On the Caccetta-Häggkvist Conjecture with Forbidden Subgraphs

Alexander A. Razborov*

June 21, 2011

Abstract

The Caccetta-Häggkvist conjecture made in 1978 asserts that every orgraph on n vertices without oriented cycles of length $\leq \ell$ must contain a vertex of outdegree at most $\frac{n-1}{\ell}$. It has a rather elaborate set of (conjectured) extremal configurations.

In this paper we consider the case $\ell = 3$ that received quite a significant attention in the literature. We identify three orgraphs on four vertices that are missing as an induced subgraph in all known extremal examples and prove the Caccetta-Häggkvist conjecture for orgraphs missing as induced subgraphs any of these orgraphs, along with $\vec{C}_3$. Using a standard trick, we can also lift the restriction of being induced, although this makes graphs in our list slightly more complicated.

1. Introduction

One prominent way to attack a difficult problem in extremal combinatorics is by better understanding the nature of its (conjectured) extremal configurations. What one would hope for is to find some property P, as "natural" as possible that is shared by all known extremal configurations, and then solve

*University of Chicago, razborov@cs.uchicago.edu. Part of this work was done while the author was at Steklov Mathematical Institute, supported by the Russian Foundation for Basic Research, and at Toyota Technological Institute, Chicago.

the extremal problem in question for all configurations possessing this property P. Arguably but conceivably, this may shed some light on the nature of difficulties surrounding the problem in question and perhaps even open up a possibility to solve the problem by gradually lifting constraints defining the property P. For the famous Turan's (3,4)-problem this approach was recently undertaken by the author in [Raz10a, Raz10b]; another good example of this sort is the recent solution of the local Sidorenko conjecture by Lovász [Lov10].

In this paper we address along these lines another major open problem in the area, Caccetta-Häggkvist conjecture, that is nearly as famous as those mentioned above. Recall that we are given an oriented graph (i.e., a digraph without loops, parallel or anti-parallel edges) G on n vertices that does not contain (oriented) cycles of length $\leq \ell$ or, in other words, has girth $\geq \ell + 1$. Behzad, Chartrand and Wall [BCW70] asked the following question: if G is additionally known to be bi-regular, how large can be its degree? They conjectured that the answer is $\lfloor \frac{n-1}{\ell} \rfloor$ and presented a simple construction attaining this bound. Eight years later, Caccetta and Häggkvist [CH78] proposed to lift in this conjecture the restriction of bi-regularity and, moreover, restrict attention to minimal *outdegree* only. In other words, they asked if every orgraph without oriented cycles of length $\leq \ell$ must contain a vertex of out-degree $\leq \frac{n-1}{\ell}$, and it is this question that became known as the *Caccetta-Häggkvist conjecture*. It turned out to be notoriously difficult, too.

The case of higher values of ℓ was studied in [CS83, Ham87, HR87, Nis72, She00, She02].

In this paper we concentrate on the case $\ell = 3$, as much of the previous work in this area did. Let c be the minimal constant for which the asymptotic upper bound $(c + o(1))n$ on the minimal outdegree in $\vec{C}_3$-free orgraphs holds. Caccetta and Häggkvist themselves proved in [CH78] that $c \leq \frac{3-\sqrt{5}}{2} \approx 0.382$. This was improved in the series of papers [Bon97, She98, HHK07] to the current record of $c \leq 0.3465n$ [HKN09].

As we already noticed, the first example of a graph G on n vertices without copies of $\vec{C}_3$ and minimal degree $\lfloor \frac{n-1}{3} \rfloor$ was given in the paper [BCW70]. It is quite simple: assuming that $n = 3h + 1$ for some integer h, we let $\mathbb{Z}_{3h+1}$ be the set of vertices, and we connect i to j if and only if $j - i \leq h \mod (3h + 1)$. But this example is definitely not unique: another one was provided in [Cha05, Section 5.1], and it seems to be a part of the common knowledge that even more examples of orgraphs with this property can be

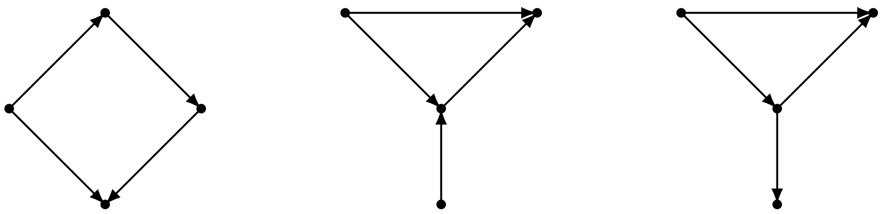

Figure 1: Forbidden orgraphs.

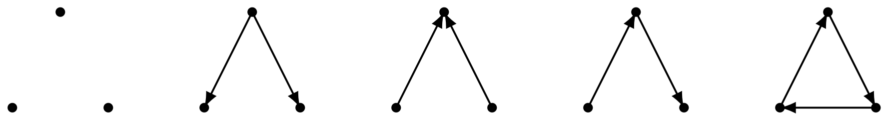

Figure 2: Some orgraphs on 3 vertices.

constructed by combining the ideas of these two.

