a new proof of the Abramson–Harrington theorem [AH78], or in other words, the Nešetřil–Rödl theorem without any forbidden substructures [NR77b]. (This will appear in full detail elsewhere.)

Using the Abramson–Harrington theorem one can give a simple proof of the partite lemma of Nešetřil and Rödl [NR89]. This goes as follows: We can consider \mathbf{A} -partite structures (say, induced, but the non-induced variant can also be done in this way), as structures with unary marks on vertices which describe the projection to \mathbf{A} . The constraints of \mathbf{A} -partiteness are saying that some particular combination of unaries is forbidden to be in a relation together (i.e. they have a projection to \mathbf{A} where the relation is not present).

Given an \mathbf{A} -partite structure \mathbf{B} let \mathbf{B}^- be its reduct forgetting the unaries. By the Abramson–Harrington theorem there is \mathbf{C}^- such that $\mathbf{C}^- \rightarrow (\mathbf{B}^-)_2^{\mathbf{A}}$. Let \mathbf{C} be an \mathbf{A} -partite structure with vertex set $(C^- \times A)$ where the unaries are given by the projection to A and a tuple $((u_i, x_i))_{i \in n}$ is in a relation $R_{\mathbf{C}}$ if and only if $(u_i)_{i \in n} \in R_{\mathbf{C}^-}$ and $(x_i)_{i \in n} \in R_{\mathbf{A}}$.

A substructure of \mathbf{C} is *transversal* if it does not contain any two vertices of the form $(u, x), (u, y)$ for some $u \in C^-$ and $x \neq y \in A$. Note that a substructure of a transversal structure is again transversal. By a similar argument as in the proof of Theorem 9.5.1 we can show that a colouring of transversal copies of \mathbf{A} in \mathbf{C} gives a transversal copy of \mathbf{B} in which all transversal (and thus actually all) copies of \mathbf{A} have the same colour, thereby proving the partite lemma. Note that conditions of being \mathbf{A} -partite can be described by a set of forbidden substructures which are covered by a relation. From this point of view, Theorem 9.1.1 can be considered as a big Ramsey generalisation of the partite lemma.

Acknowledgments

We would like to thank ChatGPT for improving the abstract, Stevo Todorčević for discussions about the infinite-edge-coloured random graph, and Andy Zucker for helpful comments regarding the presentation of our results.

10. Characterisation of the big Ramsey degrees of the generic partial order

MARTIN BALKO, DAVID CHODOUNSKÝ, NATASHA DOBRINEN, JAN HUBIČKA, MATĚJ KONEČNÝ, LLUIS VENA, ANDY ZUCKER

Abstract As a result of 33 intercontinental Zoom calls, we characterise big Ramsey degrees of the generic partial order. This is an infinitary extension of the well known fact that finite partial orders endowed with linear extensions form a Ramsey class (this result was announced by Nešetřil and Rödl in 1984 with first published proof by Paoli, Trotter and Walker in 1985). Towards this, we refine earlier upper bounds obtained by Hubička based on a new connection of big Ramsey degrees to the Carlson–Simpson theorem and we also introduce a new technique of giving lower bounds using an iterated application of the upper-bound theorem.

10.1 Introduction

Given structures \mathbf{A} and \mathbf{B} , we denote by $\binom{\mathbf{B}}{\mathbf{A}}$ the set of all embeddings from \mathbf{A} to \mathbf{B} . We write $\mathbf{C} \rightarrow (\mathbf{B})_{k,\ell}^{\mathbf{A}}$ to denote the following statement:

For every colouring χ of $\binom{\mathbf{C}}{\mathbf{A}}$ with k colours, there exists an embedding $f: \mathbf{B} \rightarrow \mathbf{C}$ such that χ takes no more than ℓ values on $\binom{f(\mathbf{B})}{\mathbf{A}}$.

For a countably infinite structure \mathbf{B} and a finite substructure \mathbf{A} , the *big Ramsey degree* of \mathbf{A} in \mathbf{B} is the least number $\ell \in \mathbb{N} \cup \{\infty\}$ such that $\mathbf{B} \rightarrow (\mathbf{B})_{k,\ell}^{\mathbf{A}}$ for every $k \in \mathbb{N}$. Similarly, if \mathcal{K} is a class of finite structures and $\mathbf{A} \in \mathcal{K}$, the *small Ramsey degree* of \mathbf{A} in \mathcal{K} is the least number $\ell \in \mathbb{N} \cup \{\infty\}$ such that for every $\mathbf{B} \in \mathcal{K}$ and $k \in \mathbb{N}$ there exists $\mathbf{C} \in \mathcal{K}$ such that $\mathbf{C} \rightarrow (\mathbf{B})_{k,\ell}^{\mathbf{A}}$.

A structure is *homogeneous* if every isomorphism between two of its finite substructures extends to an automorphism. It is well known that up to isomorphism there is a unique homogeneous partial order $\mathbf{P} = (P, \leq_{\mathbf{P}})$ such that every countable partial order has an embedding into \mathbf{P} . The order \mathbf{P} is called the *generic partial order* (see e.g. [Mac11]). We refine the following recent result.

Theorem 10.1.1 (Hubička [Hub20a]). *The big Ramsey degree of every finite partial order in the generic partial order \mathbf{P} is finite.*

We characterise the big Ramsey degrees of (finite substructures of) \mathbf{P} in a similar fashion as the authors did in [BCD⁺21b] for binary finitely constrained free amalgamation classes. Similarly to that case, our characterisation is in terms of a tree-like object we call a “poset diary” where each level has exactly one critical event. It follows that any poset diary which codes the generic poset is a big Ramsey structure for \mathbf{P} (Definition 1.3 of [Zuc19]).

Relative to the small Ramsey degree case (see e.g. [NVT15, Bod15, Hub20b]), there are fewer classes of structures for which big Ramsey degrees are fully understood. The following is the current state of the art:

1. The Ramsey theorem implies that the big Ramsey degree of every finite linear order in the order of ω is 1.
2. In 1979, Devlin refined upper bounds by Laver and characterised big Ramsey degrees of the order of rationals [Dev79, Tod10].
3. In 2006 Laflamme, Sauer and Vuksanović characterised big Ramsey degrees of the Rado (or random) graph and related random structures in binary languages [LSV06]. Actual numbers were counted by Larson [Lar08].
4. In 2008 Nguyen Van Thé characterised big Ramsey degrees of the ultrametric Urysohn spaces [NVT08].
5. In 2010 Laflamme, Nguyen Van Thé and Sauer [LNVT10] characterised the big Ramsey degrees of the dense local order, $\mathbf{S}(2)$ and the \mathbb{Q}_n structures
6. In 2020 Coulson, Dobrinen, and Patel in [CDP22a] and [CDP22b] characterised the big Ramsey degrees of homogeneous binary relational structures which satisfy SDAP^+ , recovering work in [LSV06] and [LNVT10] and characterising the big Ramsey degrees of the ordered versions of structures in [LSV06], the generic n -partite and generic ordered n -partite graphs, and the $(\mathbb{Q}_{\mathbb{Q}})_n$ hierarchy.
7. In 2020 Barbosa characterised big Ramsey degrees of the directed circular orders $\mathbf{S}(k)$, $k \geq 2$, [Bar20].
8. In 2020 a characterisation of big Ramsey degrees of the triangle-free Henson graph was obtained by Dobrinen [Dob20b] and independently by the remaining authors of this article.
9. Joining efforts, the authors were able to fully characterise big Ramsey degrees of free amalgamation classes in finite binary languages described by finitely many forbidden cliques [BCD⁺21b].

If one draws analogy to the small Ramsey results, many of the aforementioned characterisations can be seen as infinitary generalisations of special cases of the Nešetřil–Rödl theorem [NR77a]. Small Ramsey degrees (or the *Ramsey expansions* satisfying the *expansion property*, see [NVT15]) are known for many other classes including the class of partial orders with linear extensions [NR84, PTW85, Sok12, Maš18, SZ17, NR18], metric spaces [Neš07, DR12, Maš18] and other examples derived by rather general structural conditions [HN19].

In this paper, we characterize big Ramsey degrees of the generic partial order. This class represents an important new example since its finitary counter-part is not a consequence of the Nešetřil–Rödl theorem [NR77a], and it is not of the form covered by [CDP22b]. Towards this direction we needed to further refine the new proof technique for upper bounds, based on an application of the Carlson–Simpson theorem, introduced in [Hub20a], and also find a completely

new approach to lower bounds. The techniques introduced in this paper generalize to other classes as we briefly outline in the final section. However, to keep the paper simple, we did not attempt to state the results in the greatest possible generality.

Our construction uses the following partial order introduced in [Hub20a] (which is closely tied to the order of 1-types within a fixed enumeration of \mathbf{P} discussed in Section 10.3). Let

$$\Sigma = \{\mathbf{L}, \mathbf{X}, \mathbf{R}\}$$

be an *alphabet* ordered by $<_{\text{lex}}$ as

$$\mathbf{L} <_{\text{lex}} \mathbf{X} <_{\text{lex}} \mathbf{R}.$$

We denote by Σ^* the set of all finite words in the alphabet Σ , by \leq_{lex} their lexicographic order, and by $|w|$ the length of the word $w = w_0w_1 \cdots w_{|w|-1}$. We denote the empty word by \emptyset . Given words $w, w' \in \Sigma^*$, we write $w \sqsubseteq w'$ if w is an initial segment of w' .

Definition 10.1.1 (Partial order (Σ^*, \preceq)). For $w, w' \in \Sigma^*$, we set $w \prec w'$ if and only if there exists i such that:

1. $0 \leq i < \min(|w|, |w'|)$,
2. $(w_i, w'_i) = (\mathbf{L}, \mathbf{R})$, and
3. $w_j \leq_{\text{lex}} w'_j$ for every $0 \leq j < i$.

We call the least i satisfying the condition above the *witness* of $w \prec w'$ and denote it by $i(w, w')$. We say that $w \preceq w'$ if either $w \prec w'$ or $w = w'$.

Proposition 10.1.2 ([Hub20a]). *The pair (Σ^*, \preceq) is a partial order and $(\Sigma^*, \leq_{\text{lex}})$ is a linear extension of it.*

This partial order will serve as the main tool for characterising the big Ramsey degrees of \mathbf{P} . The intuition for its definition is described in Section 10.3.

Proof of Proposition 10.1.2. It is easy to see that \preceq is reflexive and anti-symmetric. We verify transitivity. Let $w \prec w' \prec w''$ and put $i = \min(i(w, w'), i(w', w''))$.

First assume that $i = i(w, w')$. Then we have $w_i = \mathbf{L}, w'_i = \mathbf{R}$ which implies that $w''_i = \mathbf{R}$. For every $0 \leq j < i$ it holds that $w_j \leq_{\text{lex}} w'_j \leq_{\text{lex}} w''_j$. It follows, by the transitivity of \leq_{lex} , that $w \prec w''$ and $i(w, w'')$ exists with $i(w, w'') \leq i$.

Now assume that $i = i(w', w'')$. Then we have $w'_i = \mathbf{L}, w''_i = \mathbf{R}$, and as $w'_i = \mathbf{L}$ we also have that $w_i = \mathbf{L}$. Again, for every $0 \leq j < i$ it holds that $w_j \leq_{\text{lex}} w'_j \leq_{\text{lex}} w''_j$. Similarly as above, it also follows that $w \preceq w''$ and $i(w, w'') \leq i$. \square

Given a word w and an integer $i \geq 0$, we denote by $w|_i$ the *initial segment* of w of length i . For $S \subseteq \Sigma^*$, we let \overline{S} be the set $\{w|_i : w \in S, 0 \leq i \leq |w|\}$. Given $\ell \geq 0$, we let $S_\ell = \{w \in S : |w| = \ell\}$ and call it the *level ℓ of S* . When writing \overline{S}_ℓ , we always mean level ℓ of \overline{S} . A word $w \in S$ is called a *leaf* of S if there is no word $w' \in S$ extending w . Given a word w and a character $c \in \Sigma$, we denote by $w \frown c$ the word created from w by adding c to the end of w . We also set $S \frown c = \{w \frown c : w \in S\}$.

In addition to the partial order \preceq we define the following relation on each Σ_ℓ^* , $\ell \geq 0$, where Σ_ℓ^* denotes the set of words of length ℓ in the alphabet Σ .

Definition 10.1.2 (Partial orders $(\Sigma_\ell^*, \trianglelefteq)$). Given $\ell > 0$ and words $w, w' \in \Sigma_\ell^*$, we write $w \trianglelefteq w'$ if $w_i \leq_{\text{lex}} w'_i$ for every $0 \leq i < \ell$ (this is the usual element-wise partial order). We write $w \perp w'$ if w and w' are \trianglelefteq -incomparable (that is, if neither of $w \trianglelefteq w'$ nor $w' \trianglelefteq w$ holds). We call w and w' *related* if one of expressions $w \preceq w'$, $w' \preceq w$ or $w \perp w'$ holds, otherwise they are *unrelated*.

Intuitively, $s \trianglelefteq t$ describes those pairs of nodes on the same level where it is possible to extend s and t to nodes with $s' \prec t'$. However, observe that (somewhat counter-intuitively) it can happen that both \preceq and \perp hold at the same time: For example, we have both that $\text{LR} \preceq \text{RL}$ as well as $\text{LR} \perp \text{RL}$. Later, we will introduce the notion of compatibility to help us handle these situations. While it is easy to check that $(\Sigma_\ell^*, \trianglelefteq)$ is a partial order for every $\ell \geq 0$, we do not extend it to all of Σ^* .

To characterise big Ramsey degrees of \mathbf{P} , we introduce the following technical definition, our main theorem then characterizes the big Ramsey degree of a given poset in \mathbf{P} as the number of poset-diaries of a certain kind. This definition has a good intuition behind it which will be explained in Section 10.3.

Definition 10.1.3 (Poset-diaries). A set $S \subseteq \Sigma^*$ is called a *poset-diary* if no member of S extends any other (i.e., S is an antichain in (Σ^*, \sqsubseteq)) and precisely one of the following four conditions is satisfied for every level ℓ with $0 \leq \ell < \sup_{w \in S} |w|$:

1. **Leaf:** There is $w \in \bar{S}_\ell$ related to every $u \in \bar{S}_\ell \setminus \{w\}$ and

$$\bar{S}_{\ell+1} = (\bar{S}_\ell \setminus \{w\}) \frown X.$$

2. **Splitting:** There is $w \in \bar{S}_\ell$ such that

$$\begin{aligned} \bar{S}_{\ell+1} = & \{z \in \bar{S}_\ell : z <_{\text{lex}} w\} \frown X \\ & \cup \{w \frown X, w \frown R\} \\ & \cup \{z \in \bar{S}_\ell : w <_{\text{lex}} z\} \frown R. \end{aligned}$$

3. **New \perp :** There are unrelated words $v <_{\text{lex}} w \in \bar{S}_\ell$ such that the following is satisfied.

(A) For every $u \in \bar{S}_\ell$, $v <_{\text{lex}} u <_{\text{lex}} w$ implies that at least one of $u \perp v$ or $u \perp w$ holds.

Moreover:

$$\begin{aligned} \bar{S}_{\ell+1} = & \{z \in \bar{S}_\ell : z <_{\text{lex}} v\} \frown X \\ & \cup \{v \frown R\} \\ & \cup \{z \in \bar{S}_\ell : v <_{\text{lex}} z <_{\text{lex}} w \text{ and } z \perp v\} \frown X \\ & \cup \{z \in \bar{S}_\ell : v <_{\text{lex}} z <_{\text{lex}} w \text{ and } z \not\perp v\} \frown R \\ & \cup \{w \frown X\} \\ & \cup \{z \in \bar{S}_\ell : w <_{\text{lex}} z\} \frown R. \end{aligned}$$

4. **New \prec :** There are unrelated words $v <_{\text{lex}} w \in \bar{S}_\ell$ such that the following is satisfied.




Figure 10.1: Possible levels in poset-diaries




Figure 10.2: Diaries of \mathbf{A}_2 .

- (B1) For every $u \in \bar{S}_\ell$ such that $u <_{\text{lex}} v$, at least one of $u \preceq w$ or $u \perp v$ holds.
- (B2) For every $u \in \bar{S}_\ell$ such that $w <_{\text{lex}} u$, at least one of $v \preceq u$ or $w \perp u$ holds.

Moreover:

$$\begin{aligned}
\bar{S}_{\ell+1} = & \{z \in \bar{S}_\ell : z <_{\text{lex}} v \text{ and } z \perp v\} \frown X \\
& \cup \{z \in \bar{S}_\ell : z <_{\text{lex}} v \text{ and } z \not\perp v\} \frown L \\
& \cup \{v \frown L\} \\
& \cup \{z \in \bar{S}_\ell : v <_{\text{lex}} z <_{\text{lex}} w\} \frown X \\
& \cup \{w \frown R\} \\
& \cup \{z \in \bar{S}_\ell : w <_{\text{lex}} z \text{ and } w \perp z\} \frown X \\
& \cup \{z \in \bar{S}_\ell : w <_{\text{lex}} z \text{ and } w \not\perp z\} \frown R.
\end{aligned}$$

See also Figure 10.1.

Given a countable partial order \mathbf{Q} , we let $T(\mathbf{Q})$ be the set of all poset-diaries S such that (S, \preceq) is isomorphic to \mathbf{Q} .

Example 10.1.1. Denote by \mathbf{A}_n the anti-chain with n vertices and by \mathbf{C}_n the chain with n vertices.

$$\begin{aligned}
T(\mathbf{A}_1) &= T(\mathbf{C}_1) = \{\emptyset\}, \\
T(\mathbf{A}_2) &= \{\{XR, RXX\}, \{XRX, RX\}\}, \\
T(\mathbf{C}_2) &= \{\{XL, RRX\}, \{XLX, RR\}\}.
\end{aligned}$$

In all diaries in $T(\mathbf{A}_2) \cup T(\mathbf{C}_2)$, level 0 does splitting, level 1 adds a new \perp or \prec , and levels 2 and 3 are leaves.

Poset-diaries of small partial orders can be determined by an exhaustive search tool.¹ We determined that $|T(\mathbf{C}_3)| = 52$, $|T(\mathbf{C}_4)| = 11000$ and $|T(\mathbf{A}_3)| = 84$, $|T(\mathbf{A}_4)| = 75642$. Overall there are:

¹<https://github.com/janhubicka/big-ramsey>




Figure 10.3: Poset-diaries of \mathbf{C}_2 .

1. 1 poset-diary of the (unique) partial order of size 1: $S = \{\emptyset\}$,
2. 4 poset-diaries of partial orders of size 2: $T(\mathbf{A}_2) \cup T(\mathbf{C}_2)$,
3. 464 poset-diaries of partial orders of size 3,
4. 1874880 poset-diaries of partial orders of size 4.

As our main result, we determine the big Ramsey degrees of \mathbf{P} and show that \mathbf{P} admits a big Ramsey structure; while we refer to [Zuc19] for the precise definition, a big Ramsey structure for \mathbf{P} is an expansion \mathbf{P}^* of \mathbf{P} which encodes the exact big Ramsey degrees for all finite substructures simultaneously in a coherent fashion.

Theorem 10.1.3. *For every finite partial order \mathbf{Q} , the big Ramsey degree of \mathbf{Q} in the generic partial order \mathbf{P} equals $|T(\mathbf{Q})| \cdot |\text{Aut}(\mathbf{Q})|$. Furthermore, any $\mathbf{P}^* \in T(\mathbf{P})$ is a big Ramsey structure for \mathbf{P} . Consequently, the topological group $\text{Aut}(\mathbf{P})$ admits a metrizable universal completion flow.*

Note that the number of poset-diaries is multiplied by the size of the automorphism group since we define big Ramsey degrees with respect to embeddings (as done, for example, in [Zuc19]). Big Ramsey degrees are often defined with respect to substructures (see, for example, [Dev79, LSV06, Lar08, NVT09, LNVTS10]) and in that case the degree would be $|T(\mathbf{Q})|$.

Example 10.1.2.

$$\begin{aligned}
 |T(\mathbf{A}_1)| \cdot |\text{Aut}(\mathbf{A}_1)| &= |T(\mathbf{C}_1)| \cdot |\text{Aut}(\mathbf{C}_1)| = 1, \\
 |T(\mathbf{A}_2)| \cdot |\text{Aut}(\mathbf{A}_2)| &= 4, \\
 |T(\mathbf{C}_2)| \cdot |\text{Aut}(\mathbf{C}_2)| &= 2, \\
 |T(\mathbf{A}_3)| \cdot |\text{Aut}(\mathbf{A}_3)| &= 504, \\
 |T(\mathbf{C}_3)| \cdot |\text{Aut}(\mathbf{C}_3)| &= 52, \\
 |T(\mathbf{A}_4)| \cdot |\text{Aut}(\mathbf{A}_4)| &= 1816128, \\
 |T(\mathbf{C}_4)| \cdot |\text{Aut}(\mathbf{C}_4)| &= 11000.
 \end{aligned}$$

10.2 Preliminaries

10.2.1 Relational structures

We use the standard model-theoretic notion of relational structures. Let L be a language with relational symbols $R \in L$ each equipped with a positive natural

number called its *arity*. An L -structure \mathbf{A} on A is a structure with *vertex set* A and relations $R_{\mathbf{A}} \subseteq A^r$ for every symbol $R \in L$ of arity r . If the set A is finite we call \mathbf{A} a *finite structure*. We consider only structures with finitely many or countably infinitely many vertices.

Since we work with structures in multiple languages, we will list the vertex set along with the relations of the structure, e.g., (P, \leq) for partial orders.

10.2.2 Trees

For us, a *tree* is a (possibly empty) partially ordered set $(T, <_T)$ such that, for every $t \in T$, the set $\{s \in T : s <_T t\}$ is finite and linearly ordered by $<_T$. All nonempty trees we consider are *rooted*, that is, they have a unique minimal element called the *root* of the tree. An element $t \in T$ of a tree T is called a *node* of T and its *level*, denoted by $|t|_T$, is the size of the set $\{s \in T : s <_T t\}$. Note that the root has level 0. Given a tree T and nodes $s, t \in T$, we say that s is a *successor* of t in T if $t \leq_T s$. A *subtree* of a tree T is a subset T' of T viewed as a tree equipped with the induced partial ordering.

Given words $w, w' \in \Sigma^*$, we write $w \sqsubseteq w'$ if w is an initial segment of w' . With this partial order we obtain tree (Σ^*, \sqsubseteq) and the notation on words introduced in Section 10.1 can be viewed as a special case of the notation introduced here.

10.2.3 Parameter words

To obtain the upper bounds on big Ramsey degrees of \mathbf{P} we apply a Ramsey theorem for parameter words which we briefly review now.

Given a finite alphabet Σ and $k \in \omega \cup \{\omega\}$, a k -*parameter word* is a (possibly infinite) string W in the alphabet $\Sigma \cup \{\lambda_i : 0 \leq i < k\}$ containing all symbols $\lambda_i : 0 \leq i < k$ such that, for every $1 \leq j < k$, the first occurrence of λ_j appears after the first occurrence of λ_{j-1} . The symbols λ_i are called *parameters*. We will use uppercase characters to denote parameter words and lowercase characters for words without parameters. Let W be an n -parameter word and let U be a parameter word of length $k \leq n$, where $k, n \in \omega \cup \{\omega\}$. Then $W(U)$ is the parameter word created by *substituting* U to W . More precisely, $W(U)$ is created from W by replacing each occurrence of λ_i , $0 \leq i < k$, by U_i and truncating it just before the first occurrence of λ_k in W .

We apply the following Ramsey theorem for parameter words, which is an easy consequence of the Carlson–Simpson theorem [CS84], see also [Tod10, Kar13]:

Theorem 10.2.1. *Let Σ be a finite alphabet. If Σ^* is coloured with finitely many colours, then there exists an infinite-parameter word W such that $W[\Sigma^*] = \{W(s) : s \in \Sigma^*\}$ is monochromatic.*

10.3 Tree of 1-types

Poset-diaries, which can be compared to *Devlin embedding types* (see Chapter 6.3 of [Tod10]) or the *diagonal diaries* introduced in [BCD⁺21b], have an intuitive meaning when understood in the context of the tree of 1-types of \mathbf{P} . We now

introduce this tree and its enrichment to an *aged coding tree* and discuss how poset-diaries can be obtained as a suitable abstraction of the aged coding tree.

