Relation categories and graph homomorphisms

T.H.Marshall *

September 25, 2002

Abstract

Let $\mathcal{C}$ be a class of graphs. A graph H is $\mathcal{C}$*-universal* if it admits a homomorphism from every graph in $\mathcal{C}$. For certain classes of graphs–which we designate *finitely constructible* classes–we derive an algorithm to test a given graph H for $\mathcal{C}$-universality. This algorithm also works in the category of (n, m)-mixed colored graphs, which include (for different values of m and n) oriented and edge coloured graphs.

The algorithm uses a finite category related to $\mathcal{C}$, whose objects are sets of functions and whose arrows are relations with composition.

We study how this algorithm works for two sorts of finitely constructible classes of oriented graphs, grids of fixed width and certain subclasses of the planar graphs.

1 Introduction

A mixed graph is a triple $G = (V, E, A)$ where V is a set and E and A are respectively sets of unordered pairs and ordered pairs from V. We refer to the V, E and A as, respectively, the vertices edges and arcs of G. We assume that the underlying graph of G is always simple, that is, there are no loops and, if (u, v) is an arc, then (v, u) is not and $\{u, v\}$ is not an edge (digraphs with this property are also often referred to as an *oriented* graphs).

*This work is supported by a DIMATIA postdoctoral fellowship. Written at Centrum Dimatia 11800, Praha 1, Czech Republic

An (n,m)-mixed colored graph (for brevity an (n,m)-graph) is a mixed graph with arcs coloured from the set of colours $\{1, 2, \dots, n\}$ and edges coloured from the set $\{1', 2', \dots, m'\}$. We include the possibility that n or m is zero, in which case (n,m)-graph are (edge-coloured) undirected graphs and digraphs respectively. Further, $(0,1)$-graphs $(1,0)$-graphs may be regarded as ordinary uncoloured graphs and digraphs respectively, indeed our main motivation in considering mixed graphs is that it allows these two cases to be dealt with at once. Most literature about homomorphisms considers only uncoloured graphs and digraphs. Homomorphisms of coloured undirected graphs were introduced by Alon and Marshall in [1] and of mixed graphs by Nešetřil and Raspaud in [7].

A *homomorphism* between two (n,m)-graphs G_1 and G_2 is defined to be a mapping $\phi : V_1 \longmapsto V_2$ with the property that if $\{u, v\}$ (resp. (u, v)) is an edge (resp. arc) of a given colour then $\{\phi(u), \phi(v)\}$ (resp. $(\phi(u), \phi(v))$) is an edge (resp. arc) of the same colour.

Let $\mathcal{C}$ be a class of (uncoloured undirected simple) graphs. An $(n,m,\mathcal{C})$-*universal* graph is an (n,m)-graph G into which every (n,m)-graph whose underlying graph lies in $\mathcal{C}$ maps homomorphically (there is no requirement that G itself have underlying graph in $\mathcal{C}$). We may abbreviate to "$\mathcal{C}$-universal" or just "universal" when the parameters are clear from the context.

Given some class $\mathcal{C}$, the most fundamental question about $\mathcal{C}$-universal graphs is: do they exist? For unoriented uncoloured graphs this is clearly the case exactly when the chromatic numbers of graphs in $\mathcal{C}$ are bounded. In the other cases, Nešetřil and Raspaud [7] have shown that, if there is a bound on the acyclic chromatic numbers of the graphs in $\mathcal{C}$, then $(n,m,\mathcal{C})$-universal graphs exist for all (n,m). Conversely Kostochka, Sopena and Zhu [5] show that the existence of a $(1,0,\mathcal{C})$-universal graph bounds the acyclic chromatic numbers of the graphs in $\mathcal{C}$.

Given that $\mathcal{C}$-universal graphs do exist, two further problems follow naturally; what is the smallest possible order of a $\mathcal{C}$-universal graph and whether or not there is an algorithm to test a given oriented graph for $\mathcal{C}$-universality. The first question has been the subject of much recent work, mostly for oriented graphs (see for example [1], [3], [6], [7], [8] and [9]). The problem of testing a *given* graph for $\mathcal{C}$-universality has been much studied only in the case where $\mathcal{C}$ is a singleton, that is the problem, given oriented graphs G and H, of deciding whether or not there a homomorphism from G to H. Of course this question is trivially decidable and the interesting problem is to find the complexity of the decision procedure. In the undirected case this problem has been completely solved by Hell and Nešetřil [4]. In the directed case only partial answers are known ([9] and references therein).

If $\mathcal{C}$ is infinite the initial question is no longer one of complexity but whether or not there exists an algorithm to test $\mathcal{C}$-universality at all. In general this appears to be a very difficult problem. If $\mathcal{P}$ is the class of planar graphs then, even in the undirected case, an algorithmic test for $\mathcal{P}$-universality is equivalent to a proof of the four colour theorem (By the 4CT, a graph is $\mathcal{P}$-universal precisely when it contains an imbedded K_4.). We study this relationship more closely in Section 5.

In the directed case there is no obvious algorithm even for such a simple subclass of $\mathcal{P}$ as the class of grids. The difficulties appear to be related to the fact that, for all these classes the treewidth of the graphs in $\mathcal{C}$ is unbounded.