The first (minor) contribution of our paper consists in an attempt to systemize this common knowledge and present in Section 2, without any claims as to its novelty, what we believe to be the complete set of currently known extremal configurations for the Caccetta-Häggkvist conjecture.

All these examples (for the case $\ell = 3$) have the property that they are missing (as induced subgraphs) the three orgraphs shown on Figure 1. As our main result, we prove the CH-conjecture (for $\ell = 3$) for any $\vec{C}_3$-free orgraph with this additional property (Theorem 3.1). While this is the first result of this kind pertaining to *all* known extremal configurations, we would like to mention some previous (unpublished) work regarding forbidden orgraphs on 3 vertices that are missing in the original "cyclic" configuration by Behzad, Chartrand and Wall. On Figure 2, these are represented by $I_3, \vec{K}_{1,2}$ and $\vec{K}_{2,1}$.

The CH-conjecture (as always, for $\ell = 3$) is an easy exercise for orgraphs missing $\vec{C}_3$ and $\vec{K}_{1,2}$ as induced subgraphs. Under the additional assumption of out-regularity, Chudnovsky and Seymour [CS06] did the case when $\vec{C}_3$ and I_3 are missing; to the best of our knowledge, the question is still open without the restriction of out-regularity. Seymour [Sey06] proved the CH-conjecture for orgraphs missing $\vec{C}_3$ and $\vec{K}_{2,1}$ (which is substantially more difficult than the dual case of $\vec{C}_3$ and $\vec{K}_{1,2}$).

Potential usefulness of Theorem 3.1 (at least, of the sort we can think of) is undermined by the fact that it involves the notion of an *induced* subgraph. We include a very simple observation (Theorem 3.2) showing that this restriction can be removed at the expense of the forbidden family becoming slightly more complicated (see Figure 3).

The proof of Theorem 3.1 was found in the framework of flag algebras [Raz07]. But the final calculation (see the crucial Claim 5.6) does not use multiplicative structure (and, in particular, does not use Cauchy-Schwarz inequality). This makes working in the limit framework unnecessary, and in this paper we adopt a compromise approach. Namely, we exclusively work with finite objects but still use basic elements of the apparatus of flag algebras that in our case boils down to two conventions:

- systematic and consistent notation for various sets based upon types and flags
- systematic measurement of all necessary quantities in terms of their "densities" rather than absolute size.

We would like to note that even with this compromise approach lower-order terms do begin to accumulate in the proof of Claim 5.6, and we can not simply dismiss them due to the inductive nature of the argument. Fortunately, the proof is over before they become a real nuisance.

2. Extremal configurations

We let $[n] \stackrel{\text{def}}{=} \{1, 2, \dots, n\}$.

An *oriented graph*, or an *orgraph*, is a directed graph without loops and such that every pair of vertices is connected by at most one edge, regardless of direction. $V(\Gamma)$ is the set of vertices of an orgraph Γ, and $E(\Gamma)$ is its set of edges. For an edge $\langle v, w\rangle$, w is its *head* and v is its *tail*. $\mathrm{od}_\Gamma(v)$ is the out-degree of the vertex v. Vertices v and w are *independent* in Γ if neither $\langle v, w\rangle$ nor $\langle w, v\rangle$ is an edge. For an orgraph Γ and $V \subseteq V(\Gamma)$, $\Gamma|_V$ is the orgraph induced by V.

$\vec{C}_3$ is the cycle on 3 vertices, and an orgraph Γ is $\vec{C}_3$*-free* if it does not contain copies of $\vec{C}_3$.

Conjecture 1 (Caccetta-Häggkvist Conjecture) Any $\vec{C}_3$-free orgraph Γ on n vertices contains a vertex v with $\mathrm{od}_\Gamma(v) \leq \frac{n-1}{3}$.

We will sometimes refer to this as to "the CH-conjecture".

In this section we review what we believe to be the complete list of known configurations attaining this value; as we noted in Introduction, this is a rather straightforward generalization of the examples found in [BCW70, Cha05]. There are two legitimate frameworks in which this question can be addressed; one of them is exact (i.e., describing finite orgraphs matching the bound in Conjecture 1 precisely). And another is asymptotic: it can be best described in terms of (or)graphons [LS06] or flag algebras [Raz07], but intuitively it corresponds to "convergent" sequences of orgraphs $\{\Gamma_m\}$ with $\min_v \mathrm{od}_{\Gamma_m}(v) \geq \left(\frac{1}{3} - o(1)\right) |V(\Gamma_m)|$. We treat them simultaneously.

Let S^1 be the unit circle, and define the (infinite) orgraph Γ_0 with $V(\Gamma_0) \stackrel{\mathrm{def}}{=} S^1$ and $E(\Gamma_0) \stackrel{\mathrm{def}}{=} \{\langle x, y\rangle \mid y - x < 1/3 \mod 1\}$. Note that Γ_0 is $\vec{C}_3$-free. Let $\Omega \stackrel{\mathrm{def}}{=} (S^1)^\infty$ be the infinite-dimensional torus. We let Γ_{CH} be the orgraph with $V(\Gamma_{\mathrm{CH}}) = \Omega$ that is the lexicographic product of countably many copies of Γ_0. In other words, for any two different vertices $\mathsf{x} = (x_1, x_2, \ldots, x_n, \ldots)$, $\mathsf{y} = (y_1, \ldots, y_n \ldots) \in \Omega$ we choose the minimal d for which $x_d \neq y_d$ and let $\langle \mathsf{x}, \mathsf{y}\rangle \in E(\Gamma_{\mathrm{CH}})$ if and only if $\langle x_d, y_d\rangle \in E(\Gamma_0)$.