An *enumerated structure* is simply a structure \mathbf{A} whose underlying set is $|\mathbf{A}|$. Fix a countably infinite enumerated structure \mathbf{A} . Given vertices u, v and an integer n satisfying $\min(u, v) \geq n \geq 0$, we write $u \sim_n^{\mathbf{A}} v$, and say that u and v are of the same (quantifier-free) type over $0, 1, \dots, n-1$, if the structure induced by \mathbf{A} on $\{0, 1, \dots, n-1, u\}$ is identical to the structure induced by \mathbf{A} on $\{0, 1, \dots, n-1, v\}$ after renaming vertex v to u . We write $[u]_n^{\mathbf{A}}$ for the $\sim_n^{\mathbf{A}}$ -equivalence class of vertex u .

Definition 10.3.1 (Tree of 1-types). Let \mathbf{A} be a countably infinite (relational) enumerated structure. Given $n < \omega$, write $\mathbb{T}_{\mathbf{A}}(n) = \omega / \sim_n^{\mathbf{A}}$. A (quantifier-free) *1-type* is any member of the disjoint union $\mathbb{T}_{\mathbf{A}} := \bigsqcup_{n < \omega} \mathbb{T}_{\mathbf{A}}(n)$. We turn $\mathbb{T}_{\mathbf{A}}$ into a tree as follows. Given $x \in \mathbb{T}_{\mathbf{A}}(m)$ and $y \in \mathbb{T}_{\mathbf{A}}(n)$, we declare that $x \leq_{\mathbb{T}_{\mathbf{A}}}^{\mathbb{T}_{\mathbf{A}}} y$ if and only if $m \leq n$ and $x \supseteq y$.

In the case that we have a Fraïsséclass \mathcal{K} in mind (which for us will always be the class of finite partial orders), we can extend the definition to a finite enumerated $\mathbf{A} \in \mathcal{K}$ as follows. Fix an enumerated Fraïssélimit \mathbf{K} of \mathcal{K} which has \mathbf{A} as an initial segment. We then set $\mathbb{T}_{\mathbf{A}} = \mathbb{T}_{\mathbf{K}}(<|\mathbf{A}|)$. This does not depend on the choice of \mathbf{K} .

In the case that \mathbf{A} is a structure in a finite binary relational language, we can encode $\mathbb{T}_{\mathbf{A}}$ as a subtree of $k^{<\omega}$ for some $k < \omega$ as follows. Given two enumerated structures \mathbf{B} and \mathbf{C} , an *ordered embedding* of \mathbf{B} into \mathbf{C} is any embedding of \mathbf{B} into \mathbf{C} which is an increasing injection of the underlying sets. Write $\text{OEmb}(\mathbf{B}, \mathbf{C})$ for the set of ordered embeddings of \mathbf{B} into \mathbf{C} . Fix once and for all an enumeration $\{\mathbf{B}_i : i < k\}$ of the set of enumerated structures of size 2 which admit an enumerated embedding into \mathbf{A} . Given $x \in \mathbb{T}_{\mathbf{A}}(m)$, we define $\sigma(x) \in k^m$ where given $j < m$, we set $x(j) = i$ iff for some (equivalently every) $n \in x$ there is $f \in \text{OEmb}(\mathbf{B}_i, \mathbf{A})$ with $\text{Im}(f) = \{m, n\}$. The map $\sigma : \mathbb{T}_{\mathbf{A}} \rightarrow k^{<\omega}$ is then an embedding of trees. We write $\text{CT}^{\mathbf{A}} = \sigma[\mathbb{T}_{\mathbf{A}}]$ and call this the *coding tree* of \mathbf{A} . Typically we also endow $\text{CT}^{\mathbf{A}}$ with *coding nodes*, where for each n , the n^{th} coding node of $\text{CT}^{\mathbf{A}}$ is defined to be $c^{\mathbf{A}}(n) := \sigma([n]_n^{\mathbf{A}})$.

Understanding of the tree of 1-types of a given structure is useful for constructing unavoidable colorings as well as for showing upper bounds on big Ramsey degrees; see for instance [LSV06, Dob20a, Dob23, CDP22a, CDP22b, Zuc22, Hub20a, BCH⁺19, BCH⁺22, BCD⁺21b]. We therefore fix an enumerated generic partial order \mathbf{P} , and we put $(\mathbb{T}, \leq_{\mathbb{T}}) = (\mathbb{T}_{\mathbf{P}}, \leq_{\mathbb{T}_{\mathbf{P}}})$ and $\text{CT} = \text{CT}^{\mathbf{P}}$. By identifying the symbols $\{L, X, R\}$ with $\{0, 1, 2\}$, we can identify CT as a subtree of (Σ^*, \sqsubseteq) . More concretely, given $x \in \mathbb{T}(m)$ and $j < m$, we have:

$$\sigma(x)_j = \begin{cases} L & \text{if } a <_{\mathbf{P}} j \text{ for every (some) } a \in x, \\ R & \text{if } j <_{\mathbf{P}} a \text{ for every (some) } a \in x, \\ X & \text{otherwise.} \end{cases}$$

We remark that CT is a proper subset of Σ^* , no matter the enumeration we choose. For example if $LR \in \text{CT}$ then $RL \notin \text{CT}$ since that would imply an existence of vertices $a, b \in \mathbf{P}$ such that $a <_{\mathbf{P}} 0 <_{\mathbf{P}} b$ and $b <_{\mathbf{P}} 1 <_{\mathbf{P}} a$, contradicting the fact that \mathbf{P} is a partial order.




Figure 10.4: Initial part of the tree of types of an enumerated linear order $(\mathbf{Q}, \leq_{\mathbf{Q}})$ (left) and of the enumerated partial order \mathbf{P} (right). The bold node on each level corresponds to the coding node.

The relations \prec , $<_{\text{lex}}$ and \perp , introduced in Section 10.1, capture the following properties of types:

Proposition 10.3.1. *Let $x \in \mathbb{T}(m)$ and $y \in \mathbb{T}(n)$ be 1-types of \mathbf{P} .*

- (1) *If there exist $a \in x$ and $b \in y$ satisfying $a <_{\mathbf{P}} b$, then for every $\ell < \min(m, n)$ it holds that $\sigma(x)_{\ell} \leq_{\text{lex}} \sigma(y)_{\ell}$.*
- (2) *If $\sigma(x) \prec \sigma(y)$, then for every $a \in x$ and $b \in y$ it holds that $a \leq_{\mathbf{P}} b$.*
- (3) *If $\sigma(x) <_{\text{lex}} \sigma(y)$, then for every $a \in x$ and $b \in y$ it holds that $b \not\leq_{\mathbf{P}} a$.*
- (4) *If $m = n$ and $\sigma(x) \perp \sigma(y)$, then for every $a \in x$ and $b \in y$ it holds that a and b are $\leq_{\mathbf{P}}$ incomparable.*

Proof. We first verify (1) by contrapositive. Assume there is $\ell < \min(m, n)$ such that $\sigma(y)_{\ell} <_{\text{lex}} \sigma(x)_{\ell}$. First consider the case that $(\sigma(x)_{\ell}, \sigma(y)_{\ell}) = (\mathbf{X}, \mathbf{L})$. It follows for any $a \in x$ and $b \in y$ that a is $\leq_{\mathbf{P}}$ -incomparable with ℓ and $b <_{\mathbf{P}} \ell$. It follows that we cannot have $a <_{\mathbf{P}} b$. The arguments in the cases $(\sigma(x)_{\ell}, \sigma(y)_{\ell}) = (\mathbf{R}, \mathbf{L})$ and $(\sigma(x)_{\ell}, \sigma(y)_{\ell}) = (\mathbf{R}, \mathbf{X})$ are similar.

To see (2) observe that $\sigma(x) \prec \sigma(y)$ implies the existence of a vertex $\ell \in \mathbf{P}$ satisfying $\ell < \min(m, n)$ and $(\sigma(x)_{\ell}, \sigma(y)_{\ell}) = (\mathbf{L}, \mathbf{R})$. It follows that for any $a \in x$ and $b \in y$ we have $a <_{\mathbf{P}} \ell <_{\mathbf{P}} b$.

To verify (3) observe that there exists $\ell < \min(m, n)$ such that $\sigma(x)_{\ell} <_{\text{lex}} \sigma(y)_{\ell}$. Thus we cannot have $a \in x$, $b \in y$ such that $b <_{\mathbf{P}} a$, as this would contradict (1).

Finally to verify (4) observe that $\sigma(x) \perp \sigma(y)$ implies the existence of vertices $k, \ell \in \mathbf{P}$ satisfying $\max(k, \ell) < \min(m, n)$, $\sigma(x)_k <_{\text{lex}} \sigma(y)_k$ and $\sigma(y)_{\ell} <_{\text{lex}} \sigma(x)_{\ell}$. Hence the existence of $a \in x$ and $b \in y$ with either $a <_{\mathbf{P}} b$ or $b <_{\mathbf{P}} a$ contradicts (1). \square

The main difficulty while working with the tree $(\mathbb{T}, \leq_{\mathbb{T}})$ is the fact that it depends on the choice of an enumeration of \mathbf{P} . For this reason we will focus on the tree (Σ^*, \sqsubseteq) which can be seen as an amalgamation of all possible trees $(\mathbb{T}, \leq_{\mathbb{T}})$ constructed using all possible enumerations of \mathbf{P} . The next definition captures the main properties of words in CT which are independent of the choice of enumeration of \mathbf{P} .

Definition 10.3.2 (Compatibility). Words $u \leq_{\text{lex}} v \in \Sigma^*$ are *compatible* if the following two conditions are satisfied:

- (1) there is no $\ell < \min(|u|, |v|)$ such that $(u_\ell, v_\ell) = (\mathbf{R}, \mathbf{L})$, and
- (2) if there exists $\ell' < \min(|u|, |v|)$ such that $(u_{\ell'}, v_{\ell'}) = (\mathbf{L}, \mathbf{R})$, then for every $\ell'' < \min(|u|, |v|)$ it holds that $u_{\ell''} \leq_{\text{lex}} v_{\ell''}$.

Proposition 10.3.2. *For every $s, t \in \text{CT}$ it holds that s, t are compatible.*

Proof. Suppose $x, y \in \mathbb{T}$ are such that $\sigma(x) <_{\text{lex}} \sigma(y)$. To see property (1) of Definition 10.3.2, suppose there were $\ell < \min(i, j)$ such that $(\sigma(x)_\ell, \sigma(y)_\ell) = (\mathbf{R}, \mathbf{L})$. This implies that $b <_{\mathbf{P}} \ell \leq_{\mathbf{P}} a$ for every $a \in x$ and $b \in y$ which contradicts Proposition 10.3.1 (3).

Property (2) of Definition 10.3.2 is a consequence of Proposition 10.3.1 (1) and the fact that the existence of ℓ' such that $(\sigma(x)_{\ell'}, \sigma(y)_{\ell'}) = (\mathbf{L}, \mathbf{R})$ implies that $a <_{\mathbf{P}} \ell <_{\mathbf{P}} b$ for every $a \in x$ and $b \in y$. \square

In [Zuc22], using ideas implicit in the parallel 1's of [Dob20a] and pre- a -cliques of [Dob23], levels of coding trees are endowed with the structure of *aged sets*. This means that for every $m < \omega$, every set $S \subseteq \text{CT}(m)$ is equipped with a class of finite $|S|$ -labeled structures describing exactly which finite structures can be coded by coding nodes above the members of S . For the generic partial order, it will be useful to encode this information slightly differently than in [Zuc22, BCD⁺21b], in particular since we want to do this on all of Σ^* , not just on CT.

Definition 10.3.3 (Level structure). Given $\ell \geq 0$ and $S \subseteq \Sigma_\ell^*$, the *level structure* is the structure $\mathbf{S} = (S, \leq_{\text{lex}}, \preceq, \trianglelefteq)$.

Proposition 10.3.3. *For every $\ell > 0$ and $S \subseteq \Sigma_\ell^*$, the structure $\mathbf{S} = (S, \leq_{\text{lex}}, \preceq, \trianglelefteq)$ satisfies the following properties:*

- (P1) (S, \preceq) is a partial order.
- (P2) (S, \trianglelefteq) is a partial order.
- (P3) (S, \leq_{lex}) is a linear order.
- (P4) For every $u, v \in S$ it holds that $u \preceq v \implies u \leq_{\text{lex}} v$ (\leq_{lex} is a linear extension of \preceq).
- (P5) For every $u, v \in S$ it holds that $u \trianglelefteq v \implies u \leq_{\text{lex}} v$ (\leq_{lex} is a linear extension of \trianglelefteq).
- (P6) For every $u, v, w \in S$ it holds that $u \preceq v \trianglelefteq w \implies u \preceq w$ and $u \trianglelefteq v \preceq w \implies u \preceq w$.

Moreover if all words in S are compatible then

(P7) For every $u, v \in S$ it holds that $u \preceq v \implies u \trianglelefteq v$.

Proof. Properties (P1) and (P4) are Proposition 10.1.2. (P2), (P3) (P5) and (P6) follow directly from the definitions. (P7) is Definition 10.3.2 (2). \square

Level structures can be understood as approximations of a given partial order with a given linear extension from below (using order \preceq) and from above (using \trianglelefteq). This is a natural analog of the age-set structure in CT.

Remark 10.3.1. One can, perhaps surprisingly, prove that the class \mathcal{K} of all finite structures satisfying properties (P1), (P2), \dots , (P7) is an amalgamation class. As a consequence of the construction from Section 10.5 one gets that for each structure $\mathbf{A} \in \mathcal{K}$ there exists $\ell > 0$ and $S \subseteq \Sigma_\ell^*$ such that $\mathbf{S} = (S, \leq_{\text{lex}}, \preceq, \trianglelefteq)$ is isomorphic to \mathbf{A} .

Remark 10.3.2. This interesting phenomenon of constructing a class of approximations (or, using the terminology of [Zuc22, BCD⁺21b], the class of all possible aged sets that can appear on some level of the coding tree) of a given amalgamation class exists in other cases. For binary free amalgamation classes, this approximation class corresponds to the union $\bigcup_\rho P(\rho)$, where the union is taken over all possible *sorts* ρ (see [BCD⁺21b] for the definitions). However, the theory of aged coding trees and the sets $P(\rho)$ can be defined for any strong amalgamation class in a finite binary language. Note that while for free amalgamation classes the set $P(\rho)$ is always closed under intersections, this need not be the case in general (indeed, it fails for posets).

Another key difference between the free amalgamation case and partial orders is that for free amalgamation classes, we can arrange so that going up and left (that is, by a non-relation) in the coding tree is a safe move, i.e., is an embedding of the level structure from one level to another. Indeed, if this were true for the generic partial order and the coding tree CT we fixed earlier, one could prove upper bounds for the big Ramsey degrees using forcing arguments much as is done for the free amalgamation case in [Zuc22]. However, while a weakening of the idea of a “safe direction” does hold for partial orders (see Proposition 10.6.3), the proof of Lemma 3.4 from [Zuc22] breaks in the setting of the generic partial order. However, it is possible that the coding tree Milliken theorem still holds. Below, $\text{AEmb}(\text{CT}^{\mathbf{A}}, \text{CT})$ refers to the set of *aged embeddings* of the coding tree $\text{CT}^{\mathbf{A}}$ into CT, the strong similarity maps that respect coding nodes and level structures (see Definition 2.3 of [Zuc22]).

Question 10.3.4. Fix a finite partial order \mathbf{A} . Let $r < \omega$ and let

$$\gamma: \text{AEmb}(\text{CT}^{\mathbf{A}}, \text{CT}) \rightarrow r$$

be a coloring. Is there $h \in \text{AEmb}(\text{CT}, \text{CT})$ such that $h \circ \text{AEmb}(\text{CT}^{\mathbf{A}}, \text{CT})$ is monochromatic?

10.4 Poset-diaries and level structures

Given a poset-diary S , one can view \bar{S} as a binary branching tree and each level \bar{S}_ℓ as a structure $\mathbf{S}_\ell = (\bar{S}_\ell, \leq_{\text{lex}}, \preceq, \trianglelefteq)$ where the structure $\mathbf{S}_{\ell+1}$ is constructed from the structure \mathbf{S}_ℓ as described in the following proposition.

Proposition 10.4.1. *Let S be a poset-diary. Then all words in \bar{S} are mutually compatible, and for each $\ell \leq \sup_{w \in S} |w|$ the structures $\mathbf{S}_\ell = (\bar{S}_\ell, \leq_{\text{lex}}, \preceq, \triangleleft)$ and $\mathbf{S}_{\ell+1} = (\bar{S}_{\ell+1}, \leq_{\text{lex}}, \preceq, \triangleleft)$ are related as follows:*

1. *If \bar{S}_ℓ introduces a new leaf, then $\mathbf{S}_{\ell+1}$ is isomorphic to \mathbf{S}_ℓ with one vertex removed.*
2. *If \bar{S}_ℓ is splitting, then $\mathbf{S}_{\ell+1}$ is isomorphic to \mathbf{S}_ℓ with one vertex v duplicated to v', v'' with $v' <_{\text{lex}} v'', v' \not\preceq v'', v'' \not\preceq v',$ and $v' \triangleleft v''$.*
3. *If \bar{S}_ℓ has a new \perp , then $\mathbf{S}_{\ell+1}$ is isomorphic to \mathbf{S}_ℓ with one pair removed from relation \triangleleft (and thus one pair added to \perp).*
4. *If \bar{S}_ℓ has a new \preceq , then $\mathbf{S}_{\ell+1}$ is isomorphic to \mathbf{S}_ℓ extended by one pair in relation \prec .*

To prove Proposition 10.4.1, we note the following easy observation.

Observation 10.4.2. *Let $u <_{\text{lex}} v \in \Sigma_i^*$ for some $i \geq 0$ and $c, c' \in \Sigma$ such that $c \leq_{\text{lex}} c'$ and $(c, c') \neq (L, R)$. Then*

1. $u \preceq v \iff v \hat{\ } c \preceq u \hat{\ } c',$
2. $u \perp v \iff u \hat{\ } c \perp v \hat{\ } c',$
3. *if u and v are compatible then $u \hat{\ } c$ and $v \hat{\ } c'$ are compatible.*

Proof of Proposition 10.4.1. Fix a poset-diary S and level $\ell < \sup_{w \in S} |w|$ and consider individual cases.

1. Leaf vertex w : We have that $|\mathbf{S}_\ell| = |\mathbf{S}_{\ell+1}| + 1$ since w is the only vertex of \mathbf{S}_ℓ not extended to a vertex in $\mathbf{S}_{\ell+1}$. The desired isomorphism and mutual compatibility follows by Observation 10.4.2.
2. Splitting of vertex w : Here vertex w is the only vertex with two extensions. The desired isomorphism and mutual compatibility follows again by Observation 10.4.2.
3. New $v \perp w$: Since $v <_{\text{lex}} w$ are unrelated and thus $v \triangleleft w$, we know that $v \hat{\ } R$ and $w \hat{\ } X$ are compatible and $v \hat{\ } R \perp w \hat{\ } X$ holds. Since we extended by letters X and R we know that there are no new pairs in relation \preceq .

Now assume, for contradiction, that there is $u \in \bar{S}_\ell \setminus \{v, w\}$ unrelated to v but where the extension of u in $\bar{S}_{\ell+1}$ is related to $v \hat{\ } R \in \bar{S}_{\ell+1}$. Since all words lexicographically before v are extended by X and all words lexicographically after w by R , by Observation 10.4.2, we conclude that $v <_{\text{lex}} u <_{\text{lex}} w$ and u extends by X . Mutual compatibility follows by analogous argument.

4. New $v \preceq w$: Since $v <_{\text{lex}} w$ are unrelated, we know that $v \hat{\ } L$ and $w \hat{\ } R$ are compatible and $v \hat{\ } L \preceq w \hat{\ } R$ holds. To see that no additional pair to relation \perp was introduced, observe that for $u, u' \in \bar{S}_i$, $(u, u') \neq (v, w)$ to be extended to $u \hat{\ } L$, $u' \hat{\ } R$ we have, by assumptions (B1) and (B2),

$u \preceq v$ and $w \preceq u'$. Since $v <_{\text{lex}} w$ is unrelated we also have $v \triangleleft v$. By Proposition 10.3.3 (P6) $u \preceq v <_{\text{lex}} w \implies u \preceq v$ and thus also $u \preceq u'$.

It remains to consider the possibility that new pairs are added to relation \perp . We again consider individual cases.

First consider the case that u is unrelated to v but their extensions are newly in \perp . Since v extends by L we know that $u <_{\text{lex}} v$ and v extends by X. This contradicts construction of \overline{S}_{i+1} .

The case that u is unrelated to w but their extensions are newly in \perp follows by symmetry.

It thus remains to consider the case where $u <_{\text{lex}} u'$, $u, u' \notin \{v, w\}$, are unrelated in \overline{S}_i , however their extensions are related in \overline{S}_{i+1} . It is not possible for u to extend by R and u' by L. So assume that u extends by X and v extends by L (the remaining case follows by symmetry). From this we conclude that $u <_{\text{lex}} u' <_{\text{lex}} v$, $u \perp v$ and $u' \not\perp v$. Since $u \not\perp u'$ we again obtain a contradiction with Proposition 10.3.3 (P4) or (P6).

Mutual compatibility follows by analogous argument. □

10.5 A poset-diary coding \mathbf{P}

Recall that $\mathbf{P} = (\omega, \leq_{\mathbf{P}})$ denotes a fixed enumerated generic poset. We define a function $\varphi: \omega \rightarrow \Sigma^*$ by mapping $j < \omega$ to a word w of length $2j + 2$ defined by putting $(w_{2j}, w_{2j+1}) = (\text{L}, \text{R})$ and, for every $i < j$, (w_{2i}, w_{2i+1}) to (L, L) if $j \leq_{\mathbf{P}} i$, (R, R) if $i \leq_{\mathbf{P}} j$ and (X, X) otherwise. We set $T = \overline{\varphi[\omega]}$. The following result is easy to prove by induction.

Proposition 10.5.1 (Proposition 4.7 of [Hub20a]). *The function φ is an embedding $\mathbf{P} \rightarrow (\Sigma^*, \preceq)$. Moreover, $\varphi(v)$ is a leaf of T for every $v \in \mathbf{P}$, all words in T are mutually compatible, and if $v, w \in \mathbf{P}$ are incomparable, we have $\varphi(v) \perp \varphi(w)$.*

We will need the following refinement of this embedding.

Theorem 10.5.2. *There exists an embedding $\psi: \mathbf{P} \rightarrow (\Sigma^*, \preceq)$ such that $\psi[\omega]$ is a poset-diary.*

Proof. Fix the embedding φ as above and put $T = \overline{\varphi[\omega]}$. We proceed by induction on levels of T . For every ℓ , we define an integer N_ℓ and a function $\psi_\ell: T_\ell \rightarrow \Sigma_{N_\ell}^*$. We will maintain the following invariants:

1. The set $\overline{\psi_\ell[T_\ell]}$ satisfies the conditions of Definition 10.1.3 for all levels with the exception of $N_\ell - 1$.
2. If $\ell > 0$, then, for every $u \in T_\ell$, the word $\psi_\ell(u)$ extends $\psi_{\ell-1}(u|_{\ell-1})$.

We let $N_0 = 0$ and put ψ_0 to map the empty word to the empty word. Now, assume that $N_{\ell-1}$ and $\psi_{\ell-1}$ are already defined. We inductively define a sequence of functions $\psi_\ell^i: T_\ell \rightarrow \Sigma_{N_{\ell-1}+i}^*$. Put $\psi_\ell^0(u) = \psi_{\ell-1}(u|_{\ell-1})$. Now, we proceed in steps. At step j , apply the first of the following constructions that can be applied and terminate the procedure if none of them applies:

1. If ψ_ℓ^{j-1} is not injective, let $w \in T_\ell$ be lexicographically least with $\psi_\ell^{j-1}(w) = \psi_\ell^{j-1}(x)$ for some $x \in T_\ell \setminus \{w\}$. Given $u \in T_\ell$, set $\psi_\ell^j(u) = \psi_\ell^{j-1}(u) \frown X$ if $u \leq_{\text{lex}} w$, and set $\psi_\ell^j(u) = \psi_\ell^{j-1}(u) \frown R$ if $w <_{\text{lex}} u$. Then this satisfies the conditions on new splitting at $\psi_\ell^{j-1}(w)$ as given in Definition 10.1.3.
2. If there are words w and w' from T_ℓ with $w <_{\text{lex}} w'$ such that $w \perp w'$ and $\psi_\ell^{j-1}(w) \not\perp \psi_\ell^{j-1}(w')$ and condition (A) of Definition 10.1.3 is satisfied for the value range of ψ_ℓ^{j-1} , we construct ψ_ℓ^j to satisfy the conditions on new \perp for $\psi_\ell^{j-1}(w)$ and $\psi_\ell^{j-1}(w')$ as given by Definition 10.1.3.
3. If there are words w and w' from T_ℓ with $w <_{\text{lex}} w'$ such that $w \prec w'$ and $\psi_\ell^{j-1}(w) \not\prec \psi_\ell^{j-1}(w')$ and conditions (B1) and (B2) of Definition 10.1.3 are satisfied for the value range of ψ_ℓ^{j-1} , we construct ψ_ℓ^j to satisfy the conditions on new \prec for $\psi_\ell^{j-1}(w)$ and $\psi_\ell^{j-1}(w')$ as given by Definition 10.1.3.