For classes of bounded treewidth we have some more positive results. For example, if $\mathcal{C}$ is the class of trees then it is an easy exercise to see that U is $(1, 0, \mathcal{C})$-universal if and only if it contains a directed cycle. In this paper we describe an algorithm (Theorem 3.2) to test for $(n, m, \mathcal{C})$-universality when, $\mathcal{C}$ is *finitely constructible*. Roughly speaking this means that the graphs in $\mathcal{C}$ are those which can be built up by progressively "glueing" together graphs from a given finite collection of according to specified rules.

Examples of such classes include the grids of fixed width k and the classes $\mathcal{P}_k$ where $\{\mathcal{P}_k\}$ $(k \geq 3)$ is a certain increasing sequence of subclasses of the planar graphs $\mathcal{P}$ whose union is $\mathcal{P}$. These are defined precisely in the next section. The property of $\mathcal{P}_n$-universality is a useful approximation of $\mathcal{P}$-universality. In [6] we prove that a certain graph is not $\mathcal{P}$-universal, by showing that it is not $\mathcal{P}_5$-universal.

The algorithm is most easily formulated in terms of categories, although we use the language rather than the theory. The necessary terminology is given in the next section. In Section 3 we derive the algorithm itself and for the remainder of the paper consider examples and special cases.

The following notation and conventions will be used in the paper. If $\phi : G_1 \to G_2$ is a homomorphism and G is a subgraph of G_1 the *image* $\phi(G)$ of G under ϕ is defined to be the subgraph of G_2 whose vertices are the images of those of G with two vertices $\phi(u)$ and $\phi(v)$ being joined if u and v are. Hence $\phi(G)$ is not generally an induced subgraph.

We let $h(G, H)$ denote the set of homomorphisms from G to H, Γ the class of plane grids and Γ_k the class of plane grids of width k.

2 Relation categories

A *relation category* is a category whose objects are sets and for which each $\mathrm{Hom}(X, Y)$ is a set of relations from X to Y.

Let $\mathcal{S}$ be a collection of sets and $\mathcal{R}$ a collection $\{R(X, Y)|X, Y \in \mathcal{S}\}$, where each $R(X, Y)$ is a set of relations between X and Y. The relation category $C(\mathcal{S}, \mathcal{R})$ *generated* by $(\mathcal{S}, \mathcal{R})$ is the category whose objects are the sets in $\mathcal{S}$, for which each $\mathrm{Hom}(X, Y)$ is minimal, subject to containing $R(X, Y)$. Explicitly the $\mathrm{Hom}(X, Y)$ are the sets generated inductively as follows

1. The identity relation I_X is in $\mathrm{Hom}(X, X)$ $\forall X \in \mathcal{S}$
2. $R(X, Y) \subseteq \mathrm{Hom}(X, Y)$ $\forall X, Y \in \mathcal{S}$
3. If X, Y and A are in $\mathcal{S}$, $R \in \mathrm{Hom}(X, A)$, $S \in \mathrm{Hom}(A, Y)$ then $S \circ R \in \mathrm{Hom}(X, Y)$
4. $\mathrm{Hom}(X, Y)$ is as small as possible subject to the above conditions.

An important special case of the above is when $\mathcal{S}$ contains only one set, in which case $\mathrm{Hom}(X, X)$ is just the semigroup of relations generated by $R(X, X)$. In fact we can always reduce to this case just by taking the disjoint union of the vertex sets and the union of all the $R(X, Y)$ though in practice it is usually easier not to.

We will consider only cases where the sets in $\mathcal{S}$ are finite, though $\mathcal{S}$ itself may be infinite.

A number of problems concerning graph homomorphisms reduce to the following question: given a relation category generated by some given $(\mathcal{S}, \mathcal{R})$, does any $\mathrm{Hom}(X, Y)$ contain the empty relation? A more specific question is this; given a subset A of some X, is it true that $R[A]$ is nonempty for every R in every $H(X, Y)$. If so the set A is *persistent*. Note, in particular, that a persistent set must be nonempty.

These questions can also be formulated in terms of homomorphisms of directed coloured paths. Given a relation category $C(\mathcal{S}, \mathcal{R})$, where $\mathcal{S}$ contains n sets, which we may take to be mutually disjoint, form an edge coloured n-partite digraph whose vertex set is the union of the sets in $\mathcal{S}$. For each relation τ in each $R(X, Y)$, let there be an arc coloured τ from v to w if $(v, w) \in \tau$. Of course the n-partition of this digraph corresponds to the sets in $\mathcal{S}$.

A coloured directed path P of length k is *compatible* with $S \in \mathcal{S}$, if there is a sequence of sets $S = S_1, S_2, \ldots, S_{k+1}$ in $\mathcal{S}$ and relations $\tau_i \in R(S_i, S_{i+1})$, for which the edges of P are coloured succesively $\tau_1, \tau_2 \ldots \tau_k$.

Now a subset $T \subseteq S \in \mathcal{S}$ is persistent if and only if for any coloured directed path P, compatible with S, this digraph admits at least one (colour and direction preserving) homomophism from P with the initial point mapped into T.

When $|\mathcal{S}| = 1$ every path is compatible with the set in $\mathcal{S}$.

For finite $\mathcal{S}$ and given $\mathcal{R}$, the question of whether or not a given set is persistent in $C(\mathcal{S}, \mathcal{R})$ is of course computably solvable, but presumably NP-hard. When $\mathcal{S}$ is infinite this question may be very difficult, but sometimes can be "approximated" by using a finite subcategory of $C(\mathcal{S}, \mathcal{R})$. In Section 5 we show that the four colour theorem is equivalent to the statement that a certain set is persistent in an infinite relation category.

