Universal graphs with forbidden subgraphs

Gregory Cherlin *
Department of Mathematics,
Rutgers University, U.S.A.

November 5, 2005

Abstract

The graph theoretic problem of identifying the finite sets $\mathcal{C}$ of constraint graphs for which there is a countable universal $\mathcal{C}$-free graph is closely related to the problem of determining for which sets $\mathcal{C}$ the model companion $T^*_{\mathcal{C}}$ of the theory of $\mathcal{C}$-free graphs is $\aleph_0$-categorical, and this leads back to combinatorics. Little is known about these theories from any other perspective, such as stability theory.

1 Introduction

As is well known, Rado's universal graph [R64] is not just universal, but homogenous (in the Fraïssé sense), and there are analogous K_n-free universal homogeneous graphs, where K_n is the complete graph on n vertices and n is fixed [He71]. Universal homogeneous graphs can be classified in terms of their *finite induced subgraphs*, according to Fraïssé's theory [F86/00], and it turns out that one can classify all the countable universal homogeneous graphs [LW80], and there are few: countably many in all, and mainly of Henson's type, or their complements. The situation becomes more complex when one moves to more general types of combinatorial structures [He72].

It turns out that even without homogeneity, universality—or more properly, the existence of a universal object of a specific kind—is often a very restrictive condition. Pach showed [P81] there is no universal countable planar graph, and Komjáth and Pach showed [KP84] that there is no countable universal graph within the class of $K_{m,n}$-free graphs, with $K_{m,n}$ the complete bipartite graph with classes of size m, n respectively, with the exception of the stars with $\min(m, n) = 1$ and $\max(m, n) \leq 3$. Not long after, in conjunction with Mekler [KMP88], they showed that there is a countable P_n-free graph for P_n a path of

*Supported by NSF Grant DMS-0100794. Thanks to the IGD, Université Claude Bernard (Lyon-I) for their hospitality in 2004-5.

length n. In [KP84] the universality problem under consideration was treated also from a set theoretic point of view, and in [KP91] some other cases were treated, notably graphs containing on circuit C_n for $n \geq n_0$, with n_0 fixed. But our focus here is on the following portion of the larger question:

Problem 1 *For a given finite set $\mathcal{C}$ of finite connected graphs, is there a universal countable $\mathcal{C}$-free graph?*

Now a good deal more should be said about this problem, beginning with precisely what is meant by it, why I focus on this particular variant, and whether there is any particular reason to limit oneself to graphs, but I will come back to this. The point that is worth noticing, once this has all been clarified, is that this is a simple yes/no question about a finite set of data: the set $\mathcal{C}$.

Empirically, it seems that the answer has a strong tendency to be "no". So we may suspect that, or in any case wonder whether, the exceptional constraint sets $\mathcal{C}$ are "simpler" than the others in some definite way. From my perspective, the fundamental question is this:

Problem 2 *Is Problem 1 algorithmically decidable?*

The two obvious ways to solve this would be: (a) solve the problem explicitly; (b) encode a known undecidable problem into this one. Progress has been made in both directions. There has never been a really clear indication as to which way the general problem will go; as we develop tools for analyzing this type of problem we are gradually able to solve broad problems of this type, but we are regularly surprised by the answers that emerge. What is clear is that all of the problems we have solved to date turn out to have simple but somewhat unpredictable answers.

A variety of techniques from combinatorics and model theory have been brought to bear on these problems. The natural tool from model theory is not so much Fraïssé's theory of amalgamation classes (or quantifier elimination), as the Robinson theory of existential completeness. In fact the theories of graphs with "forbidden subgraphs" form a very well-behaved subclass of the universal theories, and the existence of a universal graph turns out to be closely bound up with both the $\aleph_0$-categoricity of the associated theory, and the *local finiteness* of the associated algebraic closure operation [CSS99]. In [CSS99] Shelah, Shi, and I laid out the model theoretic approach to these problems quite systematically. In my view the time had come mainly because we wanted to prove some results on *existence* of universal graphs, and the conventional machinery seemed to be reaching its limits. Indeed, it is in these positive results that mathematicians of all stripes tend to bring in a model theoretic technique, usually Fraíssé's. But that amounts in practice to proving $\aleph_0$-categoricity via quantifier elimination, which in principle is something one can *always* do, but which is frequently undesirable: the difference is between knowing that there are finitely many types, and knowing what the types actually are.

Our model theoretic analysis will lead us in §4 to an interesting class of theories, within which we would like to identify the $\aleph_0$-categorical ones. It could be interesting to pursue this model theory also from a more "modern" point of view.