Ω is a topological space (under product topology), and therefore every probability measure μ on its Borel subsets gives rise to an oriented graphon [LS06], as well as to a homomorphism $\phi \in \mathrm{Hom}^+(\mathcal{A}^0[T_{\mathrm{CH}}], \mathbb{R})$ [Raz07], where T_{CH} is the theory of $\vec{C}_3$-free orgraphs. We now describe those measures μ that lead to asymptotically extremal examples for Conjecture 1.

Fix a probability measure μ on Borel subsets of Ω. Every finite string $(a^1, \ldots, a^d) \in (S^1)^d$ defines the canonical closed set $\Omega_a = \{\mathsf{x} \in \Omega \mid x_1 = a_1, \ldots, x_d = a_d\}$. Whenever $\mu(\Omega_a) > 0$, we have the conditional measure μ_a on Ω_a ($\mu_a(X) \stackrel{\mathrm{def}}{=} \frac{\mu(X)}{\mu(\Omega_a)}$, $X \subseteq \Omega_a$) and then the pushforward measure $\widehat{\mu}_a$ on S^1 defined by projecting Ω_a onto the $(d+1)$st coordinate. Let us call the measure μ *extremal* if for every prefix a for which $\mu(\Omega_a) > 0$, this measure $\widehat{\mu}_a$ has one of the following two forms:

- uniform (Lebesgue) measure on S^1;
- uniform discrete measure on the set $\left\{\frac{0}{3h+1}, \frac{1}{3h+1}, \ldots, \frac{3h}{3h+1}\right\}$ for some integer $h \geq 1$.

A combinatorial way to visualize an extremal measure μ is by a locally finite rooted tree in which every non-leaf node has outdegree $(3h+1)$ for some h; the first case (of Lebesgue measure) corresponds to a leaf.

Claim 2.1 *For any extremal measure μ on Ω and for any $\mathsf{x} \in \Omega$,*

$$\mu(\{\mathsf{y} \in \Omega \mid \langle \mathsf{x}, \mathsf{y} \rangle \in E(\Gamma_{CH})\}) = 1/3.$$

Proof. $\{\mathsf{y} \in \Omega \mid \langle \mathsf{x}, \mathsf{y} \rangle \in E(\Gamma_{\mathrm{CH}})\}$ splits as the disjoint union $\bigcup_{d=1}^{\infty} V_d$, where V_d is the set of all y such that $x_1 = y_1, \ldots, x_{d-1} = y_{d-1}$ and $y_d - x_d \mod 1 \in (0, 1/3)$. Our restriction on the measures $\widehat{\mu}_a$ implies that $\mu(V_d) = \frac{1}{3}\left(\mu(\Omega_{x_1,\ldots,x_{d-1}}) - \mu(\Omega_{x_1,\ldots,x_d})\right)$. Summing over all d and noting that $\mu(\Omega_{x_1,\ldots,x_d}) \leq 4^{-d}$ (and hence $\lim_{d\to\infty} \mu(\Omega_{x_1,\ldots,x_d}) = 0$) gives the result.■

This collection of oriented graphons describes what we believe to be the complete set of all known extremal configurations for Conjecture 1 (more precisely, in the terminology of [Raz07, §4.1], the set of all homomorphisms $\phi \in \mathrm{Hom}(\mathcal{A}^0[T_{\mathrm{CH}}], \mathbb{R})$ with $\delta_\alpha(\phi) = 1/3$). If we additionally require the set $\{a \mid \mu(\Omega_a) > 0\}$ to be finite, we arrive at (again, to the best of our knowledge) the set of all known finite but in general *weighted* extremal orgraphs. Vertices correspond to leaves of the representing tree, and if the product of degrees is the same along all terminal paths, then the measure on leaves becomes uniform, and this gives us (apparently) all known extremal configurations that are *ordinary* (unweighted) orgraphs.

One obvious way to ensure the last uniformity property is by requiring that the tree is balanced and all outdegrees are the same on each level. But there are more sophisticated ways to arrive at a tree with the required property. For example, some vertices on the first level (we place the root onto level zero) may have $(3g+1)(3h+1)$ leaves as their children, while others may branch to a balanced tree of depth 2 with outdegrees $(3g+1)$ on the first level and $(3h+1)$ on the second (and yet another subtrees may have the same form but with the outdegrees on the first and second level exchanged). These in particular lead to extremal examples that do *not* possess a vertex-transitive group of automorphisms. Nonetheless, all these examples are bi-regular and in particular are also good for the original question asked in [BCW70].

Altogether, there are six $\vec{C}_3$-free orgraphs on four vertices missing in Γ_{CH} as an induced subgraph. Of these, we are interested only in the three shown on Figure 1, and let us first check that they are indeed missing.

Claim 2.2 *None of the three orgraphs on Figure 1 can be realized as an induced subgraph of* Γ_{CH}.