We define functions in the categorical sense, that is with an implicit dependence on a structured domain and a structured codomain. For example the set of functions from the vertex set of a graph G to that of a graph H, is really a triple comprising the set of functions together with G, H.

Since all graphs we consider are finite, we assume the vertex sets are all subsets of some fixed countable set such as the integers. Thus the categories we consider are all genuine sets.

3 Construction of Graphs

Definition An (n,m)-*pair* (briefly a *pair*) is an ordered pair $(G,A)_\phi$ of (n,m)-graphs where ϕ is an injective homomorphism of A into G. An (n,m)-*triple* (briefly a *triple*) is an ordered triple $(G,A,B)_{\psi,\psi'}$ of (n,m)-graphs where ψ and ψ' are injective homomorphisms of A and B respectively into G. When writing pairs and triples we will often omit subscripts.

Next we define a "glueing" operation between a pair and a triple. If $(G,A)_\phi$ and $(G_1,A,B)_{\psi,\psi'}$ are respectively a pair and a triple, then let $\tilde{G}$ be the quotient graph obtained from the disjoint union of G and G_1 by identifying $\phi(a)$ with $\psi(a)$ for each $a \in A$. If this graph has multiple edges then $(G,A)\sharp(G_1,A,B)$ is not defined. Otherwise it is the pair $(\tilde{G},B)_\chi$, where $\chi = \iota \circ \psi'$, where ι is the canonical injection of G_1 into $\tilde{G}$.

The "sum" $p\sharp t$ of a pair and a triple is thus defined provided that glueing in the triple t does not join any two vertices which are already joined in p.

When we consider a (non-induced) subgraph H of a graph G it will sometimes be necessary to "remember" which pairs of H were joined in G by an edge or arc, without being concerned about its colour or orientation. For this purpose we define *auxiliary edges*, which, for any given graph, will be an extra set of ordered pairs of vertices (which we refer to as being "joined" by auxiliary edges). We will require that a homomorphism of a graph with auxiliary edges does not map two vertices joined by an auxiliary edge to the same image. (In contrast to ordinary edges and arcs, there is no requirement that these images be joined by an auxiliary, or any, edge).

Definition:- For any pair $p = (G, A)_\phi$, $c(p)$ is the subgraph $\phi(A)$ of G with auxiliary edges joining pairs of vertices of $\phi(A)$ which are joined in G.

Definition:- Let U be obtained from A by adding (possibly) some auxiliary edges and let $t = (G, A, B)_{\phi,\psi}$ be a triple. We say that U is *t-compatible* if whenever vertices u and v of U are joined by an auxiliary edge in U, their images under ϕ are not joined in G.

Definition:- Let U, be t-compatible. We define a function $R_{t,U}$ on $\{U\}$ by letting $R_{t,U}(U)$ be the graph B with auxiliary edges joining vertices u and v whenever $\psi(u)$ and $\psi(v)$ are joined in G or are the images under ϕ of vertices joined by an auxiliary edge in U.

The following lemma is now an easy consequence of definitions.

Lemma 3.1 *For $p = (G, A)$ and $t = (G_1, A, B)$ be respectively a pair and a triple , $p\sharp t$ is defined if and only if $c(p)$ is t-compatible. In this case $c(p\sharp t) = R_{t,c(p)}(c(p))$.*

Definition:- Let $\mathcal{D}$ and $\mathcal{T}$ be respectively classes of pairs and of triples. Let $B(\mathcal{D}, \mathcal{T})$, the class of pairs *built* from $\mathcal{D}$ and $\mathcal{T}$, be defined inductively as follows:

- If $p \in \mathcal{D}$ then $p \in B(\mathcal{D}, \mathcal{T})$.
- If $p \in B(\mathcal{D}, \mathcal{T})$ and $t \in \mathcal{T}$ then $p\sharp t \in B(\mathcal{D}, \mathcal{T})$ whenever it is defined.

The class of graphs $G(\mathcal{D},\mathcal{T})$ is defined by

$$G(\mathcal{D},\mathcal{T}) = \{G : \exists G_1, G \text{ is a subgraph of } G_1, \quad (G_1, A) \in B(\mathcal{D},\mathcal{T})\}$$

In other words graphs in $G(\mathcal{D},\mathcal{T})$ are obtained from pairs in $B(\mathcal{D},\mathcal{T})$ by "forgetting" the homomorphic imbedding and then (possibly) taking a subgraph.

Examples

1. Let $\mathcal{D}$ contain the pairs of the form $(L, L)_I$ where L is a path of length k and I is the identity homomorphism on L and let $\mathcal{T}$ contain the triples of the form $(H, L_1, L_2)_{\phi_1,\phi_2}$ where H is a ladder of length k, L_1 and L_2 are paths of length k and ϕ_1, ϕ_2 are injections of L_1 and L_2 respectively into opposite sides of H (Figure 1).

   In this case $B(\mathcal{D},\mathcal{T})$ comprises pairs of the form (G, L) where $G \in \Gamma_k$ and L is a path of length k imbedded into one of its ends and $G(\mathcal{D},\mathcal{T}) = \Gamma_k$.