To conclude this introduction, let me say what Problem 1 actually means. In the rest of this article I will describe the results to date, and try to indicate why I am coming to believe that at least for the case of a *single* constraint, $\mathcal{C} = \{C\}$, the problem is likely to be decidable. And I will try to clarify what the model theoretic ideas are that bear on this issue, and why the first order theories involved are so nice.

We return to Problem 1: *given a finite set* $\mathcal{C}$ *of finite connected graphs, is there a universal countable* $\mathcal{C}$*-free graph?* The term "$\mathcal{C}$-free" requires elucidation. Given a "forbidden" (or "constraint") graph C, a C-free graph is one in which no *subgraph* is isomorphic to C. Now when a model theorist says "subgraph" he may mean "*induced* subgraph"; but I really do mean "subgraph". Thus for example if P_2 is the path of length 2 (three vertices), then a P_2-free graph is one of maximal vertex degree at most 1; and for any n the condition of P_n-freeness is of comparable power, as seen in [KMP88]. Of course, we call a graph $\mathcal{C}$-free if it is C-free for all $C \in \mathcal{C}$.

Another term that urgently needs clarification is universal: we are speaking of countable graphs into which every $\mathcal{C}$-free graph embeds, but there are two notions of embedding available: the model theoretic strong embedding relation, which for graph theorists is an isomorphism with an induced subgraph, and the more usual graph theoretic embedding which for model theorists is an injective homomorphism. Now for our forbidden graphs we took the weak notion of embedding, but we may speak either of strongly universal graphs or weakly universal graphs, and there is clearly something more satisfying about strongly universal graphs, when they exist. Remarkably, in every case of Problem 1 for which a weakly universal graph is known to exist, a strongly universal one also exists, and we see no reason to expect this equivalence ever to fail. So as a rule one proves existence of a strongly universal graph when on the positive side, and nonexistence of a weakly universal graph when on the negative side.

What we are trying to understand is how to interpret the existence of a (countable) universal $\mathcal{C}$-free graph as saying something definite about $\mathcal{C}$. For the case of a single constraint, $\mathcal{C} = \{C\}$, we have a rather good idea of what this should be (the Solidity Conjecture, below, captures most of this idea); and we can reformulate this also in the general case, but in the process it becomes considerably less concrete. In any case all such conjectures are currently unproved, though we are inching toward them.

A few more general comments are in order. My bibliography is narrowly focused, and may have omitted some things that should be there even with my narrow focus. There is a much larger subject here, touching on homogeneity, Ramsey theory, combinatorial set theory, and even, as we now learn (partly from Paris) topological dynamics.

Convention 1 *All graphs are at worst countably infinite; any graph being used as a constraint is finite and (with few exceptions) connected.*

2 Some results

This section is very empirical: we look at theorems concerning universal graphs with one forbidden subgraph. Taken in conjunction with some theory to be presented later on, these theorems suggest some sweeping conjectures, at least in this case, and perhaps more generally.

2.1 Single constraints

A connected graph is 2-*connected* if it remains connected after removal of a point. An edge is 2-connected, but a more typical example of a 2-connected graph woud be a cycle.

Theorem 1 ([FK97a]) *Let C be a 2-connected graph. Then there is a universal C-free graph if and only if C is complete.*

This includes for example the complete bipartite case $C = K_{m,n}$ of [KP84] except for the easy case $\min(m, n) = 1$, as well as the case of a cycle C_n considered in [CK94]. We can see that the subject has moved quickly from the anecdotal to something broad and structural.

The result of [FK97a] is still more general. The maximal 2-connected subgraphs of a graph are called its *blocks*. Every connected graph can be represented canonically as a *tree of blocks*. One takes as the vertices the blocks together with the cut vertices (whose removal disconnects the graph). One joins vertices representing cut vertices to vertices representing the blocks that contain them. (A less redundant representation might suppress cut vertices lying in only two blocks.)

Theorem 2 ([FK97a]) *Let C be a graph one of whose blocks B has the following properties.*

1. *B is 2-connected and incomplete;*
2. *B does not embed isomorphically into any other block of C*

Then there is no universal C-free graph.

In fact, I believe the following.

Conjecture 1 (Solidity Conjecture) *Let C be a graph which has an incomplete block. Then there is no universal C-free graph.*

There are some theoretical reasons in back of this conjecture, which we will give in §4 after introducing a model theoretic point of view.

At the opposite extreme from 2-connected graphs we find trees. The following was conjectured by Tallgren.

Conjecture 2 (Tree Conjecture) *Let C be a tree. Then there is a universal C-free graph if and only if C is either a path, or can be reduced to a path by removing a suitable vertex.*

A tree will be called a *nearpath* if it is not a path, but reduces to one after deletion of a suitable a vertex. I believe that Shelah and I have now proved this conjecture [CSh05]. Unfortunately the analysis involves 11 distinct cases. Still, we are pleased that the number 11 turns out to be finite. Furthermore there is a very simple idea here that can be reused, though it seems that the proliferation of cases is likely to remain a feature of the landscape.