Let J be the largest index for which for which ψ_ℓ^J is defined.

Claim 10.5.3. ψ_ℓ^J is an isomorphism $(T_\ell, \leq_{\text{lex}}, \preceq, \trianglelefteq) \rightarrow (\psi_\ell^J[T_\ell], \leq_{\text{lex}}, \preceq, \trianglelefteq)$.

Proof of claim. Suppose, to the contrary, that this is not true. If ψ_ℓ^J is not a bijection, this means that there are $w, w' \in T_\ell$ such that $\psi_\ell^J(w) = \psi_\ell^J(w')$. But then the conditions in (1) are satisfied, a contradiction with maximality of J . So ψ_ℓ^J is a bijection. Note that the steps of the construction ensure that ψ_ℓ^J respects $<_{\text{lex}}$. We also have $\psi_\ell^J(w) \perp \psi_\ell^J(w') \implies w \perp w'$ and $\psi_\ell^J(w) \preceq \psi_\ell^J(w') \implies w \preceq w'$ for $w, w' \in T_\ell$.

If there are $w, w' \in T_\ell$ such that $w <_{\text{lex}} w'$, $w \perp w'$ and $\psi_\ell^J(w) \not\perp \psi_\ell^J(w')$, pick among all such pairs one minimizing $|\{u \in T_\ell : w <_{\text{lex}} u \leq_{\text{lex}} w'\}|$. Proposition 10.3.3 implies that the conditions in (2) are satisfied, again a contradiction with maximality of J .

So there are $w, w' \in T_\ell$ such that $w <_{\text{lex}} w'$, $w \prec w'$ and $\psi_\ell^J(w) \not\prec \psi_\ell^J(w')$, and we can assume that w, w' maximize $|\{u \in T_\ell : w <_{\text{lex}} u \leq_{\text{lex}} w'\}|$. Proposition 10.3.3 implies that the conditions in (2) are satisfied, again a contradiction with maximality of J . Hence indeed ψ_ℓ^J is an isomorphism $(T_\ell, \leq_{\text{lex}}, \preceq, \trianglelefteq) \rightarrow (\psi_\ell^J[T_\ell], \leq_{\text{lex}}, \preceq, \trianglelefteq)$. \square

Finally, we put $N_\ell = |\psi_\ell^J(w)|$ for some $w \in T_\ell$ and $\psi_\ell = \psi_\ell^J$. Once all functions ψ_ℓ are constructed, we can set $\psi(i) = \psi_{2i+2}(\varphi(i))$. It is easy to verify that this is an embedding $\mathbf{P} \rightarrow (\Sigma^*, \preceq)$ such that $\psi[\omega]$ is a poset-diary (if it was not it fails at some finite level ℓ but the construction ensures that every level adheres to the conditions of Definition 10.1.3). \square

10.6 Interesting levels and sub-diaries

We now aim to prove the upper-bound for big Ramsey degrees of \mathbf{P} . Towards this direction we need to define a notion of sub-diary which corresponds to a subtree of Σ^* which preserves all important features of a given subset. Given $S \subseteq \Sigma^*$ we first determine which levels contain interesting changes and then define a sub-tree by removing the remaining ‘‘boring’’ levels from the tree. This is related to the notion of parameter-space envelopes used in [Hub20a], but sharper, making it possible to get upper bounds tight.

Definition 10.6.1 (Interesting levels). Given $S \subseteq \Sigma^*$, we call a level \bar{S}_i *interesting* if

1. the structure $\bar{\mathbf{S}}_i = (\bar{S}_i, \leq_{\text{lex}}, \preceq, \trianglelefteq)$ is not isomorphic to $\bar{\mathbf{S}}_{i+1} = (\bar{S}_{i+1}, \leq_{\text{lex}}, \preceq, \trianglelefteq)$, or
2. there exist incompatible $u, v \in \bar{S}_{i+1}$ such that $u|_i$ and $v|_i$ are compatible, or
3. there is $u \in S$ with $|u| = i$.

Remark 10.6.1. Interesting levels are the analog for subsets of Σ^* of the notion of *critical level* for a subset of coding nodes in CT; see for instance Definition 5.1 of [Zuc22] or Definition 5.1.3 of [BCD⁺21b] (we note that the two definitions are slightly different).

Given $S \subseteq \Sigma^*$ and levels $\ell < \ell'$, we call a level ℓ' a *duplicate* of ℓ if S contains no word of length ℓ and moreover for every $u \in S$ of length greater than ℓ' it holds that $u_\ell = u_{\ell'}$. By checking definitions of \leq_{lex} , $<_{\text{lex}}$, \preceq and \trianglelefteq one can derive the following simple result.

Observation 10.6.1. *For every $S \subseteq \Sigma^*$ and every $\ell < \ell'$ such that the level ℓ' is duplicate of ℓ it holds that ℓ' is not interesting.*

Definition 10.6.2 (Embedding types). Let $I(S)$ be the set of all interesting levels in S . Let $\tau_S: S \rightarrow \Sigma^*$ be the mapping assigning to each $w \in S$ the word created from w by deleting all characters with indices not in $I(S)$. Define $\tau(S) = \tau_S[S]$ and call it the *embedding type* of S .

The following observation is a direct consequence of Definition 10.1.3.

Observation 10.6.2. *For a poset-diary S and $S' \subseteq S$, $\tau(S')$ is a poset-diary.* \square

We therefore call $\tau(S')$ the *sub-diary* induced by $S' \subseteq S$.

Definition 10.6.3. Recall that for a set $A = \{u^0 <_{\text{lex}} u^1 <_{\text{lex}} \cdots <_{\text{lex}} u^{n-1}\} \subseteq \Sigma_\ell^*$ (for some $\ell > 0$) and word $e \in \Sigma_n^*$ we put $A \frown e = \{u^i \frown e_i : 0 \leq i < n\}$. We call word $e \in \Sigma_n^*$ a *boring extension* of A if level ℓ of $(S \frown e, \leq_{\text{lex}}, \preceq, \trianglelefteq)$ is not interesting. We will denote by Π_A the set of all boring extensions of A and by Π_A^* the set of all finite words over the alphabet Π_A . That is, a member of Π_A^* is a sequence $w = (w^0, w^1, \dots, w^{|w|-1})$ such that for every i we have that $w^i \in \Pi_A$.

We first prove two properties of boring extensions.

Proposition 10.6.3. *For every $\ell \geq 0$, set $S \subseteq \Sigma_\ell^*$ of mutually compatible words, and boring extension e of S , there exists a boring extension e' of Σ_ℓ^* such that $S \frown e \subseteq \Sigma_\ell^* \frown e'$.*

Proof. Fix $\ell \geq 0$, $S = \{u^0 <_{\text{lex}} u^1 <_{\text{lex}} \cdots <_{\text{lex}} u^{n-1}\}$, $\Sigma_\ell^* = \{v^0 <_{\text{lex}} v^1 <_{\text{lex}} \cdots <_{\text{lex}} v^{m-1}\}$ and a boring extension e of S . For $u \in S$ denote by $i(u)$ the integer i satisfying $u^i = u$. For $v \in \Sigma_\ell^* \setminus S$ we say that character $c \in \Sigma$ is *safe* for v if for every $0 \leq j < n$ such that u^j is compatible with v , it holds that $(\{u^j, v\}, \leq_{\text{lex}}, \preceq, \trianglelefteq)$ is isomorphic to $(\{u^j \frown e_j, v \frown c\}, \leq_{\text{lex}}, \preceq, \trianglelefteq)$.

First we check that for every $v \in \Sigma^* \setminus S$ the set of safe characters for v is non-empty. To see this, consider vertex $v \in \Sigma^* \setminus S$ such that X is not safe for v . In this case there are two options:

1. There is $u \in S$ compatible with v such that $u <_{\text{lex}} v$, $u \not\prec v$ and $e_{i(u)} = \text{R}$. In this case we argue that R is safe for v (in fact, it is the only safe character for v). For this we need to check:
 - (a) For every $w \in S$ compatible with v such that $v <_{\text{lex}} w$ and $v \not\prec w$ we have $e_{i(w)} = \text{R}$. This follows from the fact that $v \leq w$ and this needs to be preserved by the extension.
 - (b) For every $w \in S$ compatible with v such that $v <_{\text{lex}} w$ it holds that $e_{i(w)} \neq \text{L}$. From the condition on S we know that u and w are compatible and $u <_{\text{lex}} v <_{\text{lex}} w$. But then $u \wedge \text{R}$ and $w \wedge \text{R}$ are incompatible, which is in contradiction with the definition of safe extension.
 - (c) For every $w \in S$ compatible with v such that $w <_{\text{lex}} v$ and $v \not\prec w$, we have that $e_{i(w)} \neq \text{L}$. Suppose for a contradiction that $e_{i(w)} = \text{L}$. As w and u are compatible and e is boring, we know that $w \prec u$. But, by Proposition 10.3.3 (P4) and (P6), we know that $w \prec u$, $w \not\prec v$, and $u \perp v$ cannot be all satisfied at the same time, a contradiction.
2. There is $u \in S$ compatible with v such that $v <_{\text{lex}} u$, $v \not\prec u$ and $e_{i(u)} = \text{L}$. In this case we can argue symmetrically to show that L is the only safe character for v .

Now we define e'_j to be $e_{i(v^j)}$ whenever $v^j \in S$; otherwise choose the first character from X, L, R that is safe for v^j .

To verify that e' is boring, consider some $0 \leq i < j \leq m$. First observe that $\Sigma_\ell^* \wedge e'$ contains no words of length ℓ . Also, if (v^i, v^j) are incompatible, then any one-letter extensions will yield an isomorphic level structure. So we may assume that $v^i, v^j \notin S$ and v^i, v^j are compatible. We have:

1. $(e'_i, e'_j) \neq (\text{R}, \text{L})$. Assume the contrary and let $u \in S$ be the vertex that made X unsafe for v^i and $u' \in S$ be the vertex that made X unsafe for v^j . By the same analysis as above we have $u <_{\text{lex}} v^i$, $e_{i(u)} = \text{R}$, $u \not\prec v^i$, $v^j <_{\text{lex}} u'$, $e_{i(u')} = \text{L}$, $v^j \not\prec u'$. From this however we conclude that $u <_{\text{lex}} u'$, $e_{i(u)} = \text{R}$, $e_{i(u')} = \text{L}$ which contradicts the definition of boring extension and our assumption that u and u' are compatible.
2. $v^i \not\prec v^j \implies (e'_i, e'_j) \neq (\text{X}, \text{L})$. Assume the contrary and denote by $u \in S$ the vertex that made X unsafe for v^j . Clearly $v^j <_{\text{lex}} u$, $v^j \not\prec u$ and $e_{i(u)} = \text{L}$. It follows that for every $\ell' < \ell$ we have $v_{\ell'}^i \leq_{\text{lex}} v_{\ell'}^j \leq_{\text{lex}} u_{\ell'}$. From this we have that u and v^i are compatible and $u \not\prec v^i$ which makes X unsafe for v^j . A contradiction.
3. $v^i \not\prec v^j \implies (e'_i, e'_j) \neq (\text{R}, \text{X})$. Assume the contrary and denote by $u \in S$ the vertex that made X unsafe for v^i . Clearly $u <_{\text{lex}} v^j$, $u \not\prec v^i$ and $e_{i(u)} = \text{R}$. It follows that for every $\ell' < \ell$ we have $u_{\ell'} \leq_{\text{lex}} v_{\ell'}^i \leq_{\text{lex}} v_{\ell'}^j$. From this we have that u and v^j are compatible and $u \not\prec v^j$ which makes R unsafe for v^j . A contradiction.
4. $v^i \not\prec v^j \implies (e'_i, e'_j) \neq (\text{L}, \text{R})$. Assume the contrary and let $u \in S$ be the vertex that made X unsafe for v^i and $u' \in S$ be the vertex that made X unsafe for v^j . We have $v^i <_{\text{lex}} u <_{\text{lex}} u' <_{\text{lex}} v^j$ and $e_{i(u)} = \text{L}$, $e_{i(u')} = \text{R}$.

It follows that $v^i \not\preceq u$, $u \preceq u'$, $u' \not\preceq v^j$. Now for every $\ell' < \ell$ we also have $v_{\ell'}^i <_{\text{lex}} u_{\ell'} <_{\text{lex}} u'_{\ell'} <_{\text{lex}} v_{\ell'}^j$. This yields $v^i \preceq v^j$. A contradiction. \square

Proposition 10.6.4. *Let $0 \leq \ell \leq \ell'$ be integers and e be a boring extension of $\Sigma_\ell^* = \{u^0 <_{\text{lex}} u^1 <_{\text{lex}} \cdots <_{\text{lex}} u^{n-1}\}$. Put $\Sigma_{\ell'}^* = \{v^0 <_{\text{lex}} v^1 <_{\text{lex}} \cdots <_{\text{lex}} v^{m-1}\}$ and create the word e' of length m by putting, for every $0 \leq i < m$, $e'_i = e_j$ where j satisfies $v^i|_\ell = u^j$. Then e' is a boring extension of $\Sigma_{\ell'}^*$.*

Proof. The proof is just a straightforward verification of the fact that the relations \leq_{lex} , \preceq and \trianglelefteq are determined by first occurrences of certain combinations of letters which this construction does not change.

If (v^i, v^j) are incompatible, then any one-letter extensions will yield an isomorphic level structure. So we may assume that (v^i, v^j) are compatible.

Let v^i, v^j be compatible words. We then check that the structure $(\{v^i, v^j\}, \leq_{\text{lex}}, \preceq, \trianglelefteq)$ is isomorphic to $(\{v^i \widehat{e}'_i, v^j \widehat{e}'_j\}, \leq_{\text{lex}}, \preceq, \trianglelefteq)$ and that $v^i \widehat{e}'_i$ and $v^j \widehat{e}'_j$ are compatible. If $e'_i = e'_j$ (in particular, this happens whenever $v^i|_\ell = v^j|_\ell$), the result is clear.

So suppose that $e'_i \neq e'_j$; in particular, this implies that $v^i|_\ell \neq v^j|_\ell$. Note that in this case the lexicographic order of v^i and v^j is already determined by their restrictions to level ℓ (that is, $v^i <_{\text{lex}} v^j \iff v^i|_\ell <_{\text{lex}} v^j|_\ell$), hence $(\{v^i|_\ell, v^j|_\ell\}, \leq_{\text{lex}})$, $(\{v^i, v^j\}, \leq_{\text{lex}})$, and $(\{v^i \widehat{e}'_i, v^j \widehat{e}'_j\}, \leq_{\text{lex}})$ are isomorphic. Consequently, compatibility of $v^i \widehat{e}'_i$ and $v^j \widehat{e}'_j$ follows from compatibility of v^i and v^j and compatibility of $v^i|_\ell \widehat{e}'_i$ and $v^j|_\ell \widehat{e}'_j$.

Since $(\{v^i|_\ell, v^j|_\ell\}, \leq_{\text{lex}}, \preceq, \trianglelefteq)$ is isomorphic to $(\{v^i|_\ell \widehat{e}'_i, v^j|_\ell \widehat{e}'_j\}, \leq_{\text{lex}}, \preceq, \trianglelefteq)$ (by the fact that e is a boring extension), we get that either

$$(e'_i, e'_j) \notin \{(L, R), (R, L)\},$$

or \preceq is already defined on $\{v^i|_\ell, v^j|_\ell\}$. In either case it holds that $(\{v^i, v^j\}, \preceq)$ and $(\{v^i \widehat{e}'_i, v^j \widehat{e}'_j\}, \preceq)$ are isomorphic. A similar argument can be done for \trianglelefteq , and hence indeed $(\{v^i, v^j\}, \leq_{\text{lex}}, \preceq, \trianglelefteq)$ is isomorphic to $(\{v^i \widehat{e}'_i, v^j \widehat{e}'_j\}, \leq_{\text{lex}}, \preceq, \trianglelefteq)$, that is, e' is a boring extension of $\Sigma_{\ell'}^*$. \square

10.7 Upper bounds

We prove Ramsey-type theorem for the following the following kind of embeddings.

Definition 10.7.1 (Shape-preserving functions). Given $S \subseteq \Sigma^*$ we call function $f: S \rightarrow \Sigma^*$ *shape-preserving* if $\tau_S(w) = \tau_{f[S]}(f(w))$ for every $w \in S$.

We will generally consider shape-preserving functions only for those sets S such that $S = \tau(S)$ (that is for *embedding types*). However, the next observation follows directly from the definition without this extra assumption:

Observation 10.7.1. *Let $f: S \rightarrow \Sigma^*$ be shape-preserving.*

- (i) *For every shape-preserving $h: f[S] \rightarrow \Sigma^*$ it holds that $h \circ f$ is shape-preserving.*

- (ii) For all $u, v \in S$, $|u| \leq |v|$ implies that $|f(u)| \leq |f(v)|$.
- (iii) For all $u, v \in S$ with $u \sqsubseteq v$, we have $f(u) \sqsubseteq f(v)$.
- (iv) The function f is an embedding $f: (S, \leq_{\text{lex}}, \preceq, \trianglelefteq) \rightarrow (\Sigma^*, \leq_{\text{lex}}, \preceq, \trianglelefteq)$ and images of pairs of compatible words are also compatible.
- (v) Let e be a boring extension of S . Then e is a boring extension of $f[S]$.

Given $S \subseteq \Sigma^*$ and $\ell > 0$ we denote by $S_{\leq \ell} = \cup_{i \leq \ell} S_i$ the set of all words in S of length at most ℓ . We also put $S_{< \ell} = S_{\leq \ell - 1}$.

Remark 10.7.1. It is also possible to observe that if $S = \Sigma_\ell^*$ for some $\ell > 0$, it holds that shape-preserving functions correspond to special strong subtrees as used by the Milliken's tree theorem [Tod10]. However, the converse is not true: For example, the function $f: \Sigma_{< 2}^* \rightarrow \Sigma^*$ mapping the empty word to the empty word, $L \mapsto LL$, $X \mapsto XR$ and $R \mapsto RR$ is not shape-preserving because all three levels 0, 1 and 2 of $f[\Sigma_{< 2}^*]$ are interesting, while it does describe a strong subtree.

Given $S, S' \subseteq \Sigma^*$ we denote by $\text{Shape}(S, S')$ the set of all shape-preserving functions f such that $f[S] \subseteq S'$. Given integer n we also denote by $\text{Shape}_n(S, S')$ the set of all functions in $\text{Shape}(S, S')$ that are the identity when restricted to $S_{< n}$.

Let $S \subseteq \Sigma^*$. For a shape-preserving function $g: S \rightarrow \Sigma^*$, we denote by \tilde{g} the function $\{|w| : w \in S\} \rightarrow \omega$ defined by $\tilde{g}(i) = |g(w)|$ for some $w \in S$, $|w| = i$. (Note that by Observation 10.7.1 (ii) this is uniquely defined.) If S is finite, denote by $\max(S)$ the "last" level S_k where $k = \max_{w \in S} |w|$.

Observation 10.7.2. *Let $S = \bar{S} = \tau(S)$ be a finite subset of Σ^* , $\max(S) = \{u^0 <_{\text{lex}} u^1 <_{\text{lex}} \dots <_{\text{lex}} u^n\}$ and $k = \max_{w \in S} |w|$. There is a one-to-one correspondence between $\text{Shape}(S, \Sigma^*)$ and pairs (g^0, w) where $g^0 \in \text{Shape}(S_{< k}, \Sigma^*)$ and $w \in \Pi_{\max(S)}^*$ (recall Definition 10.6.3):*

1. For every $g \in \text{Shape}(S, \Sigma^*)$ there exists $w \in \Pi_{\max(S)}^*$ such that for every $0 \leq i < n$ it holds that $g(u^i) = g(u^i|_{k-1}) \frown u_{k-1}^i \frown w_i^0 \frown \dots \frown w_i^{|w|-1}$.
2. Conversely also for every $g^0 \in \text{Shape}(S_{< k}, \Sigma)$ and every word $w \in \Pi_{\max(S)}^*$ the function $g': S \rightarrow \Sigma^*$ defined by $g'(w) = g^0(w)$ for $|w| < k$ and

$$g'(u^i) = g^0(u^i|_{k-1}) \frown u_{k-1}^i \frown w_i^0 \frown \dots \frown w_i^{|w|-1}$$

is shape-preserving.

Proof. To see the first statement assume the contrary and let $g \in \text{Shape}(S, \Sigma^*)$ be a function for which there is no $w \in \Pi_{\max(S)}^*$ such that for every $0 \leq i < n$ it holds that

$$g(u^i) = g(u^i|_{k-1}) \frown u_{k-1}^i \frown w_i^0 \frown \dots \frown w_i^{|w|-1}.$$

Among all such functions g choose one which minimizes $\tilde{g}(k)$.

Because $S = \bar{S}$ we know that $\tilde{g}(k-1) \in I[g[S]]$. Because g is shape-preserving it follows that $g(u^i)_{\tilde{g}(k-1)} = u_{k-1}^i$ and $g(u^i)_{\tilde{g}(k-1)+1} = g(u^i|_{k-1}) \frown u_{k-1}^i$. It follows that $\tilde{g}(k) \geq \tilde{g}(k-1) + 2$.

Notice that $\tilde{g}(k) = \tilde{g}(k-1) + 2$: As g is shape-preserving, $I(g(S))$ contains no levels between $\tilde{g}(k-1)$ and $\tilde{g}(k)$ as otherwise we could remove them, getting a counter example g' with smaller $\tilde{g}'(k)$. If $\tilde{g}(k) = \tilde{g}(k-1) + 1$, then we can take $w = \emptyset$, and g would not be a counterexample. But now, observe that level $\tilde{g}(k) + 1$ of $g(S)$ is not interesting and thus corresponds to a boring extension in $\Pi_{\max(S)}^*$, which gives a contradiction.

The second statement follows by Proposition 10.6.4. \square

The following pigeonhole lemma is a consequence of Theorem 10.2.1.

Lemma 10.7.3. *Let $S = \overline{S} = \tau(S)$ be a finite non-empty subset of Σ^* of mutually compatible words containing at least one non-empty word. Put $k = \max_{w \in S} |w|$. Let $g^0 \in \text{Shape}(S_{<k}, \Sigma^*)$. Denote by G set of all $g \in \text{Shape}(S, \Sigma^*)$ extending g^0 and put $K = \tilde{g}^0(k-1)$. Then for every finite coloring $\chi: G \rightarrow \{0, 1, \dots, r-1\}$ there exists $f \in \text{Shape}_{K+1}(\Sigma^*, \Sigma^*)$ such that χ restricted to $\text{Shape}(S, f[\Sigma^*]) \cap G$ is constant.*

Proof. By Observation 10.7.2, the colouring $\chi: G \rightarrow \{0, 1, \dots, r-1\}$ gives rise to a colouring $\chi': \Pi_{\max(S)}^* \rightarrow \{0, 1, \dots, r-1\}$. Apply Theorem 10.2.1 to obtain W such that $W[\Pi_{\max(S)}^*]$ is monochromatic with respect to χ' . In order to avoid special cases in the upcoming construction, we will assume that W is indexed from 1 and not from 0.