Proof. Let $\mathsf{x}^i \in \Omega$ $(i = 1..4)$ be four different vertices, and let d be the first integer for which there exist $1 \leq i < j \leq 4$ with $x_d^i \neq x_d^j$. The projection onto the dth coordinate defines an equivalence relation $\approx$ on $[4]$ with more than one class and such that if $i \approx j \not\approx i' \approx j'$ then $\langle \mathsf{x}^i, \mathsf{x}^{i'} \rangle \in E(\Gamma_{\mathrm{CH}})$ iff $\langle \mathsf{x}^j, \mathsf{x}^{j'} \rangle \in E(\Gamma_{\mathrm{CH}})$. An easy inspection shows that no non-trivial equivalence relation with these properties exists for any of the orgraphs on Figure 1. Therefore, in fact all x_d^i $(i \in [4])$ are pairwise different, and we actually have an embedding into Γ_0. But Γ_0 does not contain induced copies of $\vec{K}_{1,2}, \vec{K}_{2,1}$ (see Figure 2) as an induced subgraph,while every orgraph on Figure 1 contains one of those. Contradiction.∎

3. Main results

The main result of this paper is the following.

Theorem 3.1 *Let* Γ *be an orgraph on* n *vertices that does not contain either* $\vec{C}_3$ *or any of the three orgraphs on Figure 1 as an induced subgraph. Then* Γ *contains a vertex* v *with* $od_\Gamma(v) \leq \frac{n-1}{3}$.

Before we begin the proof of Theorem 3.1, let us show how to drop the restriction of being induced.

Theorem 3.2 *Assume that the CH-conjecture holds for all* $\vec{C}_3$*-free orgraphs containing at least one of the orgraphs*[1] *on Figure 3 as a (not necessarily induced) subgraph. Then the CH-conjecture holds for all* $\vec{C}_3$*-free orgraphs.*

Proof. Let us describe first how this list of orgraphs was generated from Figure 1. For every orgraph Γ on that figure, we took all ordered pairs of independent vertices v, w such that there is no vertex x with $\langle v, x\rangle, \langle x, w\rangle \in E(\Gamma)$. And then for every such pair we added one new vertex x with precisely these edges. On Figure 3, these auxiliary vertices are encircled.

Assume now that the CH-conjecture holds for all graphs that contain at least one orgraph on Figure 3. Let Γ be an arbitrary $\vec{C}_3$-free orgraph; we

[1] Encircled on this figure are those vertices that are new w.r.t. Figure 1.

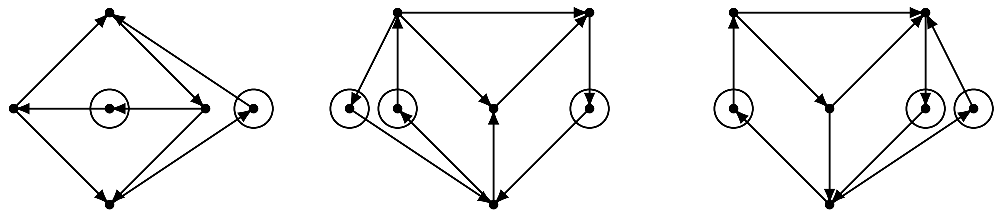

Figure 3: Another set of forbidden orgraphs

want to show the existence of a vertex v with $\mathrm{od}_\Gamma(v) \leq \frac{n-1}{3}$. W.l.o.g. we may assume that adding any new edges to Γ destroys $\vec{C}_3$-freeness, or, in other words, that every pair (v, w) of independent vertices appears as a diagonal of a copy of $\vec{C}_4$ in Γ.

If Γ does not have *induced* copies of the three orgraphs shown on Figure 1, we are done by Theorem 3.1.

Otherwise, our construction and the remark above imply that Γ must contain as a (not necessarily induced) subgraph one of the three orgraphs on Figure 3, except that some of the encircled vertices can be identical. It is, however, easy to see by inspection that identifying any two of them leads either to anti-parallel edges or a copy of $\vec{C}_3$, except for the pair of outer-most vertices on the second or the third orgraphs. But it is easy to see that the result of this identification will contain the first orgraph on Figure 3 and thus does not need a special treatment.

We have shown that Γ must contain one of the three orgraphs on Figure 3, and therefore the required vertex v exists by our assumption.■

We hope that this piece of information about the structure of hypothetical counterexamples to the CH-conjecture that separates them from all known extremal configurations, may turn out helpful, presumably in combination with inductive arguments of the kind that have been already extensively used in previous research on this problem.

The rest of the paper is entirely devoted to the proof of Theorem 3.1. Arguing by induction on the number of vertices, we fix a finite orgraph Γ_0 that does not contain either $\vec{C}_3$ or induced copies of the three orgraphs on Figure 1 and such that the CH-conjecture holds for all its proper *induced* subgraphs. Our goal is to prove it for Γ_0.

4. Flag Algebras

As indicated in Introduction, in this paper we use only a tiny fragment of the whole theory, in the amount of the first four pages of [Raz07, §2.1].