The positive part of this conjecture was known for a long time. Paths were treated in [KMP88], and nearpaths were treated by Tallgren in long unpublished work, which was finally combined with more recent work in [CT05]. On the negative side (nonexistence), some broad classes of trees were handled in [CST97] and [FK97b]. In the latter the following was proved.

Theorem 3 ([FK97b]) *Let C be a tree containing a unique vertex v_* of maximal degree d, with $d \geq 4$, and suppose that that v_* is adjacent to a leaf of C. Then there is no universal C-free graph.*

This result turns out to have much of the general case in it. In fact one can show that it implies the following directly [CSh05], which would be the first of our 11 cases.

Theorem 4 *Let C be a tree containing a unique vertex v_* of maximal degree d, with $d \geq 4$. Then there is no universal C-free graph.*

I will explain the reduction of Theorem 4 to Theorem 3 in the next section.

The restriction $d \geq 4$ is very natural since all the exceptional cases fall under $d = 3$, and the proof suggests that the main difference between the two cases comes from the structure of $(d-1)$-regular graphs.

It is tempting to extend the tree conjecture to say that any constraint graph C which allows a universal graph will have an associated tree of blocks which is either a path or a near-path, but this seems to be seriously wrong. In fact the following seems quite possible.

Conjecture 3 *Let C be a graph derived from a complete graph K_n by attaching exactly one finite path to each vertex of K_n. Then there is a universal C-free graph.*

This is not so much a conjecture as a family of problems.

What does seem to be the case is that the underlying tree of blocks should be a *star*, that is have a unique vertex of degree greater than 2.

Conjecture 4 (Generalized Tree Conjecture) *Let C be a constraint for which there is a universal C-free graph. Then the associated tree of blocks $\tilde{C}$ is a star.*

Combining this with the Solidity Conjecture we have a very restricted set of candidates for (single) constraints allowing universal graphs.

At this point we can probably see the need for an additional case study to clear the air, and there is a canonical candidate: graphs whose underlying tree structure is a path of length two, corresponding to two blocks joined by one cut vertex. This is carried out in [CT05].

Theorem 5 *Let $C = K_m \wedge K_n$ be a 2-bouquet (the sum of two complete graphs with one vertex from each identified). Then there is a universal C-free graph exactly when the following two conditions are met.*

1. $\min(m, n) \leq 5$;
2. $(m, n) \neq 5$

So we are saying in particular that there is a universal $K_5 \wedge K_n$-free graph *except* when $n = 5$. Symmetry works against existence!

With no theoretical basis at all, but simply on the grounds that we expect only "simple" constraints to allow a universal graph, we made the following false conjecture in [CSS99].

Conjecture 5 (Monotonicity Conjecture (Deceased)) *If there is a universal C-free graph, and if C_0 is an induced subgraph of C, then there is a universal C_0-free graph.*

Our case study refutes this, but it may not be far from the truth. A weaker version is given in the next section, one which has the combined virtues of being both true and useful, though it definitely is not the last word on the subject.

With the last case study in hand, one might hope for the following. (The terminology is borrowed from Shelah.)

Conjecture 6 (Noble conjecture) *Let C be a graph for which there is a universal C-free graph. Then at most one block of C has order 6 or more.*

The term "noble" here would refer to a block of maximal size. Shelah has extended the methods described in the next section to reduce the proof of this conjecture to some relatively limited cases. (This is really an inductive procedure, so if the conjecture turns out to be true his methods should lead to a proof; if there are exceptions the machinery will have to be cranked up afresh.)

My discussion here has been structural rather than historical, though I think one can more or less see the history as well. Up through [FK97a] there was no clear sign that any interesting examples of universal graphs (of this specific type) remained to be found. But the paper [K99] reversed that trend. One of the "remarks" in that paper is the following.

Theorem 6 *There is a universal $K_3 \wedge K_3$-free graph.*

This was the first indication that complete graphs and trees were not really two separate categories of constraint graphs, but that rather the mixed structure of a tree of blocks would need to be looked at. It was also a clear indication that a case study of 2-bouquets might be illuminating. I expected the favorable constraints to be more or less of the form $K_n \wedge K_3$ (with some sporadic noise), rather than the higher threshold that emerged, but on the whole the result of that study conformed to expectations, with some nuances.

Still the subject has a way of rebounding in unexpected ways, and while I have put forward some sweeping conjectures (and we have some strategies that may prove them eventually), the known facts do not yet constrain us a great deal. But at least we can say that in the case of trees the sporadic phenomena die out quickly.