For every $u \in \Sigma_{\leq K}^*$ put $f(u) = u$. We will construct the rest of f by induction on levels. Now assume that $f(\Sigma_{i-1}^*)$ is already defined for some $i > K$. Put

$$I = \begin{cases} 0 & \text{if } i = K + 1, \\ \min\{I < \omega : W_I = \lambda_{i-K-2}\} & \text{if } i > K + 1. \end{cases}$$

Let J be the minimal integer such that $W_J = \lambda_{i-K-1}$. Now define a sequence of functions $f^{i'}: \Sigma_i^* \rightarrow \Sigma_{K+I+i'}^*$ for every $I \leq i' < J$. Put $f^I(u) = f(u|_{i-1}) \frown u_{i-1}$ for every $u \in \Sigma_i^*$. Now proceed by induction on i' . Assume that $f^{i'-1}$ is constructed for some $I < i' < J$ and consider two cases:

1. $W_{i'} = \lambda_j$: Put $f^{i'}(u) = f^{i'-1}(u) \frown u_{j+k+1}$ for every $u \in \Sigma_i^*$.
2. $W_{i'} = e$ for some $e \in \Pi_S$: Let e' be the extension of Σ_{K+1}^* given by Proposition 10.6.3 for boring extension e of $g^1(S_k)$, where g^1 is defined by putting $g^1(u) \mapsto g^0(u|_{k-1}) \frown u_{k-1}$ (by Observation 10.7.1 ((v)), boringness of an extension is preserved by shape-preserving functions). Now let e'' be the extension given by Proposition 10.6.4 for extension e' and level i . Enumerate $\Sigma_i^* = \{u^0 <_{\text{lex}} u^1 <_{\text{lex}} \dots <_{\text{lex}} u^{m-1}\}$ and for $u^j \in S_i$, put $f^{i'}(u^j) = f^{i'-1}(u) \frown e''_j$ for every $0 \leq j < m$.

Finally put $f(u) = f^{J-1}(u)$.

Observe that all levels added by the rules 1 and 2 above are uninteresting since they are either constructed from boring extensions or they are duplicates of levels introduced earlier (in the sense of Observation 10.6.1). The last level is interesting because $\tau(\Sigma^*) = \Sigma^*$. From this we get $f \in \text{Shape}_{K+1}(\Sigma^*, \Sigma^*)$.

To see that χ restricted to $\text{Shape}(S, f[\Sigma^*]) \cap G$ is constant, pick an arbitrary $g \in \text{Shape}(S, f[\Sigma^*]) \cap G$. By Observation 10.7.2 we can decompose g to g^0 and

a word $w \in \Pi_{\max(S)}^*$ such that $\chi(g)$ is equal to $\chi'(w)$. From our construction it follows that $w \in W(\Sigma^*)$, and so indeed χ restricted to $\text{Shape}(S, f[\Sigma^*]) \cap G$ is constant. \square

Observation 10.7.4. *For every $S = \tau(S) \subseteq \Sigma^*$ and every $f \in \text{Shape}(S, \Sigma^*)$ there is a unique function $g \in \text{Shape}(\bar{S}, \Sigma^*)$ extending f . It is constructed by putting $g(w|_\ell) = f(w)|_{\tilde{f}(\ell)}$ for every $w \in S$ and $\ell \leq |w|$. Similarly, for every $g \in \text{Shape}(\bar{S}, \Sigma^*)$ it holds that $g \upharpoonright S \in \text{Shape}(S, \Sigma^*)$.*

Notice that g in Observation 10.7.4 is well defined by Observation 10.7.1 (iii).

Theorem 10.7.5. *For every finite set $S = \tau(S) \subseteq \Sigma^*$ of mutually compatible words and every finite coloring $\chi: \text{Shape}(S, \Sigma^*) \rightarrow \{0, 1, \dots, r-1\}$, there exists $f \in \text{Shape}(\Sigma^*, \Sigma^*)$ such that χ restricted to $\text{Shape}(S, f[\Sigma^*])$ is constant.*

Proof. By Observation 10.7.4 we can assume, without loss of generality, that $S = \bar{S}$. We will use induction on $k = \max_{w \in S} |w|$. For $k = 0$ we can interpret χ as coloring of Σ^* , apply Theorem 10.2.1 to obtain a monochromatic infinite-parameter word W , for every $w \in \Sigma^*$ put $f(w) = W(w)$, and observe that it is shape-preserving.

Now fix S such that $k = \max_{w \in S} |w| > 0$ and a finite colouring $\chi: \text{Shape}(S, \Sigma^*) \rightarrow \{0, 1, \dots, r-1\}$. We will make use of the following claim:

Claim 10.7.6. *There exists $h \in \text{Shape}(\Sigma^*, \Sigma^*)$ and a colouring $\chi': \text{Shape}(S_{<k}, \Sigma^*) \rightarrow \{0, 1, \dots, r-1\}$ such that for every finite $g \in \text{Shape}(S, \Sigma^*)$ it holds that $\chi(h \circ g) = \chi'(g \upharpoonright S_{<k})$.*

First we show that Theorem 10.7.5 follows from the claim. Let h and χ' be given by the claim. By induction hypothesis there exists $f' \in \text{Shape}(\Sigma^*, \Sigma^*)$ such that χ' is constant on $f'[\Sigma^*]$. It is easy to check that $f = h \circ f'$ is shape-preserving and χ is constant when restricted to $f[\Sigma^*]$.

It remains to prove the claim. We will obtain h as the limit of the following sequence: Pick an enumeration $\text{Shape}(S_{<k}, \Sigma^*) = \{g^0, g^1, \dots\}$ such that $0 \leq i \leq j$ it holds that $\tilde{g}^i(k-1) \leq \tilde{g}^j(k-1)$. We construct a sequence of shape-preserving functions $f^0, f^1, \dots \in \text{Shape}(\Sigma^*, \Sigma^*)$ such that for every $i > 0$ the following is satisfied:

1. $f^i[\Sigma^*] \subseteq f^{i-1}[\Sigma^*]$ and $f^i(u) = f^{i-1}(u)$ for every $u \in \Sigma_{<\tilde{g}^{i-1}(k-1)+1}^*$.
2. There exists $c^{i-1} \in \{0, 1, \dots, r-1\}$ such that $\chi(f^i \circ g) = c^{i-1}$ for every $g \in \text{Shape}(S, \Sigma^*)$ extending g^{i-1} .

Put f^0 to be the identity $\Sigma^* \rightarrow \Sigma^*$. Now assume that f^{i-1} is already constructed. Consider coloring $\chi^i: \text{Shape}(S, \Sigma^*) \rightarrow \{0, 1, \dots, r-1\}$ defined by $\chi^i(g) = \chi(f^{i-1} \circ g)$. Obtain $h^i \in \text{Shape}_{\tilde{g}^{i-1}(k-1)+1}(\Sigma^*, \Sigma^*)$ by an application of Lemma 10.7.3 for coloring χ^i and function $f^{i-1} \circ g^{i-1}$ (as g^0 in the statement) and put $f^i = f^{i-1} \circ h^i$.

Next we construct the limit shape-preserving h . For every $i > 1$ it holds that $f^i(u) = f^{i-1}(u)$ for all $u \in \Sigma^*$ where $u \leq \tilde{g}^{i-1}(k-1)$. Because $\tilde{g}^{i-1}(k-1)$ is an increasing function of i and there is no upper bound on the length of words in Σ^* it follows that $h(u) = \lim_{i \rightarrow \omega} f^i(u)$ is well-defined for every $u \in \Sigma^*$. Moreover, h is shape-preserving, because the failure of shape-preservation is witnessed on a finite set. We also put $\chi'(g^i) = c^i$ for every $i \in \omega$. \square

Now we are finally ready to prove a big Ramsey result for \mathbf{P} .

Corollary 10.7.7. *For every finite partial order \mathbf{Q} , the big Ramsey degree of \mathbf{Q} in the generic partial order \mathbf{P} is at most $|T(\mathbf{Q})| \cdot |\text{Aut}(\mathbf{Q})|$.*

Proof. Fix a finite partial order \mathbf{Q} and a coloring χ of $\binom{\mathbf{P}}{\mathbf{Q}}$. Choose an arbitrary enumeration $T(\mathbf{Q}) = \{S_0, S_1, \dots, S_{n-1}\}$ and an arbitrary embedding $\eta: (\Sigma^*, \preceq) \rightarrow \mathbf{P}$ (which exists since \mathbf{P} is universal). Observe that for every $i \in n$ it holds that χ and η induce a coloring χ_i of $\text{Shape}(S_i, \Sigma^*)$ by putting $\chi_i(g) = \chi(\eta \circ g[S_i])$. By repeated applications of Theorem 10.7.5 we construct a sequence of functions $\text{Id} = f_0, f_1, \dots, f_n \in \text{Shape}(\Sigma^*, \Sigma^*)$ such that for every $i \in n$ the following is satisfied:

1. $f_{i+1}[\Sigma^*] \subseteq f_i[\Sigma^*]$, and
2. χ_i restricted to $\text{Shape}(S_i, f_{i+1}[\Sigma^*])$ is constant.

Let $\psi: \mathbf{P} \rightarrow (\Sigma^*, \preceq)$ be obtained by Theorem 10.5.2. Observe that $f_n \circ \psi$ is the desired embedding $\mathbf{P} \rightarrow \mathbf{P}$ where color of every $h \in \binom{\mathbf{P}}{\mathbf{Q}}$ depends only on $\tau[\eta[h[Q]]] \in T(\mathbf{Q})$. \square

10.8 The lower bound

Given a finite partial order \mathbf{A} , a *labeled* poset-diary for \mathbf{A} is a pair (S, f) , where $S \subseteq \Sigma^*$ is a poset-diary in $T(\mathbf{A})$ and $f: \mathbf{A} \rightarrow (S, \prec)$ is an isomorphism. Let $T^{\text{lab}}(\mathbf{A})$ denote the set of labeled poset-diaries coding \mathbf{A} . Note that $|T^{\text{lab}}(\mathbf{A})| = |T(\mathbf{A})| \cdot |\text{Aut}(\mathbf{A})|$. Recall the embedding $\psi: \mathbf{P} \rightarrow (\Sigma^*, \preceq)$ constructed in Theorem 10.5.2. We define a function (colouring) $\chi_{\mathbf{A}}: \binom{\mathbf{P}}{\mathbf{A}} \rightarrow T^{\text{lab}}(\mathbf{A})$ by setting $\chi_{\mathbf{A}}(f) = (\tau(\psi \circ f[A]), \tau_{\psi \circ f[A]} \circ \psi \circ f)$ for every $f \in \binom{\mathbf{P}}{\mathbf{A}}$. In this section, we show that $\chi_{\mathbf{A}}$ is a *recurrent coloring* in the following sense: for every $h \in \binom{\mathbf{P}}{\mathbf{P}}$, there is $\phi \in \binom{\mathbf{P}}{\mathbf{P}}$ with $\chi_{\mathbf{A}} \circ h \circ \phi = \chi_{\mathbf{A}}$. This allows us to characterize the exact big Ramsey degrees of \mathbf{P} .

Theorem 10.8.1. *For every finite partial order \mathbf{A} and every $f \in \binom{\mathbf{P}}{\mathbf{P}}$, we have*

$$\left\{ \chi_{\mathbf{A}}[f \circ g] : g \in \binom{\mathbf{P}}{\mathbf{A}} \right\} = T^{\text{lab}}(\mathbf{A}).$$

Furthermore, for every $f \in \binom{\mathbf{P}}{\mathbf{P}}$, there is $\phi \in \binom{\mathbf{P}}{\mathbf{P}}$ with $\chi_{\mathbf{A}} \circ f \circ \phi = \chi_{\mathbf{A}}$.

Remark 10.8.1. Zucker [Zuc19] introduced the notion of *big Ramsey structure*, which is a structure capturing exact big Ramsey degrees for all finite substructures at the same time. Recurrence of the colorings $\chi_{\mathbf{A}}$ implies that any poset diary coding \mathbf{P} is a big Ramsey structure for \mathbf{P} . In fact, this recurrence property is significantly stronger than asserting that poset diaries coding \mathbf{P} are big Ramsey structures; it tells us that any two poset diaries coding the generic poset are bi-embeddable. This allows us to conclude stronger dynamical properties of the group $G := \text{Aut}(\mathbf{P})$ than indicated in [Zuc19], namely that G admits a metrizable universal completion flow which is also a *strong* completion flow (see the discussion

preceding Theorem 8.0.6 of [BCD⁺21b] for definitions). We observe that for any finite partial order \mathbf{B} , $\chi_{\mathbf{B}}$ is completely determined by the colorings $\chi_{\mathbf{A}}$ for finite partial orders \mathbf{A} with $|\mathbf{A}| \leq 4$. Thus we obtain a big Ramsey structure for \mathbf{P} in a finite relational language.

The rest of this section proves Theorem 10.8.1. First we show, by a repeated application of Lemma 10.7.3, that for every $f \in \left(\begin{smallmatrix} (\Sigma^*, \preceq) \\ (\Sigma^*, \preceq) \end{smallmatrix} \right)$ (not necessarily a shape-preserving one) there exists $g \in \text{Shape}(\Sigma^*, \Sigma^*)$ such that f preserves the main features of the tree structure on $g[\Sigma^*]$.

Lemma 10.8.2. *For every $f \in \left(\begin{smallmatrix} (\Sigma^*, \preceq) \\ (\Sigma^*, \preceq) \end{smallmatrix} \right)$ there exists $g \in \text{Shape}(\Sigma^*, \Sigma^*)$ and a sequence $(N_i)_{i \in \omega}$ satisfying:*

1. $N_0 = 0$,
2. for every $u \in \Sigma^*$ it holds that $N_{|u|} \leq |f(g(u))| < N_{|u|+1}$, and,
3. for every $u, v \in \Sigma^*$ and ℓ such that $u|_{\ell} = v|_{\ell}$, it holds that $f(g(u))|_{N_{\ell}} = f(g(v))|_{N_{\ell}}$.

See Figure 10.5.




Figure 10.5: Function $f \circ g$.

Proof. Fix an embedding $f: (\Sigma^*, \preceq) \rightarrow (\Sigma^*, \preceq)$. We define a sequence of shape-preserving functions $(g_i)_{i \in \omega}$ and a sequence $(N_i)_{i \in \omega}$ of integers satisfying, for every $i > 0$, the following three conditions:

- (A) $g_i \in \text{Shape}(\Sigma^*, \Sigma^*)$ and $g_{i-1} \upharpoonright \Sigma^*_{<i} = g_i \upharpoonright \Sigma^*_{<i}$.
- (B) For every $u \in \Sigma^*_{<i}$ it holds that $N_{|u|} \leq |f(g_i(u))| < N_{|u|+1}$.
- (C) For every $u, v \in \Sigma^*$ and $\ell \leq \min(i, |u|, |v|)$ such that $u|_{\ell} = v|_{\ell}$ it holds that $f(g_i(u))|_{N_{\ell}} = f(g_i(v))|_{N_{\ell}}$.

Put g_0 to be identity, $N_0 = 0$ and proceed by induction. That g_{i-1} and N_{i-1} are already constructed for some $i > 0$. Put

$$N_i = \max \left\{ |f(g_{i-1}(u))| : u \in \Sigma^*_{i-1} \right\} + 1.$$

Enumerate $\Sigma^*_i = w^0, w^1, \dots, w^{m-1}$. By induction we will construct a sequence of functions $g_{i-1} = g_i^0, g_i^1, \dots, g_i^m \in \text{Shape}(\Sigma^*, \Sigma^*)$. Assume that g_i^j is constructed.

Put $S_i^j = \overline{\{w^j\}}$ and define a coloring $\chi_i^j(h)$ of $\text{Shape}_i(S_i^j, \Sigma^*)$ by putting $\chi_i^j(h) = g_i^j(h(w^j))|_{N_i}$. Apply Lemma 10.7.3 on χ_i^j and obtain h_i^j . Put $g_i^{j+1} = h_i^j \circ g_i^j$. Finally, put $g_i = g_i^m$.

To see that g_i satisfies (A) note that all h_i^j 's are shape-preserving functions and that they are the identity when restricted to $\Sigma_{<i}^*$. Property (B) follows directly from the choice of N_i . It remains to verify that (C) is satisfied. By (A) it is enough to verify this for $\ell = i$. Let $u, v \in \Sigma^*$ be such that $u|_i = v|_i$ and let c_i^m be the constant value of χ_i^m for m satisfying $w^m = u|_i = v|_i$ on h_i^m . Notice that $|c_i^m| = N_i$ because there are infinitely many images of words extending w^m . It follows that $g_i(u)|_{N_i} = g_i(v)|_{N_i} = c_i^m$.

It remains to put g to be the limit of sequence $(g_i)_{i \in \omega}$. \square

We denote by $d: \Sigma^* \rightarrow \Sigma^*$ the function that repeats every letter 3 times. That is, for every $u \in \Sigma^*$ we put $d(u) = u'$ where $|u'| = 3|u|$ and for every $\ell < |u|$ it holds that $u'_{3\ell} = u'_{3\ell+1} = u'_{3\ell+2} = u_\ell$. Note that d is a shape-preserving function.

Lemma 10.8.3. *Let $f \in \left(\begin{smallmatrix} (\Sigma^*, \preceq) \\ (\Sigma^*, \preceq) \end{smallmatrix}\right)$ be an embedding. Let $g \in \text{Shape}(\Sigma^*, \Sigma^*)$ and sequence $(N_i)_{i \in \omega}$ be given by Lemma 10.8.2. Put $f' = f \circ g \circ d$. Then for every poset-diary S and every $\ell < \sup_{u \in S} |u|$ it holds that $(\overline{S}_\ell, \leq_{\text{lex}}, \preceq, \trianglelefteq)$ is isomorphic to $(\overline{f'[S]}_{N_{3\ell}}, \leq_{\text{lex}}, \preceq, \trianglelefteq)$.*

Proof. We proceed by induction on level ℓ . Since $|\overline{S}_0| = 1$ we have:

$$(\overline{S}_0, \leq_{\text{lex}}, \preceq, \trianglelefteq) = (\overline{f'[S]}_{N_0}, \leq_{\text{lex}}, \preceq, \trianglelefteq).$$

Now assume that $(\overline{S}_\ell, \leq_{\text{lex}}, \preceq, \trianglelefteq)$ is isomorphic to $(\overline{f'[S]}_{N_{3\ell}}, \leq_{\text{lex}}, \preceq, \trianglelefteq)$. We define function μ assigning every $u \in S_{\ell+1}$ word

$$\mu(u) = f'(u)|_{N_{3\ell+3}}.$$

We claim that μ is the isomorphism of $(\overline{S}_{\ell+1}, \leq_{\text{lex}}, \preceq, \trianglelefteq)$ and $(\overline{f'[S]}_{N_{3\ell+3}}, \leq_{\text{lex}}, \preceq, \trianglelefteq)$:

1. $\mu[\overline{S}_{\ell+1}] \subseteq \overline{f'[S]}_{N_{3\ell+3}}$: Recall that $f' = f \circ g \circ d$. For every $u \in \overline{S}_{\ell+1}$ we have $|d(u)| = 3\ell + 3$. Let $v \in S$ be a successor of u . Notice that $d(v)$ is a successor of $d(u)$. Now because g is given by Lemma 10.8.2 we know that both $f'(v)$ and $f'(u)$ are successors of $f'(u)|_{N_{3\ell+3}}$.
2. For every $u \in S$ with $|u| \geq \ell + 1$ it holds that $\mu(u|_{\ell+1}) \sqsubseteq f'(u)$: This follows directly from the fact that g is constructed using Lemma 10.8.2.
3. $\mu[\overline{S}_{\ell+1}] \supseteq \overline{f'[S]}_{N_{3\ell+3}}$: Let $v \in \overline{f'[S]}_{N_{3\ell+3}}$ and choose $u \in S$ such that $v \sqsubseteq f'(u)$. It follows from Lemma 10.8.2 that $|u| \geq \ell + 1$ and $\mu(u|_{\ell+1}) = v$.
4. μ is injective: Assume that there are $u <_{\text{lex}} v$ in $\overline{S}_{\ell+1}$ such that $\mu(u) = \mu(v)$. By the induction hypothesis we then know that level ℓ is splitting and u, v are the splitting words. It follows that $u = w \frown X$ and $v = w \frown R$ for $w = u|_\ell$. Put $w' = d(w) \frown L$ and observe that $d(u) \perp w'$ and $w' \prec d(v)$. Since $N_{3\ell} \leq |f \circ g(w')| < N_{3\ell+1}$ and $f \circ g$ is an embedding we know that there is level ℓ' satisfying $N_{3\ell} \leq \ell' < N_{3\ell+1}$ satisfying $f'(u)_{\ell'} = X$ and $f'(v)_{\ell'} = R$, which gives a contradiction with $\mu(u) = \mu(v)$, as $\mu(u) \sqsubseteq f'(u)$ and $\mu(v) \sqsubseteq f'(v)$.

5. $u \leq_{\text{lex}} v \implies \mu(u) \leq_{\text{lex}} \mu(v)$: By the induction hypothesis we need to consider only case where $u|_{\ell} = v|_{\ell}$. Therefore $\ell + 1$ is a splitting level, $u = w \frown X$ and $v = w \frown R$ for $w = u|_{\ell}$. Let $u' = w \frown XL$ and $v' = w \frown RR$. Since $u' \prec v'$ and thus also $f'(u') \prec f'(v')$, and since $\mu(u) \sqsubseteq f'(u')$ and $\mu(v) \sqsubseteq f'(v')$, we have $\mu(u) \leq_{\text{lex}} \mu(v)$.
6. $\mu(u) \leq_{\text{lex}} \mu(v) \implies u \leq_{\text{lex}} v$: Follows from the previous point and the fact that \leq_{lex} is a linear order.
7. $u \prec v \implies \mu(u) \prec \mu(v)$: Let ℓ' be the minimal level such that $u_{\ell'} = L$ and $v_{\ell'} = R$. Put $w = d(u|_{\ell'}) \frown LR$. Observe that $d(u) \prec w \prec d(v)$. (That is, w is a *witness* of the fact that $d(u) \prec d(v)$.) Observe also that $f'(u) = f(g(d(u))) \preceq f(g(w)) \preceq f(g(d(v))) = f'(v)$. Because $\ell' \leq \ell$ we have $|w| \leq 3\ell + 2$ and thus $|f(g(w))| < N_{3\ell+3}$. It follows that $\mu(u) \prec \mu(v)$.
8. $\mu(u) \prec \mu(v) \implies u \prec v$: Assume that $\mu(u) \prec \mu(v)$ and $u \not\prec v$. Because $\mu(u) <_{\text{lex}} \mu(v)$ we also have $u <_{\text{lex}} v$. Consequently, $u \frown LX \not\prec v \frown XL$, and so $f'(u \frown LX) \not\prec f'(v \frown XL)$. This is a contradiction with $\mu(u) \sqsubseteq f'(u \frown LX)$ and $\mu(v) \sqsubseteq f'(v \frown XL)$.
9. $u \triangleleft v \implies \mu(u) \triangleleft \mu(v)$: Since $u' \prec v' \implies u' \triangleleft v'$, we only need to consider the case where $u \triangleleft v$ and $u \not\prec v$, and so $f'(u) \not\prec f'(v)$. We have $f'(u \frown L) \prec f'(v \frown R)$, and because $\mu(u) \sqsubseteq f'(u) \sqsubseteq f'(u \frown L)$ and $\mu(v) \sqsubseteq f'(v) \sqsubseteq f'(v \frown R)$, it follows that $\mu(u) \triangleleft \mu(v)$.
10. $\mu(u) \triangleleft \mu(v) \implies u \triangleleft v$: If $u \not\triangleleft v$ then there exists a level $\ell' < \ell + 1$ such that $v_{\ell'} <_{\text{lex}} u_{\ell'}$. Similarly as in the previous cases, we can produce a witness of this fact and contradict that $\mu(u) \triangleleft \mu(v)$.

□

Proof of Theorem 10.8.1. Fix $f \in \left(\frac{\mathbf{P}}{\mathbf{P}}\right)$. Let $\psi: \mathbf{P} \rightarrow (\Sigma^*, \preceq)$ be obtained by Theorem 10.5.2. Let $\eta: (\Sigma^*, \preceq) \rightarrow \mathbf{P}$ be an embedding (which exists since \mathbf{P} is universal). Now $\psi \circ f \circ \eta$ is an embedding $(\Sigma^*, \preceq) \rightarrow (\Sigma^*, \preceq)$. Let $g: (\Sigma^*, \preceq) \rightarrow (\Sigma^*, \preceq)$ and $(N_i)_{i \in \omega}$ be obtained by the application of Lemma 10.8.2 on $\psi \circ f \circ \eta$. Put $f' = \psi \circ f \circ \eta \circ g \circ d$. We claim that for every poset-diary S it holds that $\tau(f'[S]) = S$. By Lemma 10.8.3 we know that for every $\ell < \sup_{u \in S} |u|$ it holds that $(\overline{S}_{\ell}, \leq_{\text{lex}}, \preceq, \triangleleft)$ is isomorphic to $(f'[S]_{N_{3\ell}}, \leq_{\text{lex}}, \preceq, \triangleleft)$. By Proposition 10.4.1, for every $0 < \ell < \sup_{u \in S} |u|$ there is only one difference between $(\overline{S}_{\ell-1}, \leq_{\text{lex}}, \preceq, \triangleleft)$ and $(\overline{S}_{\ell}, \leq_{\text{lex}}, \preceq, \triangleleft)$. Consequently there is only one interesting level ℓ' of $f'[S]$ between $N_{3\ell}$ and $N_{3\ell+3}$ and $(f'[S]_{N_{\ell'}}, \leq_{\text{lex}}, \preceq, \triangleleft)$ is isomorphic to $(\overline{S}_{\ell}, \leq_{\text{lex}}, \preceq, \triangleleft)$ while $(f'[S]_{N_{\ell'+1}}, \leq_{\text{lex}}, \preceq, \triangleleft)$ is isomorphic to $(\overline{S}_{\ell+1}, \leq_{\text{lex}}, \preceq, \triangleleft)$. We further note that by the construction of g , the map $\tau_{f'[S]} \circ f'$ must be the identity on S .