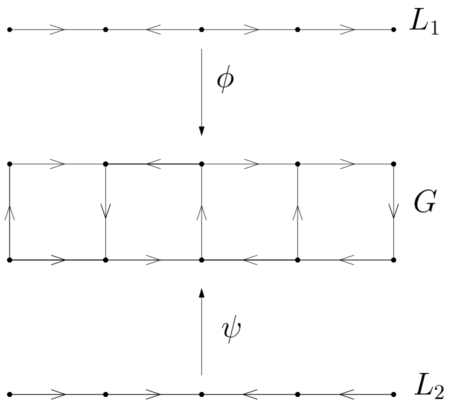

Fig. 1

2. Let $\mathcal{D} = \{(\Delta, \Delta)_I\}$ where Δ is a triangle and I is the identity homomorphism on Δ. Let $\mathcal{T}_k$ $(3 \le n \le \infty)$ contain the triples of the form $(G, C_1, C_2)_{\phi_1,\phi_2}$, where C_1 is a cycle of length at most k, G consists of the image $\phi_1(C_1)$ along with either a path of length 2 joining two

adjacent vertices of $\phi_1(C_1)$ or an edge joining two vertices at distance two in $\phi(C_1)$. Apart from $\phi_1(C_1)$ there are thus two cycles in G. We require that C_2 is a cycle mapped by ϕ_2 into the longer of these (or either if both are equal).

For $3 \leq k \leq \infty$, let $\mathcal{Q}_k = B(\mathcal{D}, \mathcal{T}_k)$ and $\mathcal{P}_k = G(\mathcal{D}, \mathcal{T}_k)$.

Since the graphs in $\mathcal{P}_\infty$ are subgraphs of those obtained from a triangle by repeatedly subdividing a face, these graphs are all planar. The converse is also true.

Theorem 3.1 *If G is a plane graph with at most one nontriangular face and C is this nontriangular face or, if none exists, any face of G then $(G, C) \in Q_\infty$.*

Proof:- Let G be a plane graph with at most one nontriangular face. We prove the theorem by induction on the number of edges in G. If G is a triangle then we are done. If not, then let C be the largest face of G (or any face if they are all triangles) and let C_1 be adjacent to C. If C and C_1 have one common edge e, remove it to obtain a plane graph G_1, which has one nontriangular face C_2 obtained by removing e. By the induction hypothesis $(G_1, C_2) \in Q$ and we recover (G, f) by glueing in a cycle with chord.

If C and C_1 have two common edges a similar argument shows that there is a $(G_3, C_3) \in Q$ from which (G, f) is obtained by glueing in a cycle with path of length two. $\square$

Clearly

$$\mathcal{P}_3 \subseteq \mathcal{P}_4 \subseteq \mathcal{P}_5 \subseteq \ldots \mathcal{P}_\infty,$$

and the proof of the previous theorem shows that each planar graph is in some $\mathcal{P}_k$ $(k < \infty)$. The classes $\mathcal{P}_k$ can thus be considered as a sequence of bounded treewidth approximations to $\mathcal{P}$, each graph in $\mathcal{P}_k$ having treewidth at most $k - 1$.

Definition:- If (G, A) is a pair then $\pi_1(G, A) = G$ and $\pi_2(G, A) = A$. Similarly for a triple (G_1, A, B), $\pi_1(G_1, A, B) = G_1$, $\pi_2(G_1, A, B) = A$ and $\pi_3(G_1, A, B) = B$.

Let $\mathcal{D}$ and $\mathcal{T}$ be given. Let $Y_1 = \pi_2(\mathcal{D}) \cup \pi_3(\mathcal{T})$. Let Y be the mixed graphs obtained from graphs in Y_1 by adding zero or more auxiliary edges.

If $y \in Y$ let $\pi(y)$ be the graph in Y_1 obtained by omitting auxiliary edges. We now derive two categories and a functor between them from $\mathcal{D}$ and $\mathcal{T}$.

Definition:-
$\mathcal{R} = \{\{R_{t,U} |\ t \in \mathcal{T}, U$ is t-compatible and $R_{t,U}(U) = V\}|\ U, V \in Y\}$ and $\mathcal{M} = \mathcal{M}(\mathcal{D}, \mathcal{T}) = C(Y, \mathcal{R})$.

Let H be a chosen and fixed (n, m)-graph. The following definitions depend implicitly on H. For $K \in Y$ let $K' = h(K, H)$.

If $U, V \in Y$, $A = \pi(U)$, $B = \pi(V)$, $t = (G, A, B)_{\phi,\psi}$ and U is t-compatible, let $R'_{t,U}$ be the relation between U' and V' given by $(\alpha, \beta) \in R'_{t,U}$ if $\beta = \gamma \circ \psi$ and γ is an extension to G of the homomorphism $\alpha \circ \phi^{-1}$ defined on $\phi(A)$. Roughly speaking β is the restriction to V of an extension of α.

Definition:-
$\mathcal{R}' = \{\{R'_{t,U} |\ t \in \mathcal{T}, U$ is t-compatible and $R_{t,U}(U) = V\}|\ U, V \in Y\}$ and $\mathcal{N} = \mathcal{N}(\mathcal{D}, \mathcal{T}) = C(Y', \mathcal{R}')$, where $Y' = \{X' \mid X \in Y\}$.

Let $F = F(\mathcal{D}, \mathcal{T})$ be the functor from $\mathcal{M}$ to $\mathcal{N}$ defined by mapping $U \in Y$ to U' and the morphism $R_{t,U}$ to $R'_{t,U}$ and extending the map to other morphisms in such a way as to make F a functor.