One can easily envision a scenario in which the decidability of Problem 1 is established, but where we cannot actually solve, or even write a specific program to solve, individual cases. For example if one has all of the preceding conjectures—the Solidity, Generalized Tree, and Noble Conjectures—then one approaches the point at which one can show that only finitely many pieces of information are missing, at which point the problem is decidable, even if one has no idea what the decision procedure is.

2.2 More general universality problems

We will continue our discussion in the next section with some ideas that are not tied to particular examples. But here we round out the discussion by considering other universality problems which are of a more general character, some of which we exclude from our subject area, and others which we do not at all exclude but rather hope to bring into very much the same framework.

First, let us review some of our assumptions about the class $\mathcal{C}$ of forbidden subgraphs: its members are finite connected graphs, and there are finitely many of them.

Without connectedness, the notion of universality is not very relevant. For example, if we take as our constraint C the disjoint union $K_1 \oplus K_2$ of a vertex and an edge, the C-free graphs are those containing no edges, and those consisting of a single edge. There is no universal graph because these two types are incompatible; there is in fact a universal *pair* of graphs. On the other hand, if we have a connected constraint C, then the *disjoint sum* of any family of C-free graphs is again C-free, and if there is a countable family of (jointly) universal C-free graphs, then there is one such. What this suggests is that if we want to allow disconnected constraints then there are two relevant thresholds to investigate: the case in which there is a finite set of jointly universal C-free graphs, and the case in which there is a countable set of jointly universal C-free graphs [KP91]. And this leads to some interesting problems where the natural approach takes one out of the category of graphs and into more general classes [CS97]. But nobody has really followed that line up systematically, and it has a different character.

What about the restriction to finitely many finite constraints? Of course, if we want to consider the decidability of the general problem, this is the natural framework to adopt. But there are other, and I think more substantial, reasons to impose these constraints. First of all, if one allows infinite constraints then one enters a very different realm, much more set theoretic in nature. But if we consider infinite families of finite constraints then we are very close to normal graph theoretic practice; we can discuss bipartite graphs or forests, for example. Furthermore one can bring model theory to bear: we are looking at the models of a universal theory T , and looking for a countable universal model. When the set of constraints is finite, there is an associated theory T^*, the so-called model companion, whose models are the "existentially complete" models of T. This condition fails in general: for example if we deal with the case of forests, the existentially complete ones will be connected, and this is not something that can be expressed by any first order theory. On a more empirical level, we observe that one can identify explicitly all the finite sets $\mathcal{C}$ of cycles (or, for that matter, 2-connected graphs) for which there is a universal $\mathcal{C}$-free graph, and they are all of an obvious type; for infinite sets of cycles we do not know such a characterization, but we do know that new examples arise. It would be very nice to understand where these new examples come from—but we leave all that aside.

So much for what we exclude. On the other hand, we have said that we allow arbitrary finite sets of appropriate constraint graphs $\mathcal{C}$, and for that matter we have no strong reason to limit ourselves to graphs. As we have intimated, there are particular occasions where some extension of the category is even useful for the treatment of graphs [CS97]. In [CS01] it is shown that all problems of our allowed type, for arbitrary relational structures, are equivalent to problems which lie just outside the category of graphs: namely, graphs equipped with a partition of the vertices into two sets. We do not know whether these problems can in fact be encoded as universality problems for ordinary graphs, but it is clear that the more general problem can be investigated by very similar methods.

Returning to the case of ordinary graphs, what does the theory look like for *sets* of constraints? New phenomena arise which cast light on the case of a single constraint. Consider for example the following result, which provides a good illustration of the state of the theory.

Theorem 7 *Let $\mathcal{C}$ be a finite collection of finite cycles. Then there is a universal $\mathcal{C}$-free graph if and only if $\mathcal{C}$ consists of all cycles of odd order up to some bound.*

Actually a generalization reveals more clearly what is at stake.

Theorem 8 *Let $\mathcal{C}$ be a finite collection of finite 2-connected graphs. Then there is a universal $\mathcal{C}$-free graph if and only if $\mathcal{C}$ is* closed under homomorphism.

One has to understand closure under homomorphism in a not overly literal sense. A homomorphism between graphs is a function taking vertices to vertices which carries adjacent vertices to adjacent vertices. We take our graphs to be loop-free, and therefore adjacent vertices cannot map to the same vertex. Now when we say that $\mathcal{C}$ is closed under homomorphism, we mean the following: if $\hat{\mathcal{C}}$ is the closure of $\mathcal{C}$ under homomorphism, then $\mathcal{C}$-free graphs are $\hat{\mathcal{C}}$-free. Taking Theorem 8 in this sense, it generalizes Theorem 7.