A *type* is a $\vec{C}_3$-orgraph σ with $V(\sigma) = [k]$ for some non-negative integer k called the *size* of σ. A *σ-flag* is a pair $F = (\Gamma, \theta)$, where Γ is a $\vec{C}_3$-free orgraph and $\theta : \sigma \longrightarrow \Gamma$ is an induced embedding. Thus, from the combinatorial point of view, a type is just a (totally) labeled orgraph, and a flag is a partially labeled one. Vertices from $V(\Gamma) \setminus \mathrm{im}(\theta)$ will be sometimes called *free*. If $F = (\Gamma, \theta)$ is a σ-flag and $V \subseteq V(\Gamma)$ contains $\mathrm{im}(\theta)$, then the sub-flag $(\Gamma|_V, \theta)$ will be also denoted by $F|_V$.

A *flag embedding* $\alpha : F \longrightarrow F'$, where $F = (\Gamma, \theta)$ and $F' = (\Gamma', \theta')$ are σ-flags for the same type σ is an induced embedding of orgraphs $\alpha : \Gamma \longrightarrow \Gamma'$ such that $\theta' = \alpha\theta$ (i.e., "label-preserving"). F and F' are *isomorphic* (denoted by $F \approx F'$) if there is a one-to-one flag embedding $\alpha : F \longrightarrow F'$. $\mathcal{F}_\ell^\sigma$ is the set of all σ-flags on ℓ vertices up to an isomorphosm.

If $F \in \mathcal{F}_\ell^\sigma$ and $F' = (\Gamma, \theta) \in \mathcal{F}_L^\sigma$ with $L \geq \ell$, the key quantity in the whole theory is the density $p(F, F')$ of induced copies of F in F' defined as follows. We choose in $V(G')$ uniformly at random a subset $\boldsymbol{V}$ of cardinality ℓ subject to the condition $\boldsymbol{V} \supseteq \mathrm{im}(\theta)$, and let $p(F, F')$ denote the probability of the event $F'|_{\boldsymbol{V}} \approx F$. In almost all calculations used in this paper, $\ell = k + 1$ and hence $\boldsymbol{V}$ can be identified with a single vertex $\boldsymbol{x}$ chosen uniformly at random from $V(\Gamma) \setminus \mathrm{im}(\theta)$.

Recall that we have fixed an orgraph Γ_0 that does not contain either $\vec{C}_3$ or induced copies of the three orgraphs on Figure 1 and such that the CH-conjecture holds for all its proper induced subgraphs. If pairwise different vertices $v_1, \ldots, v_k \in V(\Gamma_0)$ induce a copy of a type σ of size k, then, letting $\theta : [k] \longrightarrow V(\Gamma_0)$ be the corresponding embedding defined by $\theta(i) = v_i$, (Γ_0, θ) becomes a σ-flag, and for another flag F (typically, fixed and very small) we introduce the abbreviation

$$F(v_1, \ldots, v_k) \stackrel{\text{def}}{=} p(F, (\Gamma_0, \theta)). \tag{1}$$

Next, we list concrete types and flags needed for our purposes.

0 and 1 are the unique type of sizes 0 and 1, respectively. A is the type of size 2 with $E(A) \stackrel{\text{def}}{=} \{\langle 1, 2\rangle\}$, and N is the type of size 2 without any edges. P is the type of size 3 with $E(P) \stackrel{\text{def}}{=} \{\langle 1, 2\rangle, \langle 2, 3\rangle\}$.