Definition:- Let $p = (G, A)_\phi$ be a pair, $h(p)$ is the set of homomorphisms from $c(p)$ to H which are restrictions of homomorphisms from G to H.

Theorem 3.2 *Let $\mathcal{I} = \{h(p) \mid p \in \mathcal{P}\}$. Then H is $B(\mathcal{D}, \mathcal{T})$-universal if and only if each set in $\mathcal{I}$ is persistent in $\mathcal{N}$.*

The theorem follows by an easy induction from the following lemma

Lemma 3.2 *If $p \in B(\mathcal{D}, \mathcal{T})$, $t \in \mathcal{T}$ and $c(p)$ is t-compatible, then*

$$h(p\sharp t) = R'_{t,c(p)}[h(p)]$$

Proof:-. Let $p = (G, A)_\phi$, $t = (X, A, A_1)_{\phi_1,\psi}$ and $p\sharp t = (G_1, A_1)_\psi$.

$$\begin{aligned}
\beta &\in R'_{t,c(p)}[\{\alpha \circ \phi : \alpha \in h(G, H)\}] \\
&\Leftrightarrow \exists \alpha \in h(G, H), \beta = \gamma \circ \psi : \gamma \text{ an extension to } X \text{ of } \alpha|_A \\
&\Leftrightarrow \exists \alpha \in h(G, H), \beta = \gamma \circ \psi : \gamma \text{ an extension to } G_1 \text{ of } \alpha \\
&\Leftrightarrow \gamma \in h(G_1, H) \ \Box
\end{aligned}$$

Definition:- A class $\mathcal{C}$ of (n,m)-graphs is *finitely constructible* If $\mathcal{C} = \pi_1(B(\mathcal{D},\mathcal{T})) = \pi_1(B(\mathcal{D},\mathcal{T})) := \{\pi_1(p) |\ p \in B(\mathcal{D},\mathcal{T})\}$ for some finite $\mathcal{D}$ and $\mathcal{T}$. Since, in this case the category $\mathcal{N}(\mathcal{D},\mathcal{T})$ of Theorem 3.2 is also finite, we have

Corollary 3.1 *If $\mathcal{C}$ is a finitely constructible class of graphs then Theorem 3.2 gives an effective algorithm to determine whether or not a given (n,m)-graph H is $\mathcal{C}$-universal.*

Since the classes Γ_k and $\mathcal{P}_k$ given in examples 1 and 2 above are finitely constructible for each k, we have

Corollary 3.2 *Theorem 3.2 gives effective algorithms to test a given (n,m)-graph H for $\mathcal{P}_k$-universality and for Γ_k-universality.*

Observe that the graphs in any finitely constructible class $G(\mathcal{D},\mathcal{T})$ have treewidth at most $\max\{\pi_1(t)|t \in \mathcal{T}\} - 1$.

4 An example: testing $\mathcal{P}_5$-universality

In practice the algorithm of Theorem 3.2 requires only the component of $\mathcal{M}$ containing the objects in $\pi_2(\mathcal{D})$ and the image of this subcategory under F rather than $\mathcal{M}$ and $\mathcal{N}$ themselves. In the case of $B(\mathcal{D},\mathcal{T}) = \mathcal{Q}_k$, $\pi_2(\mathcal{D})$ contains only a triangular cycle. The component of $\mathcal{M}$ containing this is shown in the diagram below for $k = 5$, $(n,m) = (0,1)$, that is the oriented uncoloured case. (To keep thing manageable, we have ignored labelling of arrows and orientations of graphs. The objects in the category are really all the orientations of those pictured). Auxiliary edges are depicted by broken lines. (Figure 2)

As a detailed example of one morphism in this category, let U and $t = (G, A, B)_{\phi,\psi}$ be as shown. We let $a' = \phi(a) = \psi(a)$ etc. (Figure 3)

Because of the enormous number of homomorphisms involved, the algorithm of Theorem 3.2 seems computationally feasible only for testing highly symmetric H. (In this case we can use equivalence classes of homomorphisms under the automorphism group of H rather than the homomorphisms themselves). Examples of such graphs are the oriented tournaments QR_p and

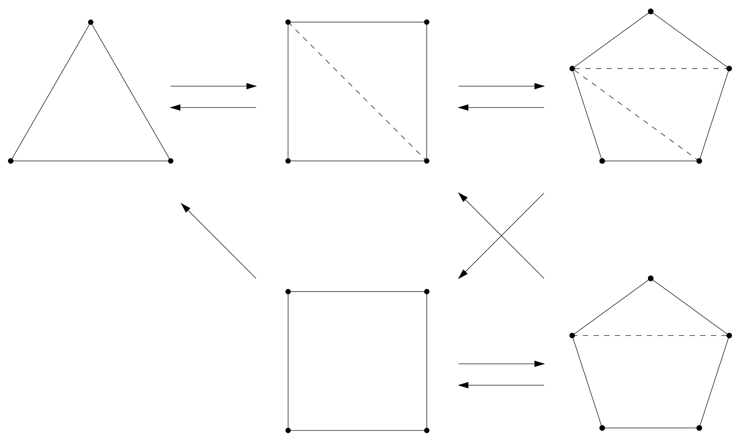

Fig. 2

the Tromp graphs T_{2p+2} defined in [2], (where in both cases p is a prime $p \equiv 3 \bmod 4$ and the subscript indicates the order of the graph).