Theorem 8 falls into two parts. One part asserts that if $\mathcal{C}$ is closed under homomorphism then there is a $\mathcal{C}$-free graph. This is a completely general fact.

Theorem 9 *Let $\mathcal{C}$ be a finite set of finite connected graphs closed under homomorphism. Then there is a universal $\mathcal{C}$-free graph.*

This holds for model theoretic reasons indicated in §4. In brief: the model companion $T^*_{\mathcal{C}}$ of the theory $T_{\mathcal{C}}$ associated with $\mathcal{C}$ is $\aleph_0$-categorical, and hence provides a canonical countable universal $\mathcal{C}$-free graph.

In the other direction, we need to show that if $\mathcal{C}$ is not closed under homoomorphism then there is no universal $\mathcal{C}$-free graph. This is more delicate, and requires a specific argument. Notice however that when $\mathcal{C}$ consists of a single 2-connected graph C, then closure under homomorphism is equivalent to the completeness of C, and we find ourselves dealing with Theorem 1. The general result can be obtained by repeating the proof in [FK97a].

There is a good deal more in this vein, and one can anticipate that what holds in the case of one constraint may also hold in a more subtle form for finite sets of constraints. But to give the proper generalization of, say the Solidity Conjecture requires more of the theory. For the record, it would read as follows.

Conjecture 7 (Generalized Solidity Conjecture) *Let $\mathcal{C}$ be a finite set of finite connected graphs, and suppose that there is universal $\mathcal{C}$-free graph. Then the associated algebraic closure operator is* unary *in the sense that*

$$\mathrm{acl}(A) = \bigcup_{a\in A} \mathrm{acl}(a)$$

We are referring here to the model theoretic notion of algebraic closure, taken relative to the model companion $T^*_{\mathcal{C}}$; we will have more to say about this notion, which can be decoded into purely graph theoretic terms, in §4.

3 Pruning

There is a simple and useful idea which has nothing to do with model theory, and everything to do with the meaning of universality: *pruning.*

We prune a tree by removing all its leaves. We can generalize this a little by removing the leaves (vertices of degree 1) from any graph; and we will generalize the idea a great deal more, to good effect.

We introduce the following notation. For C a connected graph, C^- is the graph with its leaves removed. For $\mathcal{C}$ a set of connected graphs, $\mathcal{C}^-$ is the set $\{C^- : C \in \mathcal{C}\}$.

Theorem 10 ([CSh05]) *Let $\mathcal{C}$ be a finite set of finite connected graphs, and suppose there is a universal $\mathcal{C}$-free graph. Then there is a universal $\mathcal{C}^-$-free graph.*

Notice that this is an instance of the lamented Monotonicity Conjecture. Note also that we have been vague about which version of universality we refer to here: it may be either weak or strong.

Proof. For any $\mathcal{C}$-free graph G, let G° be the induced subgraph of G on the vertices of infinite order in G. If G is $\mathcal{C}$-free then G° is $\mathcal{C}^-$-free.

Now suppose Γ is a universal $\mathcal{C}$-free graph. It suffices to show that Γ° is a universal $\mathcal{C}^-$-free graph.

Let G be $\mathcal{C}^-$-free, and extend G to a graph $\hat{G}$ by adjoining infinitely many new vertices as neighbors to each vertex of G. Then $\hat{G}$ is $\mathcal{C}$-free, so $\hat{G}$ embeds into Γ, either weakly or strongly as the case may be. Evidently this embedding carries G into Γ° and gives us the required embedding. □

Now we can see how to deduce Theorem 4, concerning trees with a unique vertex of maximal degree $d \geq 4$, from the case treated in [FK97b], where the vertex in question was taken to be adjacent to a leaf. If C is any tree with a unique vertex of maximal degree $d \geq 4$ and there is a universal C-free graph, we may prune it repeatedly until we reach the point that C^- no longer has a vertex of degree d. Then there is still a universal C-free graph, but now the hypotheses of [FK97b] are satisfied, so we have a contradiction.

Similarly, if we call a tree *critical* when it is neither a path nor nearpath, but reduces to a path or nearpath on pruning, then the proof of the Tree Conjecture reduces to the critical case.