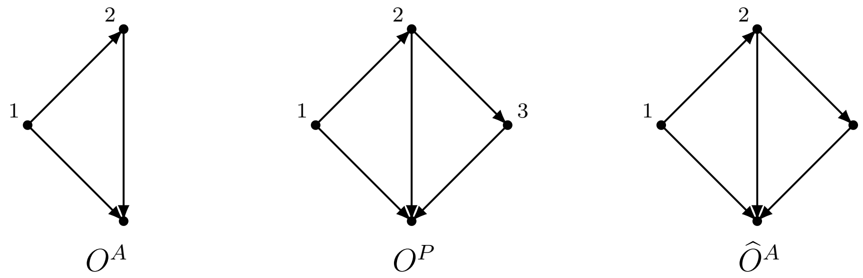

Figure 4: O-flags

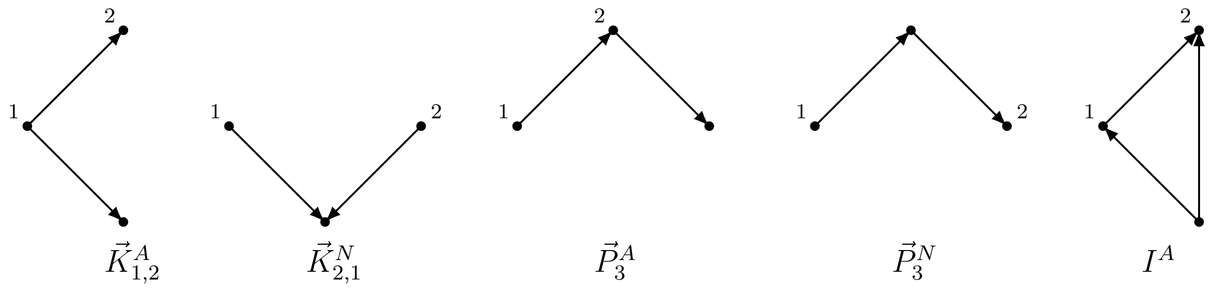

Figure 5: Miscellaneous flags

$\alpha \in \mathcal{F}_2^1$ is a directed edge in which the tail vertex is labeled by 1; our final goal is to find a vertex v with $\alpha(v) \leq 1/3$. For a type σ of size k, let $O^\sigma \in \mathcal{F}^\sigma_{k+1}$ be the flag in which the only free vertex has k incoming edges. Removing from O^P label 3, we will get a flag from $\mathcal{F}_4^A$ that we will denote by $\widehat{O}^A$, see Figure 4. We will also need a few other miscellaneous flags from $\mathcal{F}_3^A, \mathcal{F}_3^N$ shown on Figure 5.

5. Proof of Theorem 3.1

Let us call an edge $\langle v, w\rangle$ *critical* if $O^A(v, w)$ takes the minimal possible value over all edges going out of v. Combinatorially, this means that we are looking at the set $A(v)$ of all out-neighbors of v and pick $w \in A(v)$ to have the smallest possible degree in $\Gamma_0|_{A(v)}$.

Claim 5.1 *For any critical edge $\langle v, w\rangle$, $\widehat{O}^A(v, w) = 0$.*

Proof. Assume the contrary, that is Γ_0 contains a pair of other vertices $x \neq y$ such that $\{\langle v, x\rangle, \langle w, x\rangle, \langle w, y\rangle, \langle y, x\rangle\} \in E(\Gamma_0)$ while v and y are independent. We are going to prove that

$$O^A(v, x) < O^A(v, w), \tag{2}$$

and this will contradict the assumption that $\langle v, w\rangle$ is critical.

For that, let us consider an arbitrary vertex z contributing to $O^A(v, x)$ (that is, such that $\langle v, z\rangle, \langle x, z\rangle \in E(\Gamma_0)$). Since v, x, y, z do not span an In-Pendant (see Figure 1), y and z may not be independent and thus $\langle y, z\rangle \in E(\Gamma_0)$ ($\langle z, y\rangle$ would have created a copy of $\vec{C}_3$). But now since v, w, y, z do not span a Twisted Circle, w and z can not be independent and thus $\langle w, z\rangle \in E(\Gamma_0)$. Which means that z contributes to $O^A(v, w)$ as well.

Finally, let us note that x itself contributes to $O^A(v, w)$. This completes the proof of (2) and gives the desired contradiction with the criticality of $\langle v, w\rangle$.■

In what follows, we agrue by contradiction, i.e. we assume that $\alpha(v) > \frac{1}{3}$ for all $v \in V(\Gamma_0)$.

Claim 5.2 *For any critical edge $\langle v, w\rangle$, $\vec{P}_3^A(v, w) > 0$.*

Proof. Note first that $\alpha(w) = \frac{n-2}{n-1}\left(O^A(v, w) + \vec{P}_3^A(v, w)\right)$. Next, we can apply the inductive assumption to the set of all out-neighbors of v. Since w was chosen to have the minimal degree in $\Gamma_0|_{A(v)}$, $O^A(v, w) \leq \frac{1}{3}\alpha(v) \leq \frac{1}{3}$. The claim follows immediately since $\alpha(w) > 1/3$ by the assumption we have just made.■

Now we study critical paths of length 2.

Claim 5.3 *If $\langle u, v\rangle$ and $\langle v, w\rangle$ are critical edges then u and w are independent.*

Proof. Assume the contrary, that is $\langle u, w\rangle \in E(\Gamma_0)$. By Claim 5.2, there exists x such that $\langle v, x\rangle \in E(\Gamma_0)$ while u and x are independent. Since (u, v, w, x) do not induce an Out-Pendant, w and x may not be independent, and the edge $\langle x, w\rangle$ is ruled out by Claim 5.1. Therefore, $\langle w, x\rangle \in E(\Gamma_0)$.

And now we use the criticality of $\langle v, w\rangle$, and our goal, like in the proof of Claim 5.1, is to arrive at a contradiction by establishing (2). We again choose z with $\langle v, z\rangle, \langle x, z\rangle \in E(\Gamma_0)$. The edge $\langle u, z\rangle$ is again ruled out by Claim 5.1 (applied to $\{u, v, x, z\}$), therefore u and z must be independent.

And now w, z may not be independent (since otherwise $\{u, v, w, z\}$ would have formed an Out-Pendant). Therefore, $\langle w, z\rangle \in E(\Gamma_0)$, and the rest of the proof is the same as the proof of Claim 5.1.■

The flag $\vec{K}_{2,1}^N$ is shown on Figure 5.