Sopena [9] has proved that the Tromp graph T_{16} on 16 vertices is $\mathcal{T}_3$-universal, where $\mathcal{T}_3$ represents the class of 3-trees, and hence that it is $\mathcal{P}_4$-universal. On the other hand T_{16} is not $\mathcal{P}_5$-universal nor is any oriented graph of order at most 16 $\mathcal{P}$-universal [6].

5 The four colour theorem

If we are concerned only with undirected graphs all of the foregoing can be considerably simplified. In this case we can allow $(G, A)\sharp(G, A, B)$ to be defined whether or not this creates multiple edges, since the problem of two vertices being joined by edges or arcs of *different* colour or orientation does not arise. We thus dispense with auxiliary edges and set $c(p) = A$ for $p = (G, A)$, $R_{t,A}(A) = B$ when $t = (G, A, B)$ and $Y = Y_1$.

We now consider the specific problem of testing undirected graphs for $\mathcal{P}$-universality. Here of course the four colour theorem gives us prior knowledge of the answer, but the undirected version of theorem 3.2 gives a restatement of the four colour theorem in terms of persistent sets which raises some interesting new problems.

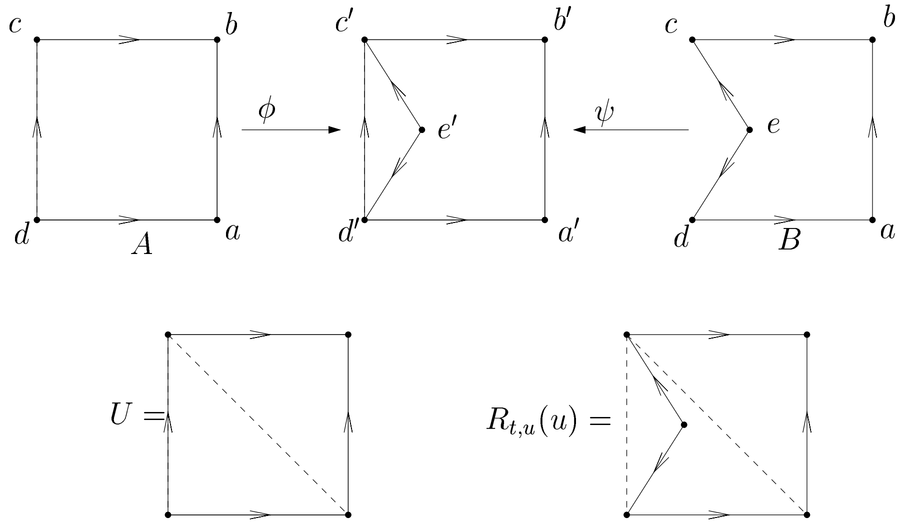

Fig. 3

When testing undirected graphs for $\mathcal{P}$ (resp. $\mathcal{P}_n$) universality, the objects of $\mathcal{M}$ are just the cycles (resp. cycles of length at most n). If C is a cycle then $R_{t,C}(C)$ is obtained from C either by replacing a path of length two by an edge or *vice versa*. We refer to the application of these two kinds of morphism as cutting a corner and adding a corner respectively and collectively as subdivisions.

By Theorem 3.2, the following is a version of the four colour theorem.

Theorem 5.1 *If $\mathcal{D}$ and $\mathcal{T}_\infty$ are defined as in Example 2 and T is a triangle, then $h(T, K_4)$ is persistent in $\mathcal{N}(\mathcal{D}, \mathcal{T}_\infty)$.*

If C is a cycle (with vertices, say, $1, 2, 3, \ldots n$), it is convenient to represent a homomorphism ϕ in $h(C, K_4)$ by listing the values of $\phi(1), \phi(2), \ldots$ in a row. A set of homomorphisms can then be represented by listing these rows in a matrix, (with the order of the rows being immaterial). Thus if C has order n then every subset of $h(C, K_4)$ gives rise to a matrix with n columns such that no two adjacent entries in a row are the same (here and subsequently the first and last columns of a matrix are understood to be "adjacent"). Conversely every such matrix determines a subset of $h(C, K_4)$. We refer to a matrix as "persistent" if its corresponding set of homomorphisms is.

Corresponding the operations of cutting and adding corners to the cycle we have operations of deletion and insertion, respectively, defined on matrices. The deletion operation D_i is defined (on matrices with at least four columns) by deleting the ith column of the matrix and then discarding repeated rows and rows in the new matrix where two adjacent elements are the same. The insertion operations I_i are defined by replacing each row of the matrix $(a_1, a_2, \ldots a_n)$ by the two rows that can be obtained by putting a new entry between a_i and a_{i+1}, which is different to both of these. A matrix is thus persistent if and only if it cannot be annihilated by any sequence of deletions and insertions.