For another application of the same idea consider the Conjecture 3 relating to complete graphs with certain paths attached; view it not as a conjecture but as a series of problems to be solved. There is a natural partial ordering on this set of constraint graphs, with $C_1 \leq C_2$ if C_2 can be pruned (repeatedly) to give C_1. If some graphs C of this type do not allow universal C-free graphs, then what we need to know are the *minimal* graphs falling on the negative side, minimal in the sense of this partial order. In particular, if this set is finite then the problem is decidable. Now it would be very nice if this partial order were already well-quasi-ordered, so that any set of incomparable graphs would be finite—this would give a soft proof of the decidability of this particular problem. But the pruning procedure is rather coarse and does not allow this conclusion. This class of graphs actually is well-quasi-ordered under strong embedding (this follows from Higman's Lemma); and perhaps one can still refine this analysis to give a "soft" treatment of the problem.

Let us now give pruning at a more suitable level of generality. Most graphs do not have leaves, but they do have the structure of a tree of blocks, and the blocks which have only one vertex in common with another block are the leaves in this tree. So let us call these *block-leaves*. A block-leaf is not just a graph; it is a graph with one vertex marked, a *pointed graph*. We could prune our graphs by taking off all the block-leaves (leaving the distinguished vertex alone, of course). But this is overkill. One may pick a particular isomorphism type of block-leaf, and then remove just the block-leaves which embed into it (with the distinguished vertex preserved). The most useful procedure is to take a minimal block-leaf and prune with respect to it.

Theorem 10 continues to hold with this more general notion of pruning, with the same proof. And Shelah has found even more general notions of pruning which can be brought into play in the context of the Noble Conjecture. But this brings our story up to, and a little beyond, the present.

4 Some model theory

If $\mathcal{C}$ is a collection of finite graphs, we make the following definitions.

1. $\mathcal{G}_\mathcal{C}$ is the collection of $\mathcal{C}$-free graphs

2. $\mathcal{G}^*_\mathcal{C}$ is the collection of existentially complete $\mathcal{C}$-free graphs

3. $T_\mathcal{C}$ is the theory of $\mathcal{G}_\mathcal{C}$

4. $T^*_\mathcal{C}$ is the theory of $\mathcal{G}^*_\mathcal{C}$

A $\mathcal{C}$-free graph G is existentially complete if for any finite induced subgraph A of G and any finite graph B containing A as an induced subgraph, if the free join (amalgam) $G \oplus_A B$ of G and B over A is C-free, then there is a strong embedding of B into G over A.

Proposition 4.1 ([CSS99]) *If $\mathcal{C}$ is a finite collection of finite connected graphs then the theory $T^*_\mathcal{C}$ is complete, any model of $T^*_\mathcal{C}$ is existentially complete, and there is a countable strongly universal $\mathcal{C}$-free graph if and only if for every n the space S_n of n-types for the theory $T^*_\mathcal{C}$ is countable (the theory is "small").*

One can adapt this to the case of weak universality using positively existentially complete structures and positive existential types, but on the model theoretic side the theory is less familiar.

As an empirical observation, when the theory $T^*_\mathcal{C}$ is small it tends to be $\aleph_0$-categorical—equivalently, the associated type spaces are finite—and this means that we then have a canonical countable universal $\mathcal{C}$-free graph, namely any model of $T^*_\mathcal{C}$. For example, when $\mathcal{C} = \{C\}$ with C a tree, we have the following:

1. T^*_C is $\aleph_0$-categorical if and only if C is a path;

2. T^*_C is small if and only if C is a path or nearpath.

So it seems that our graph theoretic problem turns out to be a mild perturbation of a problem which may have a cleaner solution: determining when $T^*_\mathcal{C}$ is $\aleph_0$-categorical. And it is this problem which lends itself to a considerable theoretical analysis, coming back to an explicit combinatorial problem. For this we need the notion of *algebraic closure*.

Definition 4.2 *Let G be an existentially complete C-free graph, A a finite set of vertices, and b a vertex. Then b is* algebraic *over A if and only if there is a finite graph B containing the induced graph on $A \cup \{b\}$ such that the set of images $f(b)$ for f a strong embedding of B into G over A is finite.*

We can define the algebraic closure of an arbitrary set as the algebraic closure of its finite subsets.

Theorem 11 ([CSS99]) *Let $\mathcal{C}$ be a finite collection of finite connected graphs. Then $T^*_\mathcal{C}$ is $\aleph_0$-categorical if and only if the algebraic closure operation is* locally finite *in the sense that the algebraic closure of any finite set is finite.*

Now this is a remarkable fact which ought not be true. It is immediate that $\aleph_0$-categoricity implies local finiteness of the algebraic closure operation—a little too immediate. The converse tends to fail. For example, if one takes any structure at all and replaces its elements by equivalence classes of infinite size, the algebraic closure operation in the modified structure will be degenerate: $\mathrm{acl}(A) = A$. But in most other respects the theory is unchanged, and in particular if it was not $\aleph_0$-categorical to begin with, it will not become so.