We can thus put $\phi = \eta \circ g \circ d$. □

Proof of Theorem 10.1.3. Given a finite poset \mathbf{A} , the fact that the big Ramsey degree of \mathbf{A} is exactly $|T^{\text{lab}}(\mathbf{A})| = |T(\mathbf{A})| \cdot |\text{Aut}(\mathbf{A})|$ follows from Corollary 10.7.7 and Theorem 10.8.1. To conclude that \mathbf{P} admits a big Ramsey structure, we observe that the colorings $\chi_{\mathbf{A}}$ as \mathbf{A} ranges over all finite posets satisfy the hypotheses of Theorem 7.1 from [Zuc19]. Theorem 1.6 from [Zuc19] then shows that $\text{Aut}(\mathbf{P})$ admits a metrizable universal completion flow. □

10.9 Concluding remarks

10.9.1 Comparision to big Ramsey degrees of the order of rationals

It is natural to ask how the characterisation of big Ramsey degrees of partial orders compares to that of linear orders. The big Ramsey degrees of the linear order correspond to *Devlin's diaries* or *types* which, in our setting, can be defined as follows:

Definition 10.9.1 (Devlin diary [Dev79], see also [Tod10], Definition 6.9). A set $S \subseteq \{L, R\}^*$ is called a *Devlin embedding type*, if no member of S extends any other and precisely one of the following two conditions is satisfied for every level ℓ with $0 \leq \ell < \sup_{w \in S} |w|$:

1. **Leaf:** There is $w \in \bar{S}_\ell$ and

$$\bar{S}_{\ell+1} = (\bar{S}_\ell \setminus \{w\}) \frown L.$$

2. **Splitting:** There is $w \in \bar{S}_\ell$ such that

$$\begin{aligned} \bar{S}_{\ell+1} = & \{z \in \bar{S}_\ell : z <_{\text{lex}} w\} \frown L \\ & \cup \{w \frown L, w \frown R\} \\ & \cup \{z \in \bar{S}_\ell : w <_{\text{lex}} z\} \frown R. \end{aligned}$$

When S is a Devlin embedding type, we call \bar{S} a *Devlin tree*. Given $n \in \omega$, we let $T'(n)$ be the set of all Devlin embedding types of size n .

Theorem 10.9.1 (Devlin [Dev79], see also [Tod10]). *For every $n \in \omega$, the big Ramsey degree of (n, \leq) in the order of rationals equals $|T'(n)|$.*

Any infinite Devlin tree whose leaves code the rational order is a big Ramsey structure for the rationals, see also [Zuc19].

Comparing Devlin trees to poset-diaries is thus relatively natural. In a Devlin tree, only two types of events happen: splitting and leaf. In a poset-diary, the splitting event is different. If w splits on its level then $w \frown X$ and $w \frown L$ are incomparable in \preceq . This adds a need for additional two events: new \prec and new \perp , deciding the poset structure between the successors of w .

10.9.2 The triangle-free graph

We outline how the techniques introduced in this paper can yield a particularly compact characterisation of the big Ramsey degrees of the generic triangle-free graph. This is a special case of the main result of [BCD⁺21b] (see Example 6.2.10). However, we now give a more compact description which is similar to Definition 10.1.3.

We fix an alphabet $\Sigma = \{0, 1\}$ and denote by Σ^* the set of all finite words in the alphabet Σ , by \leq_{lex} their lexicographic order, and by $|w|$ the length of the word w (whose characters are indexed by natural numbers starting at 0). We will denote the empty word by \emptyset .

We consider the following triangle-free graph \mathbf{G}^Δ introduced in [Hub20a].

Definition 10.9.2 (Graph \mathbf{G}^Δ).

1. Vertex set G of \mathbf{G}^Δ is Σ^* .
2. Given $u, v \in G$ satisfying $|u| < |v|$ we make u and v adjacent if and only if $v_{|u|} = 1$ and there is no i satisfying $0 \leq i < |u|$ and $u_i = v_i = 1$.
3. There are no other edges in \mathbf{G}^Δ .

Definition 10.9.3 (Relation \perp). Given $u, v \in G$ with $|u| \leq |v|$, we write $u \perp v$ if one of the following is satisfied:

1. There exists i satisfying $0 \leq i < |u|$ and $u_i = v_i = 1$.
2. There is no i satisfying $0 \leq i < |u|$ and $v_i = 1$.
3. There is no i satisfying $0 \leq i < |v|$ and $u_i = 1$.

Let \mathbf{R}_3 denote the generic triangle-free graph. To characterise big Ramsey degrees of \mathbf{R}_3 , we introduce the following technical definition.

Definition 10.9.4 (Triangle-free diaries). A set $S \subseteq \Sigma^*$ is called a *triangle-free-type* if no member of S extends any other and precisely one of the following four conditions is satisfied for every i with $0 \leq i < \sup_{w \in S} |w|$:

1. **Leaf:** There is $w \in \bar{S}_i$ with $w \neq 0^i$ such that for every distinct $u, v \in \{z \in \bar{S}_i \setminus \{w\} : z \not\perp w\}$ it holds that $u \perp v$. Moreover:

$$\bar{S}_{i+1} = \{z \in S_i \setminus \{w\} : z \perp w\} \hat{\ } 0 \cup \{z \in S_i \setminus \{w\} : z \not\perp w\} \hat{\ } 1.$$

2. **Splitting:** There is $w \in \bar{S}_i$ such that

$$\bar{S}_{i+1} = \bar{S}_i \hat{\ } 0 \cup \{w\} \hat{\ } 1.$$

3. **Non-splitting first neighbour:** $0^i \in \bar{S}_i$ and

$$\bar{S}_{i+1} = (S_i \setminus \{0^i\}) \hat{\ } 0 \cup \{0^i\} \hat{\ } 1.$$

4. **New \perp :** There are distinct words $v, w \in \bar{S}_i$ with $0^i \notin \{v, w\}$, $v \not\perp w$ such that

$$\bar{S}_{i+1} = (\bar{S}_i \setminus \{v, w\}) \hat{\ } 0 \cup \{v, w\} \hat{\ } 1.$$

Given a triangle-free graph \mathbf{H} , we let $T^\Delta(\mathbf{H})$ be the set of all triangle-free-types S such that the structure induced by \mathbf{G}^Δ on $L(S)$ is isomorphic to \mathbf{H} . We can now recover the following result of [BCD⁺21b].

Theorem 10.9.2 ([BCD⁺21b]). *For every finite triangle-free graph \mathbf{H} , the big Ramsey degree of \mathbf{H} in the generic triangle-free graph \mathbf{R}_3 is $|T^\Delta(\mathbf{H})| \cdot |\text{Aut}(\mathbf{H})|$.*

11. Big Ramsey degrees and forbidden cycles

MARTIN BALKO, DAVID CHODOUNSKÝ, JAN HUBIČKA, MATĚJ KONEČNÝ,
JAROSLAV NEŠETRIL, LLUÍS VENA

Abstract Using the Carlson–Simpson theorem, we give a new general condition for a structure in a finite binary relational language to have finite big Ramsey degrees.

11.1 Introduction

We consider standard model-theoretic (relational) structures in finite binary languages formally introduced below. Such structures may be equivalently seen as edge-labelled digraphs with finitely many labels, however the notion of structures is more standard in the area. Structures may be finite or countably infinite. Given structures \mathbf{A} and \mathbf{B} , we denote by $\binom{\mathbf{B}}{\mathbf{A}}$ the set of all embeddings from \mathbf{A} to \mathbf{B} . We write $\mathbf{C} \longrightarrow (\mathbf{B})_{k,l}^{\mathbf{A}}$ to denote the following statement: for every colouring χ of $\binom{\mathbf{C}}{\mathbf{A}}$ with k colours, there exists an embedding $f: \mathbf{B} \rightarrow \mathbf{C}$ such that χ does not take more than l values on $\binom{f(\mathbf{B})}{\mathbf{A}}$. For a countably infinite structure \mathbf{B} and its finite substructure \mathbf{A} , the *big Ramsey degree* of \mathbf{A} in \mathbf{B} is the least number $l' \in \mathbb{N} \cup \{\infty\}$ such that $\mathbf{B} \longrightarrow (\mathbf{B})_{k,l'}^{\mathbf{A}}$ for every $k \in \mathbb{N}$. We say that the *big Ramsey degrees of \mathbf{B} are finite* if for every finite substructure \mathbf{A} of \mathbf{B} the big Ramsey degree of \mathbf{A} in \mathbf{B} is finite.

We focus on structures in binary languages L and adopt some graph-theoretic terminology. Given a structure \mathbf{A} and distinct vertices u and v , we say that u and v are *adjacent* if there exists $R \in L$ such that either $(u, v) \in R_{\mathbf{A}}$ or $(v, u) \in R_{\mathbf{A}}$. A structure \mathbf{A} is *irreducible* if any two distinct vertices are adjacent. A sequence $v_0, v_1, \dots, v_{\ell-1}$, $\ell \geq 3$, of distinct vertices of a structure \mathbf{A} is called a *cycle* if v_i is adjacent to v_{i+1} for every $i \in \{0, \dots, \ell - 2\}$ as well as v_0 adjacent to v_{ℓ} . A cycle is *induced* if none of the other remaining pairs of vertices in the sequence is adjacent.

Following [HN19, Section 2], we call a homomorphism $f: \mathbf{A} \rightarrow \mathbf{B}$ (see Section 11.2) a *homomorphism-embedding* if f restricted to any irreducible substructure of \mathbf{A} is an embedding. The homomorphism-embedding f is called a (*strong*) *completion of \mathbf{A} to \mathbf{B}* provided that \mathbf{B} is irreducible and f is injective.

Our main result, which applies techniques developed by the third author in [Hub20a], gives the following condition for a given structure to have finite big Ramsey degrees.

Theorem 11.1.1. *Let L be a finite language consisting of unary and binary symbols, and let \mathbf{K} be a countably-infinite irreducible structure. Assume that every countable structure \mathbf{A} has a completion to \mathbf{K} provided that every induced cycle in \mathbf{A} (seen as a substructure) has a completion to \mathbf{K} and every irreducible substructure of \mathbf{A} of size at most 2 embeds into \mathbf{K} . Then \mathbf{A} has finite big Ramsey degrees.*

This can be seen as a first step towards a structural condition implying bounds on big Ramsey degrees, giving a strengthening of results by Hubička and Nešetřil [HN19] to countable structures.

The study of big Ramsey degrees originates in the work of Laver who in 1969 showed that the big Ramsey degrees of the ordered set of rational numbers are finite [Tod10, Chapter 6]. The whole area has been revitalized recently; see [Dob20a, Hub20a] for references. Our result can be used to identify many new examples of structures with finite big Ramsey degrees. Theorem 1 is particularly fitting to examples involving metric spaces. In particular, the following corollary may be of special interest.

Corollary 11.1.2. *The following structures have finite big Ramsey degrees:*

- (i) *Free amalgamation structures described by forbidden triangles,*
- (ii) *S -Urysohn space for finite distance sets S for which S -Urysohn space exists,*
- (iii) *Λ -ultrametric spaces for a finite distributive lattice Λ [Bra17],*
- (iv) *metric spaces associated to metrically homogeneous graphs of a finite diameter from Cherlin’s list with no Henson constraints [Che22].*

Vertex partition properties of Urysohn spaces were extensively studied in connection to oscillation stability [NVT10] and determining their big Ramsey degrees presented a long standing open problem: Corollary 11.1.2 (i) is a special case of the main result of [Zuc22], (ii) is a strengthening of [Hub20a, Corollary 6.5 (3)], (iii) strengthens [NVT09] and (iv) is a strengthening of [ABWH⁺17b] to infinite structures.

To see these connections, observe that a metric space can be also represented as an irreducible structure in a binary language having one relation for each possible distance. Possible obstacles to completing a structure in this language to a metric space are irreducible substructures with at most 2 vertices and induced *non-metric cycles*. These are cycles with the longest edge of a length exceeding the sum of the lengths of all the remaining edges; see [ABWH⁺17b].

Note that all these proofs may be modified to yield Ramsey classes of finite structures. Thus, for example, (ii) generalizes [HN19, Section 4.3.2].

Our methods yield the following common strengthening of the main results from [Zuc22] and Theorem 11.1.1. To obtain this result, which is going to appear in [BCH⁺21b], we found a new strengthening of the dual Ramsey theorem.

Theorem 11.1.3. *Let L be a finite language consisting of unary and binary symbols, and let \mathbf{K} be a countably-infinite irreducible structure. Assume that there exists $c > 0$ such that every countable structure \mathbf{A} has a completion to \mathbf{K} provided that every induced cycle in \mathbf{A} has a completion to \mathbf{K} and every irreducible substructure of \mathbf{A} of size at most c embeds into \mathbf{K} . Then \mathbf{A} has finite big Ramsey degrees.*

11.2 Preliminaries

A relational language L is a collection of (relational) symbols $R \in L$, each having its *arity*. An L -structure \mathbf{A} on A is a structure with the *vertex set* A and with

relations $R_{\mathbf{A}} \subseteq A^r$ for every symbol $R \in L$ of arity r . If the set A is finite, then we call \mathbf{A} a *finite structure*. A *homomorphism* $f: \mathbf{A} \rightarrow \mathbf{B}$ is a mapping $f: A \rightarrow B$ such that for every $R \in L$ of arity r we have $(x_1, x_2, \dots, x_r) \in R_{\mathbf{A}} \implies (f(x_1), f(x_2), \dots, f(x_r)) \in R_{\mathbf{B}}$. A homomorphism f is an *embedding* if f is injective and the implication above is an equivalence. If the identity is an embedding $\mathbf{A} \rightarrow \mathbf{B}$, then we call \mathbf{A} a *substructure* of \mathbf{B} . In particular, our substructures are always induced.

Hubička [Hub20a] connected big Ramsey degrees to an infinitary dual Ramsey theorem for parameters spaces. We now review the main notions used. Given a finite alphabet Σ and $k \in \omega \cup \{\omega\}$, a *k-parameter word* is a (possibly infinite) string W in the alphabet $\Sigma \cup \{\lambda_i : 0 \leq i < k\}$ containing all symbols $\lambda_i : 0 \leq i < k$ such that, for every $1 \leq j < k$, the first occurrence of λ_j appears after the first occurrence of λ_{j-1} . The symbols λ_i are called *parameters*. Given a parameter word W , we denote its *length* by $|W|$. The letter (or parameter) on index j with $0 \leq j < |W|$ is denoted by W_j . Note that the first letter of W has index 0. A 0-parameter word is simply a *word*. Let W be an n -parameter word and let U be a parameter word of length $k \leq n$, where $k, n \in \omega \cup \{\omega\}$. Then $W(U)$ is the parameter word created by *substituting* U to W . More precisely, $W(U)$ is created from W by replacing each occurrence of λ_i , $0 \leq i < k$, by U_i and truncating it just before the first occurrence of λ_k in W . Given an n -parameter word W and a set S of parameter words of length at most n , we define $W(S) := \{W(U) : U \in S\}$.

We let $[\Sigma] \binom{n}{k}$ be the set of all k -parameter words of length n , where $k \leq n \in \omega \cup \{\omega\}$. If k is finite, then we also define $[\Sigma]^* \binom{\omega}{k} := \bigcup_{k \leq i < \omega} [\Sigma] \binom{i}{k}$. For brevity, we put $\Sigma^* := [\Sigma]^* \binom{\omega}{0}$.

Our main tool is the following infinitary dual Ramsey theorem, which is a special case of the Carlson–Simpson theorem [CS84, Tod10].

Theorem 11.2.1. *Let $k \geq 0$ be a finite integer. If $[\emptyset]^* \binom{\omega}{k}$ is coloured by finitely many colours, then there exists $W \in [\emptyset] \binom{\omega}{k}$ such that $W([\emptyset]^* \binom{\omega}{k})$ is monochromatic.*

Definition 11.2.1 ([Hub20a]). Given a finite alphabet Σ , a finite set $S \subseteq \Sigma^*$ and $d > 0$, we call $W \in [\emptyset]^* \binom{\omega}{d}$ a *d-parametric envelope* of S if there exists a set $S' \subseteq \Sigma^*$ satisfying $W(S') = S$. In such case the set S' is called the *embedding type* of S in W and is denoted by $\tau_W(S)$. If d is the minimal integer for which a d -parameter envelope W of S exists, then we call W a *minimal envelope*.

Proposition 11.2.2 ([Hub20a]). *Let Σ be a finite alphabet and let $k \geq 0$ be a finite integer. Then there exists a finite $T = T(|\Sigma|, k)$ such that every set $S \subseteq \Sigma^*$, $|S| = k$, has a d -parameter envelope with $d \leq T$. Consequently, there are only finitely many embedding types of sets of size k within their corresponding minimal envelopes. Finally, for any two minimal envelopes W, W' of S , we have $\tau_W(S) = \tau_{W'}(S)$.*

We will thus also use $\tau(S)$ to denote the type $\tau_W(S)$ for some minimal W .

11.3 Proof of Theorem 11.1.1

The proof is condensed due to the space limitations, but we believe it gives an idea of fine interplay of all building blocks. Throughout this section we assume that \mathbf{K} and L are fixed and satisfy the assumptions of Theorem 11.1.1. Following ideas from [Hub20a, Section 4.1], we construct a special L -structure \mathbf{G} with finite big Ramsey degrees and then use \mathbf{G} to prove finiteness of big Ramsey degrees for \mathbf{K} .

Lemma 11.3.1. *Let $h: \mathbf{A} \rightarrow \mathbf{B}$ be a homomorphism-embedding. If \mathbf{B} has a completion $c: \mathbf{B} \rightarrow \mathbf{K}$, then there exists a completion $d: \mathbf{A} \rightarrow \mathbf{K}$.*

Proof. It is clearly enough to consider the case where c is the identity and h is surjective and almost identity, that is, there is a unique vertex $v \in A$ such that $h(v) \neq v$. Let \mathbf{B}' be the structure induced by \mathbf{K} on B . We create a structure \mathbf{B}'' from \mathbf{B}' by duplicating the vertex $h(v)$ to v' and leaving $h(v)$ not adjacent to v' . Since \mathbf{B}' is irreducible, it is easy to observe that all induced cycles in \mathbf{B}'' are already present in \mathbf{B}' . By the assumption on \mathbf{K} , there is a completion $c': \mathbf{B}'' \rightarrow \mathbf{K}$. Now, the completion $d: \mathbf{A} \rightarrow \mathbf{K}$ can be constructed by setting $d(v) = c'(v')$ and $d(u) = c'(u)$ for every $u \in A \setminus \{v\}$. \square

We put $\Sigma = \{\mathbf{A} : A = \{0, 1\} \text{ and there exists an embedding } \mathbf{A} \rightarrow \mathbf{K}\}$. For $U \in \Sigma^*$, we will use bold characters to refer to the letters (e.g. \mathbf{U}_0 is the structure corresponding to the first letter of U) to emphasize that Σ consists of structures.

Given $\mathbf{A}, \mathbf{B}, \mathbf{C} \in \Sigma$, there is at most one structure \mathbf{D} with the vertex set $\{u, v, w\}$ satisfying the following three conditions: (i) mapping $0 \mapsto u, 1 \mapsto v$ is an embedding $\mathbf{A} \rightarrow \mathbf{D}$, (ii) the mapping $0 \mapsto v, 1 \mapsto w$ is an embedding $\mathbf{B} \rightarrow \mathbf{D}$, and (iii) the mapping $0 \mapsto u, 1 \mapsto w$ is an embedding $\mathbf{C} \rightarrow \mathbf{D}$. If such a structure \mathbf{D} exists, we denote it by $\Delta(\mathbf{A}, \mathbf{B}, \mathbf{C})$ (since $\mathbf{A}, \mathbf{B}, \mathbf{C}$ form a triangle). Otherwise we leave $\Delta(\mathbf{A}, \mathbf{B}, \mathbf{C})$ undefined.

Definition 11.3.1. Let \mathbf{G} be the following structure.

1. The vertex set G consists of all finite words W of length at least 1 in the alphabet Σ that satisfy the following condition.
 - (A1) For all i and j with $0 \leq i < j < |W|$, the structure induced by \mathbf{W}_i on $\{1\}$ is isomorphic to the structure induced by \mathbf{W}_j on $\{1\}$.
2. Let U, V be vertices of \mathbf{G} with $|U| < |V|$ that satisfy the following condition.
 - (A2) The structure $\Delta(\mathbf{U}_i, \mathbf{V}_{|U|}, \mathbf{V}_i)$ is defined for every i with $0 \leq i < |U|$ and it has an embedding to \mathbf{K} .

Then the mapping $0 \mapsto U, 1 \mapsto V$ is an embedding of type $\mathbf{V}_{|U|} \rightarrow \mathbf{G}$.

3. There are no tuples in the relations $R_{\mathbf{G}}, R \in L$, other than the ones given by 2.

Lemma 11.3.2. *Every induced cycle in \mathbf{G} has a completion to \mathbf{K} . Since every irreducible substructure of size at most 3 embeds into \mathbf{K} there is a completion $\mathbf{G} \rightarrow \mathbf{K}$.*

Proof. Suppose for contradiction that there exists ℓ and a sequence $U^0, U^1, \dots, U^{\ell-1}$ forming an induced cycle \mathbf{C} in \mathbf{G} such that \mathbf{C} has no completion to \mathbf{K} . Without loss of generality, we assume that $|U^0| \leq |U^k|$ for every $1 \leq k < \ell$. We create a structure \mathbf{D} from \mathbf{C} by adding precisely those tuples to the relations of \mathbf{D} such that the mapping $0 \mapsto \mathbf{U}^0, 1 \mapsto \mathbf{U}^k$ is an embedding from $\mathbf{U}_{|U^0|}^k$ to \mathbf{D} for every k satisfying $2 \leq k < \ell$ and $|U^0| < |U^k|$.

For simplicity, consider first the case that we have $|U^0| < |U^k|$ for every $1 \leq k \leq \ell - 1$. In this case, we produced a triangulation of \mathbf{D} : all induced cycles are triangles containing the vertex U^0 . It follows from the construction of \mathbf{G} that, for every $2 \leq k \leq \ell$, the triangle induced by \mathbf{D} on U^0, U^k and U^{k+1} is isomorphic either to $\Delta(\mathbf{U}_{|U^0|}^k, \mathbf{U}_{|U^k|}^{k+1}, \mathbf{U}_{|U^0|}^{k+1})$ (if $|U^k| < |U^{k+1}|$) or to $\Delta(\mathbf{U}_{|U^0|}^{k+1}, \mathbf{U}_{|U^{k+1}|}^k, \mathbf{U}_{|U^0|}^k)$. By (A2) the triangle has an embedding to \mathbf{K} , hence all induced cycles in \mathbf{D} have a completion to \mathbf{K} , which implies that \mathbf{D} has a completion $c: \mathbf{D} \rightarrow \mathbf{K}$. We get completion $c: \mathbf{C} \rightarrow \mathbf{K}$, a contradiction.