Call $\phi, \psi \in h(C, K_4)$ *equivalent* if $\phi = \rho\psi$ for some permutation ρ of K_4. There is an obvious corresponding equivalence between rows of matrices. Clearly the persistence or otherwise of a matrix is determined only by the equivalence classes of the rows so we may tacitly use equivalence classes of homomorphisms rather than the homomorphisms themselves. (Thus after a deletion operation one removes all but one representative of each equivalence class). For the cycle of length 3 there is a single equivalence class in $h(C, K_4)$. Thus the four colour theorem is equivalent to the statement that

$$M = \begin{pmatrix} 1 & 2 & 3 \end{pmatrix}$$

is persistent. We now derive some more persistent matrices by performing deletions and insertions on this one. For example

$$I_3(M) = \begin{pmatrix} 1 & 2 & 3 & 4 \\ 1 & 2 & 3 & 2 \end{pmatrix}$$

Other insertions just give (modulo equivalence) cyclic permutations of the columns of this one so we ignore them. Any deletion performed on $I_3(M)$ gives M (again modulo equivalence). Note that neither of the rows of $I_3(M)$ is persistent, D_3 annihilates the second row while I_4 D_4 and D_3 applied to the top row give in turn

$$\begin{pmatrix} 1 & 2 & 3 & 4 & 2 \\ 1 & 2 & 3 & 4 & 3 \end{pmatrix}, \begin{pmatrix} 1 & 2 & 3 & 2 \end{pmatrix}$$

and the empty matrix. We refer to a matrix as *minimal persistent* if it is persistent but no submatrix obtained by removing rows is. Thus both M

and $I_3(M)$ are minimal persistent. Insertions I_1-I_4 applied to $I_3(M)$ give in turn,

$$I_1 = \begin{pmatrix} 1 & 3 & 2 & 3 & 4 \\ 1 & 4 & 2 & 3 & 4 \\ 1 & 3 & 2 & 3 & 2 \\ 1 & 4 & 2 & 3 & 2 \end{pmatrix} \quad I_2 = \begin{pmatrix} 1 & 2 & 1 & 3 & 4 \\ 1 & 2 & 4 & 3 & 4 \\ 1 & 2 & 1 & 3 & 2 \\ 1 & 2 & 4 & 3 & 2 \end{pmatrix}$$

$$I_3 = \begin{pmatrix} 1 & 2 & 3 & 1 & 4 \\ 1 & 2 & 3 & 2 & 4 \\ 1 & 2 & 3 & 1 & 2 \\ 1 & 2 & 3 & 4 & 2 \end{pmatrix} \quad I_4 = \begin{pmatrix} 1 & 2 & 3 & 4 & 2 \\ 1 & 2 & 3 & 4 & 3 \\ 1 & 2 & 3 & 2 & 3 \\ 1 & 2 & 3 & 2 & 4 \end{pmatrix}$$

One can verify that each of these matrices is also minimal. These computations suggest some natural questions.

- What is the minimum possible number of rows in a persistent matrix with k columns?
- Can every minimal persistent matrix be obtained by applying a sequence of insertions and deletions to $M = \begin{pmatrix} 1 & 2 & 3 \end{pmatrix}$?

The answer to the first question is at most 2^{k-3}, since every matrix obtained by applying $k-3$ successive insertions on M has k columns and 2^{k-3} rows. An affirmative answer to the second question would show that equality holds.

A simple criterion to determine whether or not a matrix is persistent would give a proof of the four colour theorem, so is presumably going to be difficult to find. It is not hard however to find some simple *necessary* conditions for persistency. These allow us to derive some results asserting the existence of a four colouring of a planar graph which satisfy some extra conditions.

Proposition 5.1 *A matrix is persistent only if at least one of its rows has at most three distinct entries.*

Proof:- Suppose every row of a matrix M includes each of the numbers $1, 2, 3, 4$. Insert a new column somewhere in M and then succesively delete columns starting from the left of the new one. Since every row contains at least one other entry the same as in the new column, after some sequence of deletions there must be two adjacent entries the same in any given row. Thus every row is eventually eliminated and the matrix is not persistent. □

This result is equivalent to the observation that, for every planar graph P and C a cycle bounding a face of P, there exists a four colouring of P in which the vertices of C take at most three colours (which is easily proved directly by joining a new vertex to all the vertices of C).

Here is a slightly more elaborate necessary condition for persistency. Let $\mathcal{B}$ be a partition of $\{1, 2, \ldots n\}$ into k "blocks" $\{B_i | 1 \leq i \leq k\}$, where each B_i is of the form $\{a_i, a_i + 1, a_i + 2, \ldots a_{i+1} - 1\}$ (where addition is modulo n). Define a function $\pi_{\mathcal{B}}$ from rows of length n with entries in $\{1, 2, 3, 4\}$ to a set of rows of length k by

$$\pi_{\mathcal{B}}([m_1, m_2, \ldots m_n]) = \{[n_1, n_2, \ldots n_k] |\, n_i \in \{1, 2, 3, 4\} - \{m_j |\, j \in B_i\}\}$$

Define $\pi_{\mathcal{B}}$ on a matrix by applying it rowwise.

The following is proved in essentially the same way as Proposition 5.1.

Proposition 5.2 *Let Let $\mathcal{B}$ be a partition of $\{1, 2, \ldots n\}$ into at least 3 blocks as above and let M be a matrix with n columns. The matrix M is persistent only if $\pi_{\mathcal{B}}(M)$ is.*

Proof (sketch):- It suffices to reduce M by a series of insertions and deletions to a matrix whose rows are all in $\pi_{\mathcal{B}}(M)$. This can be done by, for each i, inserting new columns between $a_{i-1} - 1$ and a_i and then deleting all the original columns, starting from the immediate left of each new column and working from right to left. □

6 Oriented chromatic numbers of grids

The *oriented chromatic number* of a oriented graph G is defined to be the minimum order of a homomorphic image of G, and of an unoriented graph H as the maximum oriented chromatic number of its orientations. The following question was raised by André Raspaud (personal communication): What is the largest possible oriented chromatic number of an oriented grid? This number is known to be at least 7 and at most 11.