This theorem depends very precisely on the fact that the constraints in $\mathcal{C}$ are being treated as subgraphs rather than induced subgraphs (otherwise we could easily force equivalence relations into the picture). For the proof of the theorem one finds that the number of 1-types over a finite algebraically closed set is finite (though possibly exponentially large relative to the size of the set). Note that even when $\mathrm{acl}(A) = A$ for all A, there may still be a fairly complicated type structure.

Now after further analysis of the algebraic closure operator in terms of the constraints in $\mathcal{C}$, one can describe the operation as generated by iterated applications of a finite number of multivalued functions, so that the operation is locally finite if and only if there is a bound on the number of times these functions can be applied to generate new elements. One can also prove the following general facts.

Proposition 4.3

1. *If $\mathcal{C}$ is closed under homomorphism then the algebraic closure operation associated with $\mathcal{C}$ is degenerate (*$\mathrm{acl}(A) = A$*) and in particular $T^*_{\mathcal{C}}$ is $\aleph_0$-categorical.*

2. *If $\mathcal{C}$ consists of solid graphs then the associated algebraic closure operation is* unary *in the sense that*

$$\mathrm{acl}(A) = \bigcup_{a \in A} \mathrm{acl}(a)$$

Here we begin to reap something substantial from the theory. Note in particular that we get some theoretical basis for the Solidity Conjecture, particularly when we look in [FK97a] to see how the failure of unarity comes into play to make the operation iterable. And we can see how to generalize that conjecture from the case of one constraint to a finite set of constraints (Conjecture 7).

Finally, returning to the case of 2-bouquets, we have the following.

Theorem 12 *Let $C = K_m \wedge K_n$ be a 2-bouquet. Then the following are equivalent.*

1. *There is a strongly universal countable C-free graph.*

2. *There is a weakly universal countable C-free graph.*

3. *The theory T^*_C is $\aleph_0$-categorical.*

4. *The algebraic closure operation associated with C is locally finite.*

5. $\min(m,n) \le 5$ *and* $(m,n) \ne (5,5)$.

The proof is combinatorial and is based on the equivalence of the last two points. One works out the algebraic closure operation explicitly, and one considers the structure of an infinite sequence of points each algebraic over the preceding (we know the operation is unary). Each step along the way has a finite set of "witnesses" to algebraicity, of bounded order, and applying the Δ-system lemma we can uniformize this over a subsequence. What we find is that for $\min(m,n) \le 5$ the uniformized witnesses are pairwise disjoint, that for $\min(m,n) \le 4$ this is impossible, while for $\min(m,n) = 5$ this leads quickly to (a) $m = n = 5$ and (b) a fairly good description of what the sequence should look like. At this point, one realizes that there is indeed such a sequence.

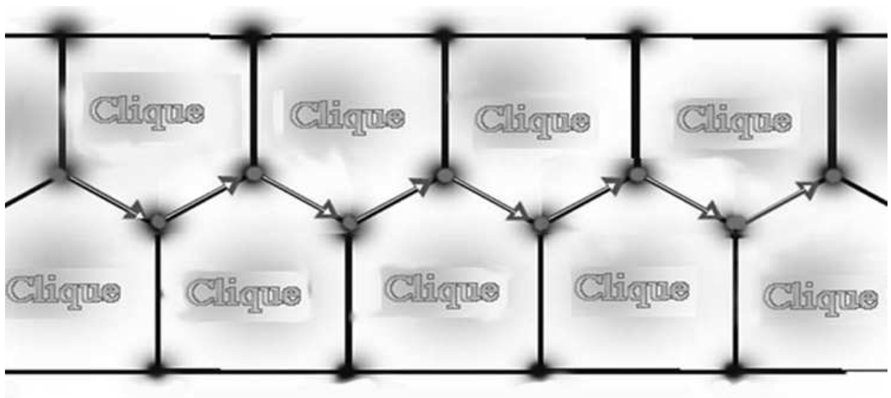

$(\mathbf{K}_5 + \mathbf{K}_5)$ **- free**

For $m,n \ge 6$ one gets the general picture more quickly and one finds a construction refuting local finiteness (this time the uniformized witnessing sets are not disjoint, and in fact the whole sequence of witnesses resembles the uniformized sequence).

All that remains after this lengthy computation and construction is to show that in the negative cases, the constructions can be adapted to blow up the type structure of T^*_C and eliminate a weakly universal C-free graph. This turns out to be easy in this context, though in other work this stage turns out to be less straightforward.

So by focusing on local finiteness we arrive at what is sometimes a combinatorial manageable problem, which we can solve completely and then adjust to solve the original graph theoretic problem. In other cases we first need a preliminary pruning to reach problems which are sufficiently tight in structure to be analyzed explicitly.