Claim 5.4 *If $\langle u, v\rangle$ and $\langle v, w\rangle$ are critical edges then $\vec{K}_{2,1}^N(u, w) = 0$.*

Proof. Assume the contrary, and let x be any vertex such that $\langle u, x\rangle, \langle x, w\rangle \in E(\Gamma_0)$. Then v and x can not be independent (as it would have created a Twisted Circle), $\langle x, v\rangle$ can not be an edge since it would have created $\vec{C}_3$ and $\langle v, x\rangle$ can not be an edge by Claim 5.1.■

Claim 5.5 *If $\langle u, v\rangle$ and $\langle v, w\rangle$ are critical edges then*

$$3O^A(u, v) \leq \vec{P}_3^N(u, w) - \frac{1}{n-2}. \tag{3}$$

Proof. Let $\vec{P}_3^N(u, w) = \frac{h}{n-2}$, and let $U \ni v$ be the corresponding set of vertices, $|U| = h$. Applying to $\Gamma_0|_U$ our inductive assumption, we find a vertex $v^* \in U$ that has degree $\leq \frac{h-1}{3}$ in $\Gamma_0|_U$ (possibly, $v^* = v$). We will prove that

$$O^A(u, v^*) \leq \frac{h-1}{3(n-2)} \tag{4}$$

from which the claim follows since $O^A(u, v) \leq O^A(u, v^*)$ due to the criticality of $\langle u, v\rangle$.

To prove (4), it suffices to show that every vertex x contributing to $O^A(u, v^*)$ in fact belongs to U, that is $\langle x, w\rangle \in E(\Gamma_0)$. But w, x can not be independent (since otherwise we would get a copy of an Out-Pendant), and the edge $\langle w, x\rangle$ is ruled out by Claim 5.4 (note that we must apply this claim to the triple (u, v, w), *not* (u, v^*, w), as we do not know that $\langle u, v^*\rangle$ is critical!).■

The following is our crucial claim that is a typical computer-assisted calculation in flag algebras, albeit much simpler than in all previous applications of the method. For an explanation of all new flags appearing in its statement and proof, we refer to Figure 5.

Claim 5.6 *If $\langle u, v\rangle$ and $\langle v, w\rangle$ are critical edges, then*

$$\left.\begin{array}{r}\alpha(u)+\alpha(v)+\alpha(w)+(O^A(u,v)+I^A(u,v)+\vec{K}^A_{1,2}(u,v))\\ -(O^A(v,w)+I^A(v,w)+\vec{K}^A_{1,2}(v,w))\leq 1.\end{array}\right\} \quad (5)$$

Proof. Subtracting the inequality in Claim 5.5 and re-grouping terms, it suffices to prove that

$$\left.\begin{array}{r}\alpha(u)+\alpha(v)+\alpha(w)+(I^A(u,v)+\vec{K}^A_{1,2}(u,v)-2O^A(u,v))\\ -(O^A(v,w)+I^A(v,w)+\vec{K}^A_{1,2}(v,w))+\vec{P}^N_3(u,w)\leq 1+\dfrac{1}{n-2}\end{array}\right\} \quad (6)$$

Let us pick $\boldsymbol{x}\in V(\Gamma_0)\setminus\{u,v,w\}$ uniformly at random and let us re-calculate all quantities in the left-hand side of (6) with respect to that distribution. Denoting these re-calculated quantities by $\widetilde{\alpha}(u),\ldots,\widetilde{\vec{P}^N_3}(u,w)$, we claim that

$$\left.\begin{array}{r}\widetilde{\alpha}(u)+\widetilde{\alpha}(v)+\widetilde{\alpha}(w)+(\widetilde{I^A}(u,v)+\widetilde{\vec{K}^A_{1,2}}(u,v)-2\widetilde{O^A}(u,v))\\ -(\widetilde{O^A}(v,w)+\widetilde{I^A}(v,w)+\widetilde{\vec{K}^A_{1,2}}(v,w))+\widetilde{\vec{P}^N_3}(u,w)\leq 1.\end{array}\right\} \quad (7)$$

Indeed, it is easy to see by inspection that every $x\not\in\{u,v,w\}$ contributes at most 1 to the left-hand side (among all P-flags of size 4 based on $\vec{C}_3$-free orgraphs, the only flag that contributes 2 is based on the Twisted Circle[2]).

[2]The reader better familiar with the formalism of flag algebras may notice that we are simply proving the inequality $\sum_{i=1}^3\pi^{P,i}(\alpha)+\pi^{P,[1,2]}(O^A+I^A+\vec{K}^A_{1,2})-\pi^{P,[2,3]}(O^A+I^A+\vec{K}^A_{1,2})+\pi^{P,[1,3]}(\vec{P}^N_3)-3\pi^{P,[1,2]}(O^A)\leq 1$. All the forthcoming argument would have become unnecessary with more systematic application of that formalism.

On the other hand, we obviously have

$$\begin{aligned}
\alpha(u) &= \frac{n-3}{n-1}\widetilde{\alpha}(u) + \frac{1}{n-1};\\
\alpha(v) &= \frac{n-3}{n-1}\widetilde{\alpha}(v) + \frac{1}{n-1};\\
\alpha(w) &= \frac{n-3}{n-1}\widetilde{\alpha}(w);\\
\vec{P}_3^N(u,w) &= \frac{n-3}{n-2}\widetilde{\vec{P}_3^N}(u,w) + \frac{1}{n-2};\\
F(y,z) &= \frac{n-3}{n-2}\widetilde{F}(y,z) \text{ for any other term } F(y,z) \text{ in (6), (7).}
\end{aligned}$$