Let QR_7 be the tournament with vertex set Z_7 and an arc from u to v if $v - u \in \{1, 2, 4\}$–the nonzero quadratic residues in Z_7. It is easy to check that this is a well defined tournament and it is highly symmetrical, the automorphisms being of the form $v \to av + b$, where $a \in \{1, 2, 4\}$.

If QR_7 is $(1, 0, \Gamma)$-universal then of course every grid has oriented chromatic number at most 7 and the original problem is settled.

This, apparently difficult, question can be approximated by substituting the class Γ_k for Γ, that is we consider the question of whether or not, for a given k, QR_7 is Γ_k-universal. Using Theorem 3.2, we formulate this as a question about relation categories and prove that it is true at least for $n = 2$.

Recall that $\Gamma_k = G(\mathcal{D}, \mathcal{T})$ where $\mathcal{D}$ and $\mathcal{T}$ are as defined in Example 1 (the dependence on k is implicit). The application of Theorem 3.2 is fairly simple in this case. Auxiliary edges never appear and the objects of $\mathcal{M}$ are just the 2^n oriented paths of length n. Between any two of these objects there are 2^{n+1} relations of the form $R_{t,L}$ (corresponding to the 2^{n+1} possible orientations of the rungs in $\pi_1(L)$). The set $\mathcal{I}$ of Theorem 3.2, in this case comprises all the sets in $\mathcal{N}$ so we have

Proposition 6.1 *Let $\mathcal{D}$ and $\mathcal{T}$ be as defined in Example 1. A graph H is Γ_k-universal if and only if, for each oriented path L of length k, each $h(L, H)$ is persistent in $\mathcal{N}(\mathcal{D}, \mathcal{T})$.*

For fixed H this test for Γ_k-universality is exponential time in n but it should be feasible using a computer to test the Γ_k-universality of QR_7 for small values of n with a view to finding either either a counterexample to the assertion that QR_7 is Γ-universal or some evidence to believe that it is true (and perhaps some raw material for a proof). For double ladders ($n = 2$) the assertion is easily proved with virtually no computation required. In fact we have the stronger statement

Proposition 6.2 *Let $\mathcal{T}$ be as defined in Example 1 for $n = 2$. For any oriented path L of length 2, $\phi \in h(L, QR_7)$ and triple $t = (G, L, L_1) \in \mathcal{T}$,*

$$R'_{t,L}(\phi) \neq \emptyset$$

hence QR_7 is Γ_2-universal

The proof requires a very simple lemma.

Lemma 6.1 *Let P be any oriented path of length two with endpoints a and b, and let u and v be distinct vertices in QR_7, then there is a homomorphism from P to QR_7 which maps a to u and b to v.*

Proof:- By symmetry we may assume $u = 0$ and $v = 1$ or *vice versa.* It is now easy to verify the lemma. □

Proof of proposition 6.2 :- Let G be a grid with vertex set $\{0, 1, 2\} \times \{0, 1\}$, where (i, j) is adjacent to (i', j') exactly when $|i - i'| + |j - j'| = 1$.

We must show that, any homomorphism ϕ of the path induced by the vertices $\{(0, i) | 0 \leq i \leq 2\}$ to QR_7 extends to G. We extend ϕ to the points $(0, 1)$, $(1, 1)$ and $(2, 1)$ in turn. First there are three choices of image for $(0, 1)$, which extend ϕ to the subgraph induced by the vertices of P and $(0, 1)$. At least two of these points are distinct from $\phi(1, 0)$ so, by the preceding lemma, there are least two choices of image for $(1, 1)$, which extend ϕ to include this vertex. The same argument applies again to show that there is at least one choice of image for $(2, 1)$ which extends ϕ to the whole grid. □

References

[1] N. Alon and T. H. Marshall, "Homomorphisms of edge-colored graphs and Coxeter groups", *J. Algebraic Comb.* **8** (1998) 5-13.

[2] V. Albiero and E. Sopena, "Combinatorial construction of timestamp systems and interpolation systems", *Proc.FPSAC'95, Formal power series and algebraic combinatorics, Marne-la-vallée, France*, (1995) 11-18.

[3] O.V. Borodin, A.V. Kostochka, J. Nešetřil, A. Raspaud and E. Sopena "On universal graphs for planar oriented graphs of a given girth", *Discrete Math.* **188** (1998) 73-85.

[4] P. Hell and J. Nešetřil "On the complexity of H-coloring", *J. Combin. Theory B* **48** (1990) 92-110.

[5] A. V. Kostochka, E. Sopena and X. Zhu, "Acyclic and oriented chromatic numbers of graphs", *J. Graph Theory* **24** (1997), 331-340.

[6] T.H. Marshall, "On $\mathcal{P}$-universal graphs", preprint, 2001.

[7] J. Nešetřil and A. Raspaud "Colored homomorphisms of colored mixed graphs", *J. Combin. Theory B* **80** (2000) 147-155.

[8] J. Nešetřil, A. Raspaud and E. Sopena, "Colorings and girth of oriented planar graphs", preprint, 1995.

[9] E. Sopena, “The chromatic number of oriented graphs”, *J. Graph Theory* **25** (1997), 191-205.