One of the complications standing in the way of an explicit resolution of Problem 1 is the wealth of examples that arise from a slight generalization of Proposition 4.3 (1), as follows.

Proposition 4.4 *Let $\mathcal{C}, \mathcal{C}'$ be two sets of constraints such that $T^*_{\mathcal{C}}$ is $\aleph_0$-categorical and $\mathcal{C}'$ is closed under homomorphism. Then $T^*_{\mathcal{C}\cup\mathcal{C}'}$ is $\aleph_0$-categorical.*

More explicitly: $\mathcal{C}$ and $\mathcal{C} \cup \mathcal{C}'$ have the same associated algebraic closure operation. I don't think we have real grounds to expect the full problem to be decidable, but it possible that what we are looking for is a combination of constraint sets closed under homomorphism with some exceptional "small" cases.

From a graph theoretic point of view, the main thing we would like to know about the model theory of the theory $T^*_{\mathcal{C}}$ is whether this theory is $\aleph_0$-categorical. But now that this class of theories has been seen to lend itself to systematic analysis, it would be interesting to develop their model theory from a more modern perspective, to try to identify the stable or simple theories in the class, and to see what their fine structure is. We do not see any clear "dichotomies" on the graph theoretic side.

References

[CK94] G. Cherlin and P. Komjáth There is no universal countable pentagon-free graph. *J. Graph Theory*, 18:337–341, 1994.

[CSh05] G. Cherlin and S. Shelah Universal graphs with a forbidden subtree. Preprint 2005. Shb:850.

[CSS99] G. Cherlin, S. Shelah, and N. Shi Universal graphs with forbidden subgraphs and algebraic closure Advances in Applied Mathematics 22:454–491, 1999.

[CS96] G. Cherlin and N. Shi. Graphs omitting a finite set of cycles. *J. Graph Theory*, 21:351–355, 1996.

[CS97] G. Cherlin and N. Shi. Graphs omitting sums of complete graphs, *J. Graph Theory*, 24:237–247, 1997.

[CS01] G. Cherlin and N. Shi. Forbidden subgraphs and forbidden substructures. *J. Symbolic Logic*, 66:1342–1352, 2001.

[CST97] G. Cherlin, N. Shi, and L. Tallgren, Graphs omitting a bushy tree. *J. Graph Theory*, 26:203–210, 1997.

[CT05] G. Cherlin and L. Tallgren, Graphs omitting a near-path or 2-bouquet. Submitted, 2005.

[F86/00] Roland Fraïssé, *Theory of relations*, Revised edition. With an appendix by Norbert Sauer. Studies in Logic and the Foundations of Mathematics 145, North-Holland Publishing Co., Amsterdam, 2000. ii+451 pp. ISBN 0-444-50542-3 First edition, translated by Peter Clote: *Studies in Logic and the Foundations of Mathematics* 118:, North-Holland Publishing Co., Amsterdam, 1986, xii+397 pp., ISBN 0-444-87865-3.

[FK97a] Z. Füredi and P. Komjáth. On the existence of countable universal graphs. *J. Graph Theory* 25:53–58, 1997.

[FK97b] Z. Füredi and P. Komjáth. Nonexistence of universal graphs without some trees, *Combinatorica*, 17:163–171, 1997.

[GK96] M. Goldstern and M. Kojman. Universal arrow free graphs, *Acta Math. Hungary* 73:319–326, 1996.

[He71] C. Ward Henson A family of countable homogeneous graphs. *Pacific J. Math.* 38:69–83, 1971.

[He72] C. Ward Henson Countable homogeneous relational structures and $\aleph_0$-categorical theories, *J. Symbolic Logic* 37:494–500, 1972.

[K99] P. Komjáth. Some remarks on universal graphs. *Discrete Math.* 199:259–265, 1999.

[KMP88] P. Komjáth, A. Mekler and J. Pach. Some universal graphs, *Israel J. Math.* 64:158–168, 1988.

[KP84] P. Komjáth and J. Pach. Universal graphs without large bipartite subgraphs. *Mathematika*, 31:282–290, 1984.

[KP91] P. Komjáth and J. Pach. Universal elements and the complexity of certain classes of infinite graphs. *Discrete Math.* 95:255–270, 1991.

[LW80] Lachlan, A. H.; Woodrow, Robert E. Countable ultrahomogeneous undirected graphs, *Trans. Amer. Math. Soc.* 262:51–94, 1980.

[P81] J. Pach. A problem of Ulam on planar graphs. *Eur. J. Comb.* 2:357–361, 1981.

[R64] R. Rado. Universal graphs and universal functions. *Acta Arith.* 9:331–340, 1964.