It remains to consider the case that there are multiple vertices of \mathbf{D} of length $|U^0|$. We then set $M := \{U^k : |U^k| = |U^0|\}$. By the construction of \mathbf{G} , the vertices in M are never neighbours. Moreover, for every $U, V \in M$, the structure induced on $\{U\}$ by \mathbf{C} is isomorphic to structure induced on $\{V\}$ by \mathbf{C} , which, by (A2), is isomorphic to the structure induced on $\{0\}$ by $\mathbf{W}_{|U^0|}$ for every $W \in C \setminus M$. Consequently, it is possible to construct a structure \mathbf{E} from \mathbf{D} by identifying all vertices in M and to obtain a homomorphism-embedding $f: \mathbf{D} \rightarrow \mathbf{E}$. Observe that the structure \mathbf{E} is triangulated and every triangle is known to have a completion to \mathbf{K} . By Lemma 11.3.1, \mathbf{D} also has a completion to \mathbf{K} . \square

The following result follows directly from the definition of substitution.

Observation 11.3.3. *For every $W \in [\emptyset] \binom{\omega}{\omega}$ and all $U, V \in G$, the structure induced by \mathbf{G} on $\{U, V\}$ is isomorphic to the structure induced by \mathbf{G} on $\{W(U), W(V)\}$.*

Without loss of generality we assume that $K = \omega \setminus \{0\}$. Let \mathbf{K}' be the structure \mathbf{K} extended by the vertex 0 such that there exists an embedding $\mathbf{K}' \rightarrow \mathbf{K}$. Such a structure \mathbf{K}' exists, because duplicating the vertex 1 does not introduce new induced cycles. We define the mapping $\varphi: \omega \setminus \{0\} \rightarrow G$ by setting $\varphi(i) = U$, where U is a word of length i defined by setting, for every $0 \leq j < i$, \mathbf{U}_j as the unique structure in Σ such that $0 \mapsto j, 1 \mapsto i$ is an embedding $\mathbf{U}_j \rightarrow \mathbf{K}'$. It is easy to check that φ is an embedding $\varphi: \mathbf{K} \rightarrow \mathbf{G}$. We prove Theorem 11.1.1 in the following form.

Theorem 11.3.4. *For every finite $k \geq 1$ and every finite colouring of subsets of G with k elements, there exists $f \in \binom{\mathbf{G}}{\mathbf{G}}$ such that the colour of every k -element subset S of $f(\mathbf{G})$ depends only on $\tau(S) = \tau(f^{-1}[S])$.*

By Proposition 11.2.2, we obtain the desired finite upper bound on the number of colours. By the completion $c: \mathbf{G} \rightarrow \mathbf{K}$ given by Lemma 11.3.2, the colouring of substructures of \mathbf{K} yields a colouring of irreducible substructures of \mathbf{G} . Embedding $f \in \binom{\mathbf{G}}{\mathbf{G}}$ can be restricted $f' \in \binom{\mathbf{G}}{\mathbf{K}}$ and gives $c \circ f' \in \binom{\mathbf{K}}{\mathbf{K}}$ and thus Theorem 11.3.4 indeed implies Theorem 11.1.1.

Sketch. Fix k and a finite colouring χ of the subsets of G of size k . Proposition 11.2.2 bounds number of embedding types of subsets of G of size k . Apply Theorem 11.2.1 for each embedding type. By Observation 11.3.3, we obtain the desired embedding; see [Hub20a, proof of Theorem 4.4] for details. \square

12. Type-respecting amalgamation and big Ramsey degrees

ANDRÉS ARANDA, SAUMUEL BRAUNFELD, DAVID CHODOUNSKÝ, JAN HUBIČKA, MATĚJ KONEČNÝ, JAROSLAV NEŠETŘIL, ANDY ZUCKER

Abstract We give an infinitary extension of the Nešetřil–Rödl theorem for category of relational structures with special type-respecting embeddings.

12.1 Introduction

We use the standard model-theoretic notion of structures allowing functions to be partial. Let L be a *language* with relational symbols $R \in L$ and functional symbols $f \in L$ each having its *arity*. An L -structure \mathbf{A} on A is a structure with *vertex set* A , relations $R_{\mathbf{A}} \subseteq A^r$ for every relation symbol $R \in L$ of arity r and partial functions $F_{\mathbf{A}}: A^s \rightarrow A$ for every function symbol $F \in L$ of arity s . If the set A is finite say that \mathbf{A} is *finite* (it may still have infinitely many relations if L is infinite). We consider only L -structures with finitely many or countably infinitely many vertices. Language L is *relational* if it contains no function symbols. We say that \mathbf{A} is a *substructure* of \mathbf{B} and write $\mathbf{A} \subseteq \mathbf{B}$ if the identity map is an embedding $\mathbf{A} \rightarrow \mathbf{B}$. Let \mathcal{K} be a class of L -structures. We say that \mathcal{K} is *hereditary* if it is closed for substructures. We say that L -structure $\mathbf{U} \in \mathcal{K}$ is \mathcal{K} -*universal* if every L -structure in \mathcal{K} embeds to \mathbf{U} .

Given L -structures \mathbf{A} and \mathbf{B} , we denote by $\binom{\mathbf{B}}{\mathbf{A}}$ the set of all embeddings from \mathbf{A} to \mathbf{B} . We write $\mathbf{C} \rightarrow (\mathbf{B})_{k,l}^{\mathbf{A}}$ to denote the following statement: for every colouring χ of $\binom{\mathbf{C}}{\mathbf{A}}$ with k colours, there exists an embedding $f: \mathbf{B} \rightarrow \mathbf{C}$ such that χ does not take more than l values on $\binom{f(\mathbf{B})}{\mathbf{A}}$. For a countably infinite L -structure \mathbf{B} and its finite substructure \mathbf{A} , the *big Ramsey degree* of \mathbf{A} in \mathbf{B} is the least number $D \in \mathbb{N} \cup \{\infty\}$ such that $\mathbf{B} \rightarrow (\mathbf{B})_{k,D}^{\mathbf{A}}$ for every $k \in \mathbb{N}$. We say that L -structure \mathbf{B} has *finite big Ramsey degrees* if the big Ramsey degree of every finite substructure \mathbf{A} of \mathbf{B} is finite. In general, we are interested in the following question: Given a hereditary class of L -structures \mathcal{K} , do \mathcal{K} -universal L -structures in \mathcal{K} have finite big Ramsey degrees? Notice that if one \mathcal{K} -universal L -structure in \mathcal{K} has finite big Ramsey degrees then all of them do. The study of big Ramsey degrees originated in 1960's Laver's unpublished proof that the big Ramsey degrees of the order of rationals are finite. This result was refined and a precise formula was obtained by Devlin [Dev79]. This area has recently been revitalized with a rapid progress regarding big Ramsey degrees of structures in finite binary languages (see e.g. recent survey [Dob21]).

We call an L -structure \mathbf{A} *irreducible* if for every pair of vertices $u, v \in A$ there exist a relational symbol $R \in L$ and a tuple $\vec{t} \in R_{\mathbf{A}}$ such that $u, v \in \vec{t}$. Given set of L -structures \mathcal{F} , L -structure \mathbf{A} is \mathcal{F} -*free* if there is no $\mathbf{F} \in \mathcal{F}$ with an embedding $\mathbf{F} \rightarrow \mathbf{A}$. The class of all (finite and countably infinite) \mathcal{F} -free L -

structures is denoted by $\text{Forb}_{\text{he}}(\mathcal{F})$. With these definitions we can state a recent result:

Theorem 12.1.1 (Zucker [Zuc22]). *Let L be a finite binary relational language, and \mathcal{F} a finite set of finite irreducible L -structures. Then every $\text{Forb}_{\text{he}}(\mathcal{F})$ -universal L -structure has finite big Ramsey degrees. (In other words, for every finite substructure \mathbf{A} of \mathbf{U} there exists finite $D = D(\mathbf{A})$ such that $\mathbf{U} \rightarrow (\mathbf{U})_{k,D}^{\mathbf{A}}$ for every $k > 1$.)*

This result can be seen as an infinitary variant of well known Nešetřil–Rödl theorem (one of the fundamental results of structural Ramsey theory) which can be stated as follows:

Theorem 12.1.2 (Nešetřil–Rödl theorem [NR77a, NR89]). *Let L be a relational language, \mathcal{F} a set of finite irreducible L -structures. Then for every finite $\mathbf{A} \in \text{Forb}_{\text{he}}(\mathcal{F})$ there exists a finite integer $d = d(\mathbf{A})$ such that for every finite $\mathbf{B} \in \text{Forb}_{\text{he}}(\mathcal{F})$ and finite $k > 0$ there exists a finite $\mathbf{C} \in \text{Forb}_{\text{he}}(\mathcal{F})$ satisfying $\mathbf{C} \rightarrow (\mathbf{B})_{k,d}^{\mathbf{A}}$.*

To see the correspondence of Theorems 12.1.1 and 12.1.2 choose \mathcal{F} as in Theorem 12.1.1 and a finite \mathcal{F} -free L -structure \mathbf{A} . By Theorem 12.1.1 there is a finite $D = D(\mathbf{A})$ such that every $\text{Forb}_{\text{he}}(\mathcal{F})$ -universal L -structure \mathbf{U} satisfies $\mathbf{U} \rightarrow (\mathbf{U})_{k,D}^{\mathbf{A}}$ for every $k > 0$. By $\text{Forb}_{\text{he}}(\mathcal{F})$ -universality of \mathbf{U} for every \mathcal{F} -free L -structure \mathbf{B} we have $\mathbf{U} \rightarrow (\mathbf{B})_{k,D}^{\mathbf{A}}$ and by compactness there exists a finite substructure \mathbf{C} of \mathbf{U} such that $\mathbf{C} \rightarrow (\mathbf{B})_{k,D}^{\mathbf{A}}$. In general, $D(\mathbf{A})$ (characterised precisely in [BCD⁺21b]) differs from $d(\mathbf{A})$, the number of linear orderings of A . However, the proof of Theorem 12.1.1 can be used to recover precise bounds on $d(\mathbf{A})$.

Comparing Theorems 12.1.1 and 12.1.2, it is natural to ask whether the assumptions about finiteness of \mathcal{F} , finiteness of the language L , and relations being only binary can be dropped from Theorem 12.1.1. It is known that the first condition can not be omitted: Sauer [Sau02] has shown that there exist infinite families \mathcal{F} of finite irreducible L -structures where $\text{Forb}_{\text{he}}(\mathcal{F})$ -universal structures have infinite big Ramsey degrees of vertices. This is true even for language L containing only one binary relation (digraphs). The latter two conditions remain open.

Until recently, most bounds on big Ramsey degrees were for L -structures in binary languages only. Techniques to give bounds on big Ramsey degrees of 3-uniform hypergraphs have been announced in Eurocomb 2019 [BCH⁺19] and published recently [BCH⁺22]; they were later extended to languages with arbitrary relational symbols [BCdR⁺23]. Extending links between the Hales–Jewett theorem [HJ63], Carlson–Simpson theorem [CS84] and big Ramsey degrees established in [Hub20a], a Ramsey-type theorem for trees with successor operations has been introduced [BCD⁺23b] which extends to all known big Ramsey results on L -structures. However, the following problem remains open:

Problem 12.1.3. *Let $L = \{E, H\}$ be a language with one binary relation E and one ternary relation H . Let \mathbf{F} be the L -structure where $F = \{0, 1, 2, 3\}$, $R^F = \{(1, 0), (1, 2), (1, 3)\}$, $H = \{(0, 2, 3)\}$. Denote by \mathcal{K} the class of all L -structures \mathbf{A} such that there is no monomorphism $\mathbf{F} \rightarrow \mathbf{A}$. Do \mathcal{K} -universal L -structures have finite big Ramsey degrees?*

This problem demonstrates unforeseen obstacles on giving a natural infinitary generalization of the Nešetřil–Rödl theorem. We give a new approach which avoids this issue and which suggests perhaps the proper setting for big Ramsey degrees.

Finite structural Ramsey results are most often proved by refinements of the Nešetřil–Rödl partite construction [NR89]. This technique does not generalize to infinite structures due to essential use of backward induction. Upper bounds on big Ramsey degrees are based on Ramsey-type theorems on trees (e.g. the Halper–Läuchli theorem [HL66], Milliken’s tree theorem [Mil79], the Carlson–Simpson theorem [CS84], and their various refinements [Zuc22, Dob20a, Dob23]). This proof structure may seem unexpected at first glance but is justified by the existence of unavoidable colourings (based on idea of Sierpiński) which are constructed by assigning colors according to subtrees of the tree of 1-types (see e.g. [BCD⁺23a, Dob21]). The exact characterisations of big Ramsey degrees can then be understood as an argument that this proof structure is in a very specific sense the only possible: the trees used to give upper bounds are also encoded in the precise characterisations of big Ramsey degrees.

We briefly review the construction of tree of 1-types. Recall that a (model-theoretic) *tree* is a partial order (T, \leq) where the down-set of every $x \in T$ is a finite chain. An *enumerated L -structure* is simply an L -structure \mathbf{U} whose underlying set is the ordinal $|\mathbf{U}|$. Fix a countably infinite enumerated L -structure \mathbf{U} . Given vertices u, v and an integer n satisfying $\min(u, v) \geq n \geq 0$, we write $u \sim_n^{\mathbf{U}} v$, and say that u and v are of the same (quantifier-free) type over $0, 1, \dots, n-1$, if the L -structure induced by \mathbf{U} on $\{0, 1, \dots, n-1, u\}$ is identical to the L -structure induced by \mathbf{U} on $\{0, 1, \dots, n-1, v\}$ after renaming vertex v to u . We write $[u]_n^{\mathbf{U}}$ for the $\sim_n^{\mathbf{U}}$ -equivalence class of vertex u .

Definition 12.1.1 (Tree of 1-types). Let \mathbf{U} be an infinite (relational) enumerated L -structure. Given $n < \omega$, write $\mathbb{T}_{\mathbf{U}}(n) = \omega / \sim_n^{\mathbf{U}}$. A (quantifier-free) *1-type* is any member of the disjoint union $\mathbb{T}_{\mathbf{U}} := \bigsqcup_{n < \omega} \mathbb{T}_{\mathbf{U}}(n)$. We turn $\mathbb{T}_{\mathbf{U}}$ into a tree as follows. Given $x \in \mathbb{T}_{\mathbf{U}}(m)$ and $y \in \mathbb{T}_{\mathbf{U}}(n)$, we declare that $x \leq_{\mathbb{T}_{\mathbf{U}}}^{\mathbf{U}} y$ if and only if $m \leq n$ and $x \supseteq y$.

One can associate every vertex of $v \in \mathbf{U}$ with its corresponding equivalence class in $\simeq_v^{\mathbf{U}}$. This way every substructure $\mathbf{A} \subseteq \mathbf{U}$ corresponds to a subset of nodes of the tree $\mathbb{T}_{\mathbf{U}}$. Sierpiński-like colourings can be then constructed by considering shapes of the meet closures of nodes corresponding to each given copy. Every type in $x \in \omega / \sim_n^{\mathbf{U}}$ can be described as an L -structure \mathbf{T} with vertex set $T = \{0, 1, \dots, n-1, t\}$ such that for every $v \in x$ it holds that L -structure induced by \mathbf{U} on $\{0, 1, \dots, n-1, v\}$ is \mathbf{T} after renaming t to v . This is very useful in the setting where types originating from multiple enumerated L -structures are considered (see for instance [Hub20a, BCD⁺23a]).

The concept of the tree of 1-types was implicit in early proofs (such as in Devlin’s thesis) and became explicit later. The tree of 1-types itself is, however, not sufficient to give upper bounds on big Ramsey degrees for L -structures in languages containing symbols of arity 3 and more. Upper bounds in [BCH⁺22] and [BCdR⁺23] are based on the product form of the Milliken tree theorem which in turn suggests the following notion of a weak type.

For the rest of this note, fix a relational language L containing a binary symbol \leq . For all L -structures, \leq will always be a linear order on vertices which is either

finite or of order-type ω . This will describe the enumeration. All embeddings will be monotone.

Definition 12.1.2 (Weak type). We denote by L^f the language L extended by unary function symbol f . An L^f -structure \mathbf{T} is a *weak type* of level ℓ if

1. $T = \{0, 1, \dots, \ell - 1, t_0, t_1, \dots\}$ where vertices t_i are called *type vertices*.
2. For every $R \in L$ and $\vec{t} \in R_{\mathbf{T}}$ it holds that $\vec{t} \cap \{t_0, t_1, \dots\}$ is a (possibly empty) initial segment of type vertices (i.e. set of the form $\{t_i : i \in k\}$ for some $k \in \omega$) and $\vec{t} \cap \{0, 1, \dots, \ell - 1\} \neq \emptyset$.
3. For every $i > 0$ we put $F_{\mathbf{T}}(t_i) = t_{i-1}$, $F_{\mathbf{T}}(t_0) = t_0$, and $F_{\mathbf{T}}$ is undefined otherwise.

Weak types thus give less information than standard model-theoretic k -types (see e.g. [Hod93] for definitions). Function f is added to type vertices to distinguish them from normal vertices. This will be useful in later constructions. Notice that while technically weak type has infinitely many types vertices, thanks to condition 2 of Definition 12.1.2, if the language L contains no relations of arity $r + 1$ or more, vertices t_{r-1}, t_r, \dots will be isolated. In particular:

Observation 12.1.4. *If L contains only unary and binary symbols then there is one-to-one correspondence between 1-types and weak types because only type vertex t_0 carries interesting structure.*

1-types describes one vertex extensions of an initial part of the enumerated L -structure. The weak-type equivalent of this is the following:

Definition 12.1.3 (Weak type of a tuple). Let \mathbf{A} be an enumerated L -structure, \mathbf{T} a weak type of level $\ell \in A \subseteq \omega$ and $\vec{a} = (a_0, a_1, \dots, a_{k-1})$ an increasing tuple of vertices from $A \setminus \ell$. We say that \vec{a} *has type* \mathbf{T} on level ℓ if the function $h: T \rightarrow A$ given by:

$$h(x) = \begin{cases} x & \text{if } x \in \ell, \\ a_i & \text{if } x = t_i \text{ for some } i < k \end{cases}$$

has the property that for every $R \in L$ and \vec{b} a tuple of vertices in $\{0, 1, \dots, \ell - 1, t_0, t_1, \dots, t_{k-1}\}$ such that $\vec{b} \cap \{t_0, t_1, \dots\}$ is an initial segment of type vertices and $\vec{b} \cap \{0, 1, \dots, \ell - 1\} \neq \emptyset$ it holds that $\vec{b} \in R_{\mathbf{T}} \iff h(\vec{b}) \in R_{\mathbf{A}}$.

Definition 12.1.4 (Tree of weak types). Given an enumerated L -structure \mathbf{U} , its *tree of weak types* consists of all L^f -structures \mathbf{T} that are weak types of some tuple of \mathbf{U} on some level $\ell \in U$ ordered by \subseteq .

Given an enumerated L -structure \mathbf{A} and a weak type \mathbf{T} , we say that \mathbf{T} *extends* \mathbf{A} if $\mathbf{T} \setminus \{t_0, t_1, \dots\} = \mathbf{A}$. Given two types \mathbf{T} and \mathbf{T}' that extend \mathbf{A} , and $n \geq 0$, we say that \mathbf{T} and \mathbf{T}' agree as n -types if $\mathbf{T} \upharpoonright (A \cup \{t_0, t_1, \dots, t_{n-1}\}) = \mathbf{T}' \upharpoonright (A \cup \{t_0, t_1, \dots, t_{n-1}\})$.

A standard technique for proving infinite Ramsey-type theorems is to work with finite approximations of the embeddings considered. See e.g. Todorćević's axiomatization of Ramsey spaces [Tod10]. Initial approximations of our embeddings will be described as follows:

Definition 12.1.5 (Structure with types). Given a finite enumerated L -structure \mathbf{A} , \mathbf{A}^+ denotes the L -structure created from the disjoint union of all weak types extending \mathbf{A} by

1. identifying all copies of \mathbf{A} , and,
2. identifying the copy of vertex t_i of weak type \mathbf{T} and with the copy of t_i of weak type \mathbf{T}' whenever \mathbf{T} and \mathbf{T}' agree as $i + 1$ types.

Observe that thanks to the function f added to weak types, for any two L -structures with types \mathbf{A}^+ and \mathbf{B}^+ , every embedding $h: \mathbf{A}^+ \rightarrow \mathbf{B}^+$ is also a map from weak types of \mathbf{A} on level $|A|$ to weak types of \mathbf{B} of level $|B|$.

Given an L -structure \mathbf{A} and a vertex v , we denote by $\mathbf{A}(<v)$ the L -structure induced by \mathbf{A} on $\{a \in A; a < v\}$ and call it the *initial segment* of \mathbf{A} . The key notion for our approach is to restrict attention to embedding which behave well with respect to weak types. That is, for every initial segment of the L -structure, the rest of the embedding can be summarized via embedding of weak types extending the initial segments.

Definition 12.1.6 (Type-respecting embeddings of L -structures). Given enumerated L -structures \mathbf{A} and \mathbf{B} and an embedding $h: \mathbf{A} \rightarrow \mathbf{B}$, we say that h is *type-respecting* if for every $v \in A$ there exists an embedding $h^v: \mathbf{A}(<v)^+ \rightarrow \mathbf{B}(<h(v))^+$ such that the weak types of tuples in \mathbf{B} on level $h(v)$ consisting only of vertices of $h[A]$ are all in the image $h^v[\mathbf{A}]$.

Definition 12.1.7 (\mathcal{K} -type-respecting embeddings of initial segments). Let \mathbf{A} and \mathbf{B} be two finite enumerated L -structures. Embedding $h: \mathbf{A}^+ \rightarrow \mathbf{B}^+$ is *type-respecting* if for every (possibly infinite) L -structure \mathbf{A}' with initial segment \mathbf{A} there exists an L -structure \mathbf{B}' with initial segment \mathbf{B} and a type-respecting embedding $g: \mathbf{A} \rightarrow \mathbf{B}$ finitely approximated by h . That is $g \upharpoonright A = h \upharpoonright A$ and every weak type in \mathbf{B}' of a tuple consisting of vertices of $g[A]$ of level $g(\max A)$ is in $h[A^+]$.

Given class \mathcal{K} of L -structures we say that $h: \mathbf{A}^+ \rightarrow \mathbf{B}^+$ is *\mathcal{K} -type-respecting* if for every L -structure $\mathbf{A}' \in \mathcal{K}$ with initial segment \mathbf{A} there exists an structure $\mathbf{B}' \in \mathcal{K}$ with initial segment \mathbf{B} and a type-respecting embedding $g: \mathbf{A} \rightarrow \mathbf{B}$ finitely approximated by h .

Definition 12.1.8 (Type-respecting amalgamation property). Let \mathcal{K} be a hereditary class of enumerated L -structures. We say that \mathcal{K} has *type-respecting amalgamation property* if given three finite enumerated L -structures $\mathbf{A}, \mathbf{B}, \mathbf{B}' \in \mathcal{K}$ such that $B' \setminus B = \{\max B'\}$ and $\mathbf{B}' \upharpoonright \mathbf{B} = \mathbf{B}$, two \mathcal{K} -type-respecting embeddings $f: \mathbf{A}^+ \rightarrow \mathbf{B}^+$, $f': \mathbf{A}^+ \rightarrow \mathbf{B}'^+$ and a type-respecting (but not necessarily \mathcal{K} -type-respecting) embedding $g: \mathbf{B}^+ \rightarrow \mathbf{B}'^+$ such that $g \upharpoonright B$ is the identity and $g \circ f = f'$, there exists a \mathcal{K} -type-respecting embedding $g': \mathbf{B}^+ \rightarrow \mathbf{B}'^+$ such that $g' \circ f = f'$ and $g' \upharpoonright B = \text{Id}$.

Given a class of L -structures \mathcal{K} , finite $\mathbf{A} \in \mathcal{K}$ and $\mathbf{B} \in \mathcal{K}$, we denote by $\left(\frac{\mathbf{B}}{\mathbf{A}}\right)^\mathcal{K}$ the set of all \mathcal{K} -type-respecting embeddings $\mathbf{A}^+ \rightarrow \mathbf{B}'^+$ for \mathbf{B}' an initial segment of \mathbf{B} . We write $\mathbf{C} \xrightarrow{\mathcal{K}} \left(\frac{\mathbf{B}}{\mathbf{A}}\right)^\mathcal{K}$ to denote the following statement: for every colouring χ of $\left(\frac{\mathbf{C}}{\mathbf{A}}\right)^\mathcal{K}$ with k colours, there exists a type-respecting embedding $f: \mathbf{B} \rightarrow \mathbf{C}$

such that χ does not take more than l values on $\left(\frac{f(\mathbf{B})}{\mathbf{A}}\right)^{\mathcal{K}}$. For a countably infinite L -structure \mathbf{B} and its finite suborder \mathbf{A} , the *big Ramsey degree of \mathbf{A} in \mathcal{K} -type-respecting embeddings* of \mathbf{A} in \mathbf{B} is the least number $D \in \mathbb{N} \cup \{\infty\}$ such that $\mathbf{B} \longrightarrow^{\mathcal{K}} (\mathbf{B})_{k,D}^{\mathbf{A}}$ for every $k \in \mathbb{N}$.