Multiplying (7) by $\frac{n-3}{n-2}$ and adding $\frac{1}{n-1}$ to both sides, we get

$$\left.\begin{array}{l}
\dfrac{n-3}{n-2}\left(\widetilde{\alpha}(u)+\widetilde{\alpha}(v)+\widetilde{\alpha}(w)\right) + \left(I^A(u,v) + \vec{K}_{1,2}^A(u,v) - 2O^A(u,v)\right)\\
\quad -\left(O^A(v,w) + I^A(v,w) + \vec{K}_{1,2}^A(v,w)\right) + \vec{P}_3^N(u,w) \le 1.
\end{array}\right\} \quad (8)$$

Also, since $\alpha(u)+\alpha(v)+\alpha(w) > 1$ by our assumption, we have

$$\begin{aligned}
\frac{n-3}{n-2}\left(\widetilde{\alpha}(u)+\widetilde{\alpha}(v)+\widetilde{\alpha}(w)\right) &= \frac{n-1}{n-2}\left(\alpha(u)+\alpha(v)+\alpha(w)\right) - \frac{2}{n-2}\\
&\ge \alpha(u)+\alpha(v)+\alpha(w) - \frac{1}{n-2}.
\end{aligned}$$

Substituting this into (8) gives us (6).■

Now the proof of Theorem 3.1 is completed by an easy averaging argument. Since for every vertex v there exists at least one critical edge going out of v, there exists a cycle $C = (v_1, v_2, \ldots, v_\ell)$ consisting of critical edges. After summing up the inequalities (5) along this cycle, the terms $O^A(u,v), \ldots, \vec{K}_{1,2}^A(v,w)$ get canceled and we arrive at

$$\sum_{i\in\mathbb{Z}_\ell} \alpha(v_i) \le \ell/3.$$

Therefore, there exists at least one i with $\alpha(v_i) \le 1/3$. Theorem 3.1 is proved.

That would be interesting to improve our result by removing some (and preferably all) forbidden subgraphs on Figure 1. We have tried it for a while, but all three constraints are very essential in our proof, and removing any one of them immediately creates a new level of difficulties that we have not been able to surpass.

Acknowledgement

I am grateful to Jan Hladký for several useful remarks.

References

[BCW70] M. Behzad, G. Chartrand, and C. E. Wall. On minimal regular digraphs with given girth. *Fundamenta Mathematicae*, 69:227–231, 1970.

[Bon97] J. A. Bondy. Counting subgraphs: A new approach to the Caccetta-Häggkvist conjecture. *Discrete Math.*, 165/166:71–80, 1997.

[CH78] L. Caccetta and R. Häggkvist. On minimal digraphs with given girth. *Congressus Numerantium*, 21:181–187, 1978.

[Cha05] P. Charbit. *Circuits in Graphs and Digraphs via Embeddings.* PhD thesis, L'Université Claude Bernard – Lyon I, 2005.

[CS83] V. Chvátal and E. Szemerédi. Short cycles in directed graphs. *J. Combin. Theory Ser. B*, 35(3):323–327, 1983.

[CS06] M. Chudnovsky and P. Seymour, 2006. Personal communication.

[Ham87] Y. O Hamidoune. A note on minimal directed graphs with given girth. *J. Combin. Theory Ser. B*, 43(3):343–348, 1987.

[HHK07] P. Hamburger, P. Haxell, and A. Kostochka. On directed triangles in digraphs. *Electronic Journal of Combinatorics*, 14(1):Note 19, 2007.

[HKN09] J. Hladký, D. Král', and S. Norine. Counting flags in triangle-free digraphs. Manuscript, available at `http://arxiv.org/abs/0908.2791`, 2009.

[HR87] C. T. Hoáng and B. Reed. A note on short cycles in digraphs. *Discrete Mathematics*, 66(1-2):103–107, 1987.

[Lov10] L. Lovász. Subgraph densities in signed graphons and the local Sidorenko conjecture. Technical Report 1004.3026v1 [math.CO], arXiv, 2010.

[LS06] L. Lovász and B. Szegedy. Limits of dense graph sequences. *Journal of Combinatorial Theory, Series B*, 96(6):933–957, 2006.

[Nis72] T. Nishimura. Short cycles in digraphs. *Discrete Mathematics*, 1988(1-3):295–298, 72.

[Raz07] A. Razborov. Flag algebras. *Journal of Symbolic Logic*, 72(4):1239–1282, 2007.

[Raz10a] A. Razborov. On 3-hypergraphs with forbidden 4-vertex configurations. *SIAM Journal on Discrete Mathematics*, 24(3):946–963, 2010.

[Raz10b] A. Razborov. On the Fon-der-Flaass interpretation of extremal examples for Turan's (3,4)-problem. Technical Report 1008.4707v1 [math.CO], arXiv, 2010.

[Sey06] P. Seymour, 2006. Personal communication.

[She98] J. Shen. Directed triangles in graphs. *Journal of Combinatorial Theory Ser. B*, 74(2):405–407, 1998.

[She00] J. Shen. On the girth of digraphs. *Discrete Mathematics*, 211(1-3):167–181, 2000.

[She02] J. Shen. On the Caccetta-Haӗgkvist conjecture. *Graphs and Combinatorics*, 18(3):645–654, 2002.