For type-respecting embeddings we can prove the Ramsey property in the full generality (showing that, in this situation, Problem 12.1.3 is not a problem).

Theorem 12.1.5. *Let L be a finite relational language. Let \mathcal{F} be a finite family of finite irreducible enumerated L -structures. Denote by $\mathcal{K}_{\mathcal{F}}$ the class of all finite or countably-infinite enumerated L -structures \mathbf{A} where $\leq_{\mathbf{A}}$ is either finite or of order-type ω such that for every $\mathbf{F} \in \mathcal{F}$ there no embedding $\mathbf{F} \rightarrow \mathbf{A}$. Assume that $\mathcal{K}_{\mathcal{F}}$ has the type-respecting amalgamation property. Then for every universal L -structure $\mathbf{U} \in \mathcal{K}_{\mathcal{F}}$ and every finite $\mathbf{A} \in \mathcal{K}_{\mathcal{F}}$ there is a finite $D = D(\mathbf{A})$ such that $\mathbf{U} \longrightarrow^{\mathcal{K}} (\mathbf{U})_{k,D}^{\mathbf{A}}$ for every $k \in \mathbb{N}$.*

We show the following:

Proposition 12.1.6. *Let L be a finite language consisting of binary and unary relational symbols only. Let \mathcal{F} be a finite family of enumerated irreducible L -structures. Then $\mathcal{K}_{\mathcal{F}}$ has the type-respecting amalgamation property. Moreover, Theorem 12.1.5 implies Theorem 12.1.1.*

Proof. Fix L , \mathcal{F} and $\mathcal{K}_{\mathcal{F}}$. Let $\mathbf{A}, \mathbf{B}, \mathbf{B}' \in \mathcal{K}_{\mathcal{F}}$, $f: \mathbf{A}^+ \rightarrow \mathbf{B}^+$, $f': \mathbf{A}^+ \rightarrow \mathbf{B}'^+$ and $g: \mathbf{B}^+ \rightarrow \mathbf{B}'^+$ be as in Definition 12.1.8. By Observation 12.1.4, in order to specify g' , it is only necessary to give, for every weak type \mathbf{T} extending \mathbf{A} , an image of its type vertex $t_0 \in T$. Let $t' \in B^+$ be a vertex corresponding to t_0 . We consider two cases. (1) If $t' \in f[A^+]$ then we put $g'(t') = g(t')$. (2) If $t' \notin f[A^+]$ we put $g'(t') = t''$ where t'' is the only possible image of t' such that there is no relational symbol $R \in L$ such that $R_{\mathbf{B}^+}$ contains a tuple with both t' and $\max B'$.

To verify that g' is $\mathcal{K}_{\mathcal{F}}$ -type-respecting, choose $\mathbf{A}' \in \mathcal{K}_{\mathcal{F}}$ with initial segment \mathbf{B} . Construct \mathbf{A}'' from \mathbf{A} by inserting a new vertex v after $\max B$ and extending $\leq_{\mathbf{A}''}$. Add the needed tuples to relations to make \mathbf{B}' the initial segment of \mathbf{A}'' . Finally, for every $R \in L$ and $u \in A'$ with $u > v$, put $(u, v) \in R_{\mathbf{A}''}$ if and only if $(g'(t), u) \in R_{\mathbf{B}^+}$ where t is the type vertex of \mathbf{B} corresponding to the type of u in \mathbf{A}' . Add tuples $(v, u) \in R_{\mathbf{A}''}$ analogously.

To verify that $\mathbf{A}'' \in \mathcal{K}_{\mathcal{F}}$, assume to the contrary that there is $\mathbf{F} \in \mathcal{F}$ and embedding $e: \mathbf{F} \rightarrow \mathbf{A}''$. Because $\mathbf{A}' \in \mathcal{K}_{\mathcal{F}}$, clearly $v \in e[F]$. Because $\mathbf{B}' \in \mathcal{K}_{\mathcal{F}}$ we also know that $e[F]$ contains vertices of $\mathbf{A}'' \setminus \mathbf{B}'$. Since \mathbf{F} is irreducible, all such vertices must have types created by condition (1) above. This contradicts that f' is $\mathcal{K}_{\mathcal{F}}$ -type-respecting.

To see the moreover part we have to construct a universal \mathbf{U} which is a substructure of some $\mathbf{U}' \in \mathcal{K}_{\mathcal{F}}$ with the property that for every $n \in \mathbb{N}$ there exists $N \in \mathbb{N}$ such that for every $\mathbf{A} \in \mathcal{K}_{\mathcal{F}}$ with n vertices and every embedding $e: \mathbf{A} \rightarrow \mathbf{U}'$ there exist a structure $\mathbf{E} \in \mathcal{K}_{\mathcal{F}}$ (called an *envelope*) with at most N vertices and a $\mathcal{K}_{\mathcal{F}}$ -type-respecting embedding $h: \mathbf{E} \rightarrow \mathbf{U}$ such that $e[A] \subseteq h[E]$. This follows from Section 4 of [Zuc22], because \mathcal{K} -type-respecting embeddings in this setup are precisely the aged embeddings from [Zuc22]. \square

Proposition 12.1.7. *Let $L' = \{E, H, \leq\}$ and let \mathbf{F}' be an L' -structure created by expanding the L -structure \mathbf{F} from Problem 12.1.3 by the natural order of vertices. Denote by \mathcal{K} the class of all enumerated L' -structures \mathbf{A} for which there*

is no monomorphism $\mathbf{F} \rightarrow \mathbf{A}$. The class \mathcal{K} has no type-respecting amalgamation property.

Proof. We give an explicit failure of type-respecting amalgamation showing that the use of Observation 12.1.4 in the previous proof is essential. Let \mathbf{A} be the empty L' -structure, \mathbf{B} be L' -structure with $B = \{0\}$, $E_{\mathbf{B}} = H_{\mathbf{B}} = \emptyset$ and \mathbf{B}' be L' -structure with $B' = \{0, 1\}$, $E_{\mathbf{B}'} = \{(0, 1)\}$, $H_{\mathbf{B}'} = \emptyset$. Let $\mathbf{T}_{\mathbf{A}}$ be the unique weak type extending \mathbf{A} . Let $\mathbf{T}_{\mathbf{B}}$ be weak type extending \mathbf{B} with $E_{\mathbf{T}_{\mathbf{B}}} = H_{\mathbf{T}_{\mathbf{B}}} = \emptyset$ and $\mathbf{T}'_{\mathbf{B}}$ weak type extending \mathbf{B} with $E_{\mathbf{T}'_{\mathbf{B}}} = \emptyset$ and $H_{\mathbf{T}'_{\mathbf{B}}} = \{(0, t_0, t_1)\}$. Notice that $\mathbf{T}_{\mathbf{B}}$ and $\mathbf{T}'_{\mathbf{B}}$ agree as 1-types and thus in \mathbf{B}^+ their vertices t_0 are identified. Finally, let $\mathbf{T}_{\mathbf{B}'}$ and $\mathbf{T}'_{\mathbf{B}'}$ be weak types extending \mathbf{B}' with $E_{\mathbf{T}_{\mathbf{B}'}} = H_{\mathbf{T}'_{\mathbf{B}'}} = \{(0, 1), (1, t_0)\}$, $H_{\mathbf{T}_{\mathbf{B}'}} = \emptyset$, $H_{\mathbf{T}'_{\mathbf{B}'}} = \{(0, t_0, t_1)\}$. Again $\mathbf{T}_{\mathbf{B}}$ and $\mathbf{T}_{\mathbf{B}'}$ agree as 1-types. Now let $f: \mathbf{A}^+ \rightarrow \mathbf{B}^+$ map $\mathbf{T}_{\mathbf{A}}$ to $\mathbf{T}_{\mathbf{B}}$ and $f^+: \mathbf{A}^+ \rightarrow \mathbf{B}^+$ map $\mathbf{T}_{\mathbf{A}}$ to $\mathbf{T}_{\mathbf{B}'}$. It is easy to check that these are \mathcal{K} -type-respecting. $g: \mathbf{B}^+ \rightarrow \mathbf{B}'^+$ can be constructed to be type-respecting by mapping type $\mathbf{T}_{\mathbf{B}}$ to $\mathbf{T}_{\mathbf{B}'}$ and $\mathbf{T}'_{\mathbf{B}}$ to $\mathbf{T}'_{\mathbf{B}'}$. However there is no \mathcal{K} -type-respecting $g': \mathbf{B}^+ \rightarrow \mathbf{B}'^+$. To see that, observe that any image of $\mathbf{T}'_{\mathbf{B}}$ must agree as 1-type with $\mathbf{T}_{\mathbf{B}'}$ and consider \mathbf{A}' with $A' = \{0, 1, 2\}$ and $H_{\mathbf{A}'} = \{(0, 1, 2)\}$. \mathbf{A} is an initial segment of \mathbf{A}' and there is no way to extend g' to a \mathcal{K} -type-respecting embedding of \mathbf{A}' to some L -structure in \mathcal{K} since it will always add vertex v after vertex 0 of \mathbf{A} in a way that there is a monomorphism from \mathbf{F} to $\{0, v, 1, 2\}$. \square

We conjecture that the answer to Problem 12.1.3 is in fact negative. It is possible that by concentrating on type-respecting embeddings, the study of big Ramsey degrees can find a proper setting.

Acknowledgement First six authors are supported by the project 21-10775S of the Czech Science Foundation (GAČR). This article is part of a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 810115).

Bibliography

- [ABC⁺23] Andres Aranda, Samuel Braunfeld, David Chodounský, Jan Hubička, Matěj Konečný, Jaroslav Nešetřil, and Andy Zucker. Type-respecting amalgamation and big Ramsey degrees. Extended abstract accepted to EuroComb 2023, arXiv:2303.12679, 2023.
- [ABWH⁺17a] Andres Aranda, David Bradley-Williams, Eng Keat Hng, Jan Hubička, Miltiadis Karamanlis, Michael Kompatscher, Matěj Konečný, and Micheal Pawliuk. Completing graphs to metric spaces. *Electronic Notes in Discrete Mathematics*, 61:53–60, 2017. The European Conference on Combinatorics, Graph Theory and Applications (EUROCOMB’17).
- [ABWH⁺17b] Andres Aranda, David Bradley-Williams, Jan Hubička, Miltiadis Karamanlis, Michael Kompatscher, Matěj Konečný, and Micheal Pawliuk. Ramsey expansions of metrically homogeneous graphs. Accepted to *European Journal of Combinatorics*, arXiv:1707.02612, 2017.
- [ABWH⁺21] Andres Aranda, David Bradley-Williams, Eng Keat Hng, Jan Hubička, Miltiadis Karamanlis, Michael Kompatscher, Matěj Konečný, and Micheal Pawliuk. Completing graphs to metric spaces. *Contributions to Discrete Mathematics*, 16:71–89, 2021.
- [ACM21] Daniela A. Amato, Gregory Cherlin, and Dugald Macpherson. Metrically homogeneous graphs of diameter 3. *Journal of Mathematical Logic*, 21(01):2050020, 2021.
- [AH78] Fred G. Abramson and Leo A. Harrington. Models without indiscernibles. *Journal of Symbolic Logic*, 43:572–600, 1978.
- [AKL14] Omer Angel, Alexander S. Kechris, and Russell Lyons. Random orderings and unique ergodicity of automorphism groups. *Journal of the European Math. Society*, 16:2059–2095, 2014.
- [AN19] H. Andréka and I. Németi. Extending partial isomorphisms, a small construction. private communication, 2019.
- [Ass86] Roland Assous. Enchainabilité et seuil de monomorphie des tournois n -aires. *Discrete Mathematics*, 62(2):119–125, 1986. In French.
- [Bar20] Keegan Dasilva Barbosa. A categorical notion of precompact expansion. *arXiv preprint arXiv:2002.11751*, 2020.
- [Bau16] Andreas Baudisch. Free amalgamation and automorphism groups. *The Journal of Symbolic Logic*, 81(3):936–947, 2016.
- [BCD⁺21a] Martin Balko, David Chodounský, Natasha Dobrinen, Jan Hubička, Matěj Konečný, Lluís Vena, and Andy Zucker. Big

Ramsey degrees of the generic partial order. In Jaroslav Nešetřil, Guillem Perarnau, Juanjo Rué, and Oriol Serra, editors, *Extended Abstracts EuroComb 2021*, pages 637–643, Cham, 2021. Springer International Publishing.

- [BCD⁺21b] Martin Balko, David Chodounský, Natasha Dobrinen, Jan Hubička, Matěj Konečný, Lluís Vena, and Andy Zucker. Exact big Ramsey degrees via coding trees. Submitted, arXiv:2110.08409, 2021.
- [BCD⁺23a] Martin Balko, David Chodounský, Natasha Dobrinen, Jan Hubička, Matěj Konečný, Lluís Vena, and Andy Zucker. Characterisation of the big Ramsey degrees of the generic partial order. Submitted, arXiv:2303.10088, 2023.
- [BCD⁺23b] Martin Balko, David Chodounský, Natasha Dobrinen, Jan Hubička, Jaroslav Nešetřil, Matěj Konečný, and Andy Zucker. Ramsey theorem for trees with successor operation. In preparation, 2023.
- [BCdR⁺23] Samuel Braunfeld, David Chodounský, Noé de Rancourt, Jan Hubička, Jamal Kawach, and Matěj Konečný. Big Ramsey degrees and infinite languages. Submitted, arXiv:2301.13116, 2023.
- [BCH⁺19] Martin Balko, David Chodounský, Jan Hubička, Matěj Konečný, and Lluís Vena. Big Ramsey degrees of 3-uniform hypergraphs. *Acta Mathematica Universitatis Comenianae*, 88(3):415–422, 2019. Extended abstract for Eurocomb 2019.
- [BCH⁺20] Martin Balko, David Chodounský, Jan Hubička, Matěj Konečný, and Lluís Vena. Big Ramsey degrees of unconstrained relational structures. (in preparation), 2020.
- [BCH⁺21a] Martin Balko, David Chodounský, Jan Hubička, Matěj Konečný, Jaroslav Nešetřil, and Lluís Vena. Big Ramsey degrees and forbidden cycles. In Jaroslav Nešetřil, Guillem Perarnau, Juanjo Rué, and Oriol Serra, editors, *Extended Abstracts EuroComb 2021*, pages 436–441. Springer International Publishing, 2021.
- [BCH⁺21b] Martin Balko, David Chodounský, Jan Hubička, Matěj Konečný, Jaroslav Nešetřil, and Lluís Vena. Infinitary Graham–Rothschild theorem for words of higher order. (in preparation), 2021.
- [BCH⁺22] Martin Balko, David Chodounský, Jan Hubička, Matěj Konečný, and Lluís Vena. Big Ramsey degrees of 3-uniform hypergraphs are finite. *Combinatorica*, 42(2):659–672, 2022.
- [BdRHK23] Tristan Bice, Noe de Racourt, Jan Hubička, and Matěj Konečný. Big Ramsey degrees in the metric setting. Extended abstract accepted to EuroComb 2023, arXiv:2303.14577, 2023.

- [Ber21] Jeffrey Bergfalk. The first omega alephs: From simplices to trees of trees to higher walks. *Advances in Mathematics*, 393:108083, 2021.
- [BK17] Dana Bartošová and Aleksandra Kwiatkowska. Gowers’ Ramsey theorem with multiple operations and dynamics of the homeomorphism group of the lelek fan. *Journal of Combinatorial Theory, Series A*, 150:108–136, 2017.
- [BK19] Dana Bartošová and Aleksandra Kwiatkowska. The universal minimal flow of the homeomorphism group of the lelek fan. *Transactions of the American Mathematical Society*, 371(10):6995–7027, 2019.
- [BLALM17] Dana Bartošová, Jordi Lopez-Abad, Martino Lupini, and Brice Mbombo. The Ramsey property for Banach spaces and Choquet simplices, and applications. *Comptes Rendus Mathématique*, 355(12):1242–1246, 2017.
- [BM05] Meenaxi Bhattacharjee and Dugald Macpherson. A locally finite dense group acting on the random graph. 17(3):513–517, 2005.
- [Bod15] Manuel Bodirsky. Ramsey classes: Examples and constructions. *Surveys in Combinatorics 2015*, 424:1, 2015.
- [BPT11] Manuel Bodirsky, Michael Pinsker, and Todor Tsankov. Decidability of definability. In *Proceedings of the 2011 IEEE 26th Annual Symposium on Logic in Computer Science, LICS ’11*, pages 321–328, Washington, DC, USA, 2011. IEEE Computer Society.
- [Bra16] Samuel Braunfeld. The lattice of definable equivalence relations in homogeneous n -dimensional permutation structures. *The Electronic Journal of Combinatorics*, 23(4):P4–44, 2016.
- [Bra17] Samuel Braunfeld. Ramsey expansions of Λ -ultrametric spaces. Accepted to *European Journal of Combinatorics*, arXiv:1710.01193, 2017.
- [BS71] Norman L. Biggs and Donald H. Smith. On trivalent graphs. *Bulletin of the London Mathematical Society*, 3(2):155–158, 1971.
- [BWC20] David Bradley-Williams and Peter J. Cameron. Notes on eppa witnesses. Unpublished note, 2020.
- [BYU10] Itai Ben Yaacov and Alexander Usvyatsov. Continuous first order logic and local stability. *Transactions of the American Mathematical Society*, 362(10):5213–5259, 2010.
- [Cam99] Peter J. Cameron. *Oligomorphic permutation groups*, volume 45 of *London Mathematical Society Student Texts*. Cambridge University Press, 1999.

- [Cam15] Peter J. Cameron. History of the Random Graph. <https://cameroncounts.wordpress.com/2015/08/04/history-of-the-random-graph/>, 2015.
- [CDP20] Rebecca Coulson, Natasha Dobrinen, and Rehana Patel. Canonical partitions of relational structures. In final stages of preparation, 2020.
- [CDP22a] Rebecca Coulson, Natasha Dobrinen, and Rehana Patel. Fraïssé structures with SDAP+, Part I: Indivisibility. *arXiv:2207.06393*, 2022.
- [CDP22b] Rebecca Coulson, Natasha Dobrinen, and Rehana Patel. Fraïssé structures with SDAP+, Part II: Simply characterized big Ramsey structures. *arXiv:2207.06505*, 2022.
- [CH03] Gregory Cherlin and Ehud Hrushovski. *Finite Structures with Few Types*. Princeton University Press, 2003.
- [Che] Gregory Cherlin. Some classification problems for finite homogeneity. Unpublished notes.
- [Che98] Gregory Cherlin. *The Classification of Countable Homogeneous Directed Graphs and Countable Homogeneous N -tournaments*. Number 621 in Memoirs of the American Mathematical Society. American Mathematical Society, 1998.
- [Che11] Gregory Cherlin. Two problems on homogeneous structures, revisited. *Model theoretic methods in finite combinatorics*, 558:319–415, 2011.
- [Che22] Gregory Cherlin. *Homogeneous Ordered Graphs, Metrically Homogeneous Graphs, and Beyond*, volume 1 of *Lecture Notes in Logic*. Cambridge University Press, 2022.
- [CHKN21] Gregory Cherlin, Jan Hubička, Matěj Konečný, and Jaroslav Nešetřil. Ramsey expansions of 3-hypertournaments. In *Extended Abstracts EuroComb 2021*, pages 696–701. Springer, 2021.
- [CKT20] F. Calderoni, A. Kwiatkowska, and K. Tent. Simplicity of the automorphism groups of order and tournament expansions of homogeneous structures. Preprint, arXiv:1808.05177, 2020.
- [Con19] Gabriel Conant. Extending partial isometries of generalized metric spaces. *Fundamenta Mathematicae*, 244:1–16, 2019.
- [CS84] Timothy J. Carlson and Stephen G. Simpson. A dual form of Ramsey’s theorem. *Advances in Mathematics*, 53(3):265–290, 1984.
- [Dev79] Denis Devlin. *Some partition theorems and ultrafilters on ω* . PhD thesis, Dartmouth College, 1979.

- [DK16] Pandelis Dodos and Vassilis Kanellopoulos. *Ramsey theory for product spaces*, volume 212. American Mathematical Society, 2016.
- [DLPS07] Christian Delhommé, Claude Laflamme, Maurice Pouzet, and Norbert W. Sauer. Divisibility of countable metric spaces. *European Journal of Combinatorics*, 28(6):1746–1769, 2007.
- [DLPS08] Christian Delhommé, Claude Laflamme, Maurice Pouzet, and Norbert W. Sauer. Indivisible ultrametric spaces. *Topology and its Applications*, 155(14):1462–1478, 2008.
- [DLS16] Natasha Dobrinen, Claude Laflamme, and Norbert Sauer. Rainbow Ramsey simple structures. *Discrete Mathematics*, 339(11):2848–2855, 2016.
- [Dob20a] Natasha Dobrinen. The Ramsey theory of the universal homogeneous triangle-free graph. *Journal of Mathematical Logic*, 20(02):2050012, 2020.
- [Dob20b] Natasha Dobrinen. The Ramsey theory of the universal homogeneous triangle-free graph part II: Exact big Ramsey degrees. *arXiv preprint arXiv:2009.01985*, 2020.
- [Dob20c] Natasha Dobrinen. Ramsey theory on infinite structures and the method of strong coding trees. In Adrian Rezus, editor, *Contemporary Logic and Computing, Landscapes in Logic*. College Publications, 2020.
- [Dob21] Natasha Dobrinen. Ramsey theory of homogeneous structures: current trends and open problems. *arXiv preprint arXiv:2110.00655; Survey for ICM’22*, 2021.
- [Dob23] Natasha Dobrinen. The Ramsey theory of Henson graphs. *Journal of Mathematical Logic*, 0(0):2250018, 2023. DOI: 10.1142/S0219061322500180.
- [Dod15] Pandelis Dodos. Some recent results in Ramsey theory. *Zbornik Radova*, 17:81–91, 2015.
- [DR12] Domingos Dellamonica and Vojtěch Rödl. Distance preserving Ramsey graphs. *Combinatorics, Probability and Computing*, 21(04):554–581, 2012.
- [DZ23] Natasha Dobrinen and Andy Zucker. Infinite-dimensional ramsey theory for binary free amalgamation classes. *arXiv:2303.04246*, 2023.
- [EGT16] David M. Evans, Zaniar Ghadernezhad, and Katrin Tent. Simplicity of the automorphism groups of some Hrushovski constructions. *Annals of Pure and Applied Logic*, 167(1):22–48, 2016.

- [EH74] Paul Erdős and András Hajnal. Unsolved and solved problems in set theory. In *Proceedings of the Tarski Symposium (Berkeley, Calif., 1971)*, Amer. Math. Soc., Providence, volume 1, pages 269–287, 1974.
- [EHK⁺21] David M. Evans, Jan Hubička, Matěj Konečný, Yibei Li, and Martin Ziegler. Simplicity of the automorphism groups of generalised metric spaces. *Journal of Algebra*, 584:163–179, 2021.
- [EHKN20] David M. Evans, Jan Hubička, Matěj Konečný, and Jaroslav Nešetřil. EPPA for two-graphs and antipodal metric spaces. *Proceedings of the American Mathematical Society*, 148:1901–1915, 2020.
- [EHN19] David M. Evans, Jan Hubička, and Jaroslav Nešetřil. Automorphism groups and Ramsey properties of sparse graphs. *Proceedings of the London Mathematical Society*, 119(2):515–546, 2019.
- [EHN21] David M. Evans, Jan Hubička, and Jaroslav Nešetřil. Ramsey properties and extending partial automorphisms for classes of finite structures. *Fundamenta Mathematicae*, 253:121–153, 2021.
- [ER56] Paul Erdős and Richard Rado. A partition calculus in set theory. *Bulletin of the American Mathematical Society*, 62:427–489, 1956.
- [Eva] David M. Evans. Homogeneous structures, ω -categoricity and amalgamation constructions. Notes on a Minicourse given at HIM, Bonn, September 2013.
- [EZS89] Mohamed El-Zahar and Norbert W. Sauer. The indivisibility of the homogeneous kn-free graphs. *Journal of Combinatorial Theory, Series B*, 47(2):162–170, 1989.
- [FGR87] Peter Frankl, Ronald L Graham, and Vojtech Rödl. Induced restricted Ramsey theorems for spaces. *Journal of Combinatorial Theory, Series A*, 44(1):120–128, 1987.
- [Fol70] Jon Folkman. Graphs with monochromatic complete subgraphs in every edge coloring. *SIAM Journal on Applied Mathematics*, 18(1):19–24, 1970.
- [Fra53] Roland Fraïssé. Sur certaines relations qui généralisent l’ordre des nombres rationnels. *Comptes Rendus de l’Academie des Sciences*, 237:540–542, 1953.
- [Gal68] Fred Galvin. Partition theorems for the real line. *Notices Of The American Mathematical Society*, 15:660, 1968.
- [Gal69] Fred Galvin. Errata to “Partition theorems for the real line”. *Notices Of The American Mathematical Society*, 16:1095, 1969.
- [Gar76] Anthony Gardiner. Homogeneous graphs. *Journal of Combinatorial Theory, Series B*, 20(1):94–102, 1976.

- [GGL95] Ronald L. Graham, M. Grotschel, and László Lovász. *Handbook of combinatorics*, volume 1. North Holland, 1995.
- [Gow92] W. Timothy Gowers. Lipschitz functions on classical spaces. *European Journal of Combinatorics*, 13(3):141–151, 1992.
- [GR71] Ronald L. Graham and Bruce L. Rothschild. Ramsey’s theorem for n -parameter sets. *Transactions of the American Mathematical Society*, 159:257–292, 1971.
- [GR01] Chris Godsil and Gordon Royle. *Algebraic Graph Theory*, volume 207 of *Graduate Texts in Mathematics*. volume 207 of Graduate Texts in Mathematics. Springer, 2001.
- [GRS90] Ronald L. Graham, Bruce L. Rothschild, and Joel H. Spencer. *Ramsey Theory*. A Wiley-Interscience publication. Wiley, 1990.
- [Hal49] Marshall Hall. Coset representations in free groups. *Transactions of the American Mathematical Society*, 67(2):421–432, 1949.
- [Hal50] Marshall Hall. A topology for free groups and related groups. *Annals of Mathematics*, 52:127, 1950.
- [Her95] Bernhard Herwig. Extending partial isomorphisms on finite structures. *Combinatorica*, 15(3):365–371, 1995.
- [Her98] Bernhard Herwig. Extending partial isomorphisms for the small index property of many ω -categorical structures. *Israel Journal of Mathematics*, 107(1):93–123, 1998.
- [H HLS93] Wilfrid Hodges, Ian Hodkinson, Daniel Lascar, and Saharon Shelah. The small index property for ω -stable ω -categorical structures and for the random graph. *Journal of the London Mathematical Society*, 2(2):204–218, 1993.
- [HJ63] Alfred W. Hales and Robert I. Jewett. Regularity and positional games. *Transactions of the American Mathematical Society*, 106:222–229, 1963.
- [HJKS19] Jan Hubička, Colin Jahel, Matěj Konečný, and Marcin Sabok. Extending partial automorphisms of n -partite tournaments. *Acta Mathematica Universitatis Comenianae*, 88(3):807–811, 2019. Extended abstract for Eurocomb 2021.
- [HJKS23] Jan Hubička, Colin Jahel, Matěj Konečný, and Marcin Sabok. Extension property for partial automorphisms of the n -partite and semi-generic tournaments. To appear, 2023.
- [HKK18] Jan Hubička, Michael Kompatscher, and Matěj Konečný. Forbidden cycles in metrically homogeneous graphs. Accepted to *European Journal of Combinatorics*, arXiv:1808.05177, 2018.

- [HKN18] Jan Hubička, Matěj Konečný, and Jaroslav Nešetřil. Semigroup-valued metric spaces: Ramsey expansions and EPPA. In preparation, 2018.
- [HKN19a] Jan Hubička, Matěj Konečný, and Jaroslav Nešetřil. A combinatorial proof of the extension property for partial isometries. *Commentationes Mathematicae Universitatis Carolinae*, 60(1):39–47, 2019.
- [HKN19b] Jan Hubička, Matěj Konečný, and Jaroslav Nešetřil. Ramsey properties of edge-labelled graphs via completions. *Acta Mathematica Universitatis Comenianae*, 88(3):801–805, 2019.
- [HKN21] Jan Hubička, Matěj Konečný, and Jaroslav Nešetřil. Conant’s generalised metric spaces are Ramsey. *Contributions to Discrete Mathematics*, 16:46–70, 2021.
- [HKN22] Jan Hubička, Matěj Konečný, and Jaroslav Nešetřil. All those EPPA classes (strengthenings of the Herwig–Lascar theorem). *Transactions of the American Mathematical Society*, 375(11):7601–7667, 2022.
- [HKNS20] Jan Hubička, Matěj Konečný, Jaroslav Nešetřil, and Norbert Sauer. Ramsey expansions and EPPA for S -metric spaces. To appear, 2020.
- [HL66] James D. Halpern and Hans Läuchli. A partition theorem. *Transactions of the American Mathematical Society*, 124(2):360–367, 1966.
- [HL00] Bernhard Herwig and Daniel Lascar. Extending partial automorphisms and the profinite topology on free groups. *Transactions of the American Mathematical Society*, 352(5):1985–2021, 2000.
- [HN19] Jan Hubička and Jaroslav Nešetřil. All those Ramsey classes (Ramsey classes with closures and forbidden homomorphisms). *Advances in Mathematics*, 356C:106791, 2019.
- [HO03] Ian Hodkinson and Martin Otto. Finite conformal hypergraph covers and Gaifman cliques in finite structures. *Bulletin of Symbolic Logic*, 9(03):387–405, 2003.
- [Hod93] Wilfrid Hodges. *Model theory*, volume 42. Cambridge University Press, 1993.
- [Hod02] Ian Hodkinson. Finite model property for guarded fragments. slides available at <http://www.cllc.vuw.ac.nz/LandCtalks/imhslides.pdf>, 2002.
- [HPSW19] Jingyin Huang, Micheal Pawliuk, Marcin Sabok, and Daniel T Wise. The Hrushovski property for hypertournaments and profinite topologies. *Journal of the London Mathematical Society*, 100(3):757–774, 2019.

- [Hru92] Ehud Hrushovski. Extending partial isomorphisms of graphs. *Combinatorica*, 12(4):411–416, 1992.
- [Hub20a] Jan Hubička. Big Ramsey degrees using parameter spaces. *arXiv:2009.00967*, 2020.
- [Hub20b] Jan Hubička. Structural Ramsey Theory and the Extension Property for Partial Automorphisms. Introduction to habilitation thesis, arXiv:2010.05041, 2020.
- [IS06] Trevor Irwin and Sławomir Solecki. Projective Fraïssé limits and the pseudo-arc. *Transactions of the American Mathematical Society*, 358(7):3077–3096, 2006.
- [Iva15] Aleksander Ivanov. An ω -categorical structure with amenable automorphism group. *Mathematical Logic Quarterly*, 61(4-5):307–314, 2015.
- [JLNVTW14] Jakub Jasiński, Claude Laflamme, Lionel Nguyen Van Thé, and Robert Woodrow. Ramsey precompact expansions of homogeneous directed graphs. *The Electronic Journal of Combinatorics*, 21, 2014.
- [JT22] Colin Jahel and Todor Tsankov. Invariant measures on products and on the space of linear orders. *Journal de l'École polytechnique — Mathématiques*, 9:155–176, 2022.
- [Kar13] Nikolaos Karagiannis. A combinatorial proof of an infinite version of the Hales–Jewett theorem. *Journal of Combinatorics*, 4(2):273–291, 2013.
- [Kec12] A. Kechris. Dynamics of non-archimedean polish groups. In *European Congress of Mathematics Kraków, 2 – 7 July, 2012*, pages 375–397, 2012.
- [Kon18] Matěj Konečný. Combinatorial properties of metrically homogeneous graphs. Bachelor's thesis, Charles University, 2018. arXiv:1805.07425.
- [Kon19] Matěj Konečný. Semigroup-valued metric spaces. Master's thesis, Charles University, 2019. arXiv:1810.08963.
- [Kon20] Matěj Konečný. Extending partial isometries of antipodal graphs. *Discrete Mathematics*, 343(1):111633, 2020.
- [KPR18] Mustapha Kabil, Maurice Pouzet, and Ivo G. Rosenberg. Free monoids and generalized metric spaces. *European Journal of Combinatorics*, 2018.
- [KPT05] Alexander S. Kechris, Vladimir G. Pestov, and Stevo Todorčević. Fraïssé limits, Ramsey theory, and topological dynamics of automorphism groups. *Geometric and Functional Analysis*, 15(1):106–189, 2005.

- [KR86] Péter Komjáth and Vojtech Rödl. Coloring of universal graphs. *Graphs and Combinatorics*, 2(1):55–60, 1986.
- [KR07] Alexander S. Kechris and Christian Rosendal. Turbulence, amalgamation, and generic automorphisms of homogeneous structures. *Proceedings of the London Mathematical Society*, 94(2):302–350, 2007.
- [KS19] Itay Kaplan and Pierre Simon. Automorphism groups of finite topological rank. *Transactions of the American Mathematical Society*, 372(3):2011–2043, 2019.
- [Kub14] Wiesław Kubiś. Fraïssé sequences: category-theoretic approach to universal homogeneous structures. *Annals of Pure and Applied Logic*, 165(11):1755–1811, 2014.
- [Lar08] Jean A. Larson. Counting canonical partitions in the random graph. *Combinatorica*, 28(6):659–678, 2008.
- [Lav84] Richard Laver. Products of infinitely many perfect trees. *Journal of the London Mathematical Society*, 2(3):385–396, 1984.
- [Lee73] Klaus Leeb. Vorlesungen über Pascaltheorie. *Arbeitsberichte des Instituts für Mathematische Maschinen und Datenverarbeitung*, 6, 1973.
- [Li18] Yibei Li. Simplicity of the automorphism groups of some binary homogeneous structures determined by triangle constraints. *arXiv:1806.01671*, 2018.
- [Li19] Yibei Li. Automorphism groups of homogeneous structures with stationary weak independence relations. *arXiv:1911.08540*, 2019.
- [Li20a] Yibei Li. Automorphism groups of linearly ordered homogeneous structures. *arXiv:2009.02475*, 2020.
- [Li20b] Yibei Li. *Normal Subgroups of the Automorphism Groups of Some Homogeneous Structures*. PhD thesis, Imperial College London, 2020.
- [LNVTS10] Claude Laflamme, Lionel Nguyen Van Thé, and Norbert W. Sauer. Partition properties of the dense local order and a colored version of Milliken’s theorem. *Combinatorica*, 30(1):83–104, 2010.
- [LSV06] Claude Laflamme, Norbert W. Sauer, and Vojkan Vuksanovic. Canonical partitions of universal structures. *Combinatorica*, 26(2):183–205, 2006.
- [LW80] Alistair H. Lachlan and Robert E. Woodrow. Countable ultrahomogeneous undirected graphs. *Transactions of the American Mathematical Society*, pages 51–94, 1980.

- [Mac66] George W. Mackey. Ergodic theory and virtual groups. *Mathematische Annalen*, 166(3):187–207, 1966.
- [Mac11] Dugald Macpherson. A survey of homogeneous structures. *Discrete Mathematics*, 311(15):1599–1634, 2011. Infinite Graphs: Introductions, Connections, Surveys.
- [Maš18] Dragan Mašulović. Pre-adjunctions and the Ramsey property. *European Journal of Combinatorics*, 70:268–283, 2018.
- [Maš20] Dragan Mašulović. Finite big Ramsey degrees in universal structures. *Journal of Combinatorial Theory, Series A*, 170:105137, 2020.
- [Maš21] Dragan Mašulović. The Kechris–Pestov–Todorčević Correspondence from the Point of View of Category Theory. *Applied Categorical Structures*, 29(1):141–169, Feb 2021.
- [Mil79] Keith R. Milliken. A Ramsey theorem for trees. *Journal of Combinatorial Theory, Series A*, 26(3):215–237, 1979.
- [MNVTT15] Julien Melleray, Lionel Nguyen Van Thé, and Todor Tsankov. Polish groups with metrizable universal minimal flows. *International Mathematics Research Notices*, page rnv171, 2015.
- [MT11] Dugald Macpherson and Katrin Tent. Simplicity of some automorphism groups. *Journal of Algebra*, 342(1):40–52, 2011.
- [MT14] Julien Melleray and Todor Tsankov. Extremely amenable groups via continuous logic. *arXiv:1404.4590*, 2014.
- [MZ22] Dragan Mašulović and Andy Zucker. From ramsey degrees to ramsey expansions via weak amalgamation. *arXiv:2208.07351*, 2022.
- [Neš89] Jaroslav Nešetřil. For graphs there are only four types of hereditary Ramsey classes. *Journal of Combinatorial Theory, Series B*, 46(2):127–132, 1989.
- [Neš95] Jaroslav Nešetřil. Ramsey theory. In Ronald L. Graham, Martin Grötschel, and László Lovász, editors, *Handbook of Combinatorics*, volume 2, pages 1331–1403. MIT Press, Cambridge, MA, USA, 1995.
- [Neš05] Jaroslav Nešetřil. Ramsey classes and homogeneous structures. *Combinatorics, probability and computing*, 14(1-2):171–189, 2005.
- [Neš07] Jaroslav Nešetřil. Metric spaces are Ramsey. *European Journal of Combinatorics*, 28(1):457–468, 2007.
- [NR75] Jaroslav Nešetřil and Vojtěch Rödl. Partitions of subgraphs. In Miroslav Fiedler, editor, *Recent advances in graph theory*, pages 413–423. Prague, 1975.

- [NR76a] Jaroslav Nešetřil and Vojtěch Rödl. Partitions of vertices. *Commentationes Mathematicae Universitatis Carolinae*, 17(1):85–95, 1976.
- [NR76b] Jaroslav Nešetřil and Vojtěch Rödl. The Ramsey property for graphs with forbidden complete subgraphs. *Journal of Combinatorial Theory, Series B*, 20(3):243–249, 1976.
- [NR77a] Jaroslav Nešetřil and Vojtěch Rödl. A structural generalization of the Ramsey theorem. *Bulletin of the American Mathematical Society*, 83(1):127–128, 1977.
- [NR77b] Jaroslav Nešetřil and Vojtěch Rödl. Partitions of finite relational and set systems. *Journal Combinatorial Theory, Series A*, 22(3):289–312, 1977.
- [NR79] Jaroslav Nešetřil and Vojtěch Rödl. A short proof of the existence of highly chromatic hypergraphs without short cycles. *Journal of Combinatorial Theory, Series B*, 27(2):225–227, 1979.
- [NR81] Jaroslav Nešetřil and Vojtěch Rödl. Simple proof of the existence of restricted Ramsey graphs by means of a partite construction. *Combinatorica*, 1(2):199–202, 1981.
- [NR82] Jaroslav Nešetřil and Vojtěch Rödl. Two proofs of the Ramsey property of the class of finite hypergraphs. *European Journal of Combinatorics*, 3(4):347–352, 1982.
- [NR83] Jaroslav Nešetřil and Vojtěch Rödl. Ramsey classes of set systems. *Journal of Combinatorial Theory, Series A*, 34(2):183–201, 1983.
- [NR84] Jaroslav Nešetřil and Vojtěch Rödl. Combinatorial partitions of finite posets and lattices—Ramsey lattices. *Algebra Universalis*, 19(1):106–119, 1984.
- [NR87] Jaroslav Nešetřil and Vojtěch Rödl. Strong Ramsey theorems for Steiner systems. *Transactions of the American Mathematical Society*, 303(1):183–192, 1987.
- [NR89] Jaroslav Nešetřil and Vojtěch Rödl. The partite construction and Ramsey set systems. *Discrete Mathematics*, 75(1):327–334, 1989.
- [NR90] Jaroslav Nešetřil and Vojtěch Rödl. Partite construction and Ramsey space systems. In *Mathematics of Ramsey Theory*, volume 5 of *Algorithms and Combinatorics*, pages 98–112. Springer, 1990.
- [NR12] Jaroslav Nešetřil and Vojtěch Rödl. *Mathematics of Ramsey Theory (Algorithms and Combinatorics)*. Springer, 2012.
- [NR18] Jaroslav Nešetřil and Vojtěch Rödl. Ramsey partial orders from acyclic graphs. *Order*, 35(2):293–300, 2018.

- [NVT08] Lionel Nguyen Van Thé. Big Ramsey degrees and divisibility in classes of ultrametric spaces. *Canadian Mathematical Bulletin*, 51(3):413–423, 2008.
- [NVT09] Lionel Nguyen Van Thé. Ramsey degrees of finite ultrametric spaces, ultrametric Urysohn spaces and dynamics of their isometry groups. *European Journal of Combinatorics*, 30(4):934–945, 2009.
- [NVT10] Lionel Nguyen Van Thé. *Structural Ramsey Theory of Metric Spaces and Topological Dynamics of Isometry Groups*. Memoirs of the American Mathematical Society. American Mathematical Society, 2010.
- [NVT13] Lionel Nguyen Van Thé. More on the Kechris–Pestov–Todorćević correspondence: Precompact expansions. *Fundamenta Mathematicae*, 222:19–47, 2013.
- [NVT15] Lionel Nguyen Van Thé. A survey on structural Ramsey theory and topological dynamics with the Kechris–Pestov–Todorćević correspondence in mind. *Selected Topics in Combinatorial Analysis*, 17(25):189–207, 2015.
- [NVTS09] Lionel Nguyen Van Thé and Norbert W. Sauer. The Urysohn sphere is oscillation stable. *Geometric and Functional Analysis*, 19(2):536–557, Sep 2009.
- [Ott20] Martin Otto. Amalgamation and symmetry: From local to global consistency in the finite. arXiv:1709.00031, 2020.
- [Pes08] Vladimir G. Pestov. A theorem of Hrushovski–Solecki–Vershik applied to uniform and coarse embeddings of the Urysohn metric space. *Topology and its Applications*, 155(14):1561–1575, 2008.
- [Prö85] Hans Jürgen Prömel. Induced partition properties of combinatorial cubes. *Journal of Combinatorial Theory, Series A*, 39(2):177–208, 1985.
- [Prö13] Hans Jürgen Prömel. *Ramsey Theory for Discrete Structures*. Springer International Publishing, 2013.
- [PS96] Maurice Pouzet and Norbert Sauer. Edge partitions of the rado graph. *Combinatorica*, 16:505–520, 12 1996.
- [PS18] Gianluca Paolini and Saharon Shelah. The automorphism group of Hall’s universal group. *Proceedings of the American Mathematical Society*, 146(4):1439–1445, 2018.
- [PTW85] Madeleine Paoli, William T. Trotter, and James W. Walker. Graphs and orders in Ramsey theory and in dimension theory. In Ivan Rival, editor, *Graphs and Order*, volume 147 of *NATO AST series*, pages 351–394. Springer, 1985.

- [Ram30] Frank Plumpton Ramsey. On a problem of formal logic. *Proceedings of the London Mathematical Society*, 2(1):264–286, 1930.
- [Ros11a] Christian Rosendal. Finitely approximable groups and actions part II: Generic representations. *The Journal of Symbolic Logic*, 76(04):1307–1321, 2011.
- [Ros11b] Christian Rosendal. Finitely approximate groups and actions part I: The Ribes–Zaleskiĭ property. *The Journal of Symbolic Logic*, 76(04):1297–1306, 2011.
- [RZ93] Luis Ribes and Pavel A Zaleskiĭ. On the profinite topology on a free group. *Bulletin of the London Mathematical Society*, 25(1):37–43, 1993.
- [Sab17] Marcin Sabok. Automatic continuity for isometry groups. *Journal of the Institute of Mathematics of Jussieu*, pages 1–30, 2017.
- [Sau98] Norbert W. Sauer. Edge partitions of the countable triangle free homogeneous graph. *Discrete Mathematics*, 185(1-3):137–181, 1998.
- [Sau02] Norbert Sauer. A ramsey theorem for countable homogeneous directed graphs. *Discrete Mathematics*, 253:45–61, 06 2002.
- [Sau06] Norbert W. Sauer. Coloring subgraphs of the Rado graph. *Combinatorica*, 26(2):231–253, 2006.
- [Sau12] Norbert W. Sauer. Vertex partitions of metric spaces with finite distance sets. *Discrete Mathematics*, 312(1):119–128, 2012.
- [Sau13a] Norbert W. Sauer. Distance sets of Urysohn metric spaces. *Canadian Journal of Mathematics*, 65(1):222–240, 2013.
- [Sau13b] Norbert W. Sauer. Oscillation of Urysohn type spaces. In *Asymptotic geometric analysis*, pages 247–270. Springer, 2013.
- [Sau20] Norbert W. Sauer. Colouring homogeneous structures. *arXiv:2008.02375*, 2020.
- [Sei73] Johan Jacob Seidel. A survey of two-graphs. *Colloquio Internazionale sulle Teorie Combinatorie (Rome, 1973)*, 1:481–511, 1973.
- [Sie33] Waclaw Sierpiński. Sur un problème de la théorie des relations. *Annali della Scuola Normale Superiore di Pisa - Scienze Fisiche e Matematiche*, 2e série, 2(3):285–287, 1933.
- [Sin17] Daoud Siniora. *Automorphism Groups of Homogeneous Structures*. PhD thesis, University of Leeds, March 2017.
- [Sok12] Miodrag Sokić. Ramsey properties of finite posets. *Order*, 29(1):1–30, 2012.

- [Sok15] Miodrag Sokić. Semilattices and the Ramsey property. *The Journal of Symbolic Logic*, 80(4):1236–1259, 2015.
- [Sol05] Sławomir Solecki. Extending partial isometries. *Israel Journal of Mathematics*, 150(1):315–331, 2005.
- [Sol09] Sławomir Solecki. Notes on a strengthening of the Herwig–Lascar extension theorem. Unpublished note, 2009.
- [Sol10] Sławomir Solecki. A Ramsey theorem for structures with both relations and functions. *Journal of Combinatorial Theory, Series A*, 117(6):704–714, 2010.
- [SS19] Daoud Siniora and Sławomir Solecki. Coherent extension of partial automorphisms, free amalgamation, and automorphism groups. *The Journal of Symbolic Logic*, pages 199–223, 2019.
- [SZ17] Sławomir Solecki and Min Zhao. A Ramsey theorem for partial orders with linear extensions. *European Journal of Combinatorics*, 60:21–30, 2017.
- [Tho91] Simon Thomas. Reducts of the random graph. *The Journal of symbolic logic*, 56(01):176–181, 1991.
- [Tod10] Stevo Todorćević. *Introduction to Ramsey spaces*, volume 174. Princeton University Press, 2010.
- [Tru85] J. K. Truss. The group of the countable universal graph. *Mathematical Proceedings of the Cambridge Philosophical Society*, 98(2):213–245, 1985.
- [Tru89] J. K. Truss. Infinite permutation groups II. Subgroups of small index. *Journal of Algebra*, 120(2):494–515, 1989.
- [Tru92] J. K. Truss. Generic automorphisms of homogeneous structures. *Proceedings of the London Mathematical Society*, s3-65(1):121–141, 1992.
- [TZ13a] Katrin Tent and Martin Ziegler. The isometry group of the bounded Urysohn space is simple. *Bulletin of the London Mathematical Society*, 45(5):1026–1030, 2013.
- [TZ13b] Katrin Tent and Martin Ziegler. On the isometry group of the Urysohn space. *Journal of the London Mathematical Society*, 87(1):289–303, 2013.
- [Ury27] Paul S. Urysohn. Sur un espace métrique universel. *Bulletin des Sciences Mathématiques*, 51(2):43–64, 1927.
- [Ver08] Anatoly M. Vershik. Globalization of the partial isometries of metric spaces and local approximation of the group of isometries of Urysohn space. *Topology and its Applications*, 155(14):1618–1626, 2008.

- [Voi80] Bernd Voigt. The partition problem for finite abelian groups. *Journal of Combinatorial Theory, Series A*, 28(3):257–271, 1980.
- [Zuc16] Andy Zucker. Topological dynamics of automorphism groups, ultrafilter combinatorics, and the generic point problem. *Transactions of the American Mathematical Society*, 368(9):6715–6740, 2016.
- [Zuc19] Andy Zucker. Big Ramsey degrees and topological dynamics. *Groups, Geometry, and Dynamics*, 13(1):235–276, 2019.
- [Zuc22] Andy Zucker. On big Ramsey degrees for binary free amalgamation classes. *Advances in Mathematics*, 408:108585, 2022.