

Abstract

We consider the problem of the existence of universal countable C -free graphs with C a connected finite graph. For C a tree arising by from a path by adjunction of one additional edge we show that a universal countable C -free graph exists. We determine precisely the 2-bouquets C (i.e., unions of two complete graphs with just one point in common) for which a universal countable C -free graph exists. We lay out some elements of a program for determining all the connected finite graphs C for which a countable universal C -free graph exists. One element of this program is the Tree Conjecture, which is now proved [2]. Our methods involve a mixture of model theory and combinatorics, with the Δ -system lemma playing a significant role.

KEYWORDS: graph, universal, tree, bouquet, model theory, existential completeness, delta system

Universal graphs with a forbidden near-path or 2-bouquet

Gregory Cherlin*

Department of Mathematics, Rutgers University
110 Frelinghuysen Road
Piscataway, N.J. 08854

Lasse Tallgren

Nokia Research Center, P.O. Box 407
FIN-00045 NOKIA GROUP, FINLAND

April 15, 2013

1 Introduction

Given a set of “forbidden” subgraphs \mathcal{C} , we may consider the collection of all graphs G containing no subgraph isomorphic to one of those in the class \mathcal{C} . Such graphs are said to be \mathcal{C} -free. A general question that arises in various forms is whether this class is well behaved in some appropriate sense. While this question may be formalized in many different ways, one expects that the answer will usually be negative, and what is wanted is a way of identifying the exceptionally well behaved constraint sets.

The version of the problem that we have taken up is the existence of a *universal* graph, by which we mean a countable \mathcal{C} -free graph G^* with the property that every countable \mathcal{C} -free graph is isomorphic to an *induced subgraph* of G^* . This turns out to have a very concrete meaning in terms of the original constraint set \mathcal{C} , and is closely related to the halting problem for a certain explicit computation procedure associated with \mathcal{C} .

Let us illustrate this by a trivial example. If \mathcal{C} consists of a single constraint graph C which is an n -star, that is a tree consisting of a single vertex connected to n leaves, then the \mathcal{C} -free graphs are those of bounded degree, less than n . In this case, any connected graph

*Supported by NSF Grant DMS 0100794

which is regular of degree $n - 1$ will be a connected component in any \mathcal{C} -free graph containing it, which immediately implies that there is no universal \mathcal{C} -free graph for $n > 3$.

This analysis has two ingredients, both of which are valid for *any* finite set of connected constraint graphs. The edge relation in the graph provides a notion of “neighbor” with a finiteness constraint coming from the constraint set; then “connected components” are defined by iterating the “neighbor” relation. When working with constraints other than a single star, the “neighbor” relation will not coincide with the original edge relation, but rather with a more complicated but completely explicit relation; once one has the correct “neighbor” relation, one defines the analog of connected components by iterating it, and then the main question is whether the connected components are finite or not. Let us illustrate this with a slightly less trivial case: graphs omitting a cycle of length 4. In this case, it turns out that a single vertex has no neighbors in our sense; but a finite set A of vertices does, namely all those vertices which are adjacent to at least two vertices in A ; obviously, these neighbors are tightly controlled by the forbidden subgraph. Iterating this operation one defines the “closure” of the set A , which in general may be infinite, and this leads to a proof of the nonexistence of a universal graph with this constraint. We will define this closure operator explicitly in §4 in the case of a 2-bouquet. This is based on the general analysis given in [3]. For model theorists this would just be the algebraic closure operator in the class of existentially complete \mathcal{C} -free graphs, but in any case it is necessary to work out explicitly what this means in combinatorial terms in order to make real use of it.

Experience confirms that constraint sets allowing universal graphs are rare (cf. [1-16]). The problem is to identify those exceptional choices for \mathcal{C} which allow a universal graph. A very general result of Füredi and Komjáth [8], which subsumes a number of examples considered earlier, states that in the case in which \mathcal{C} consists of a single 2-connected graph C , then the corresponding universal \mathcal{C} -free graph exists if and only if C is complete—in which case the universal graph is afforded by a general method of Fraïssé. In fact, the method of Füredi and Komjáth proves something considerably stronger, when combined with a general result from [3].

Fact 1.1 *If \mathcal{C} is a finite set of 2-connected graphs, then there is a universal \mathcal{C} -free graph if and only if the class \mathcal{C} is closed under homomorphism.*

Here a *homomorphism* collapses certain sets of independent vertices to a single vertex. *Closure under homomorphism* is defined as follows: for each constraint C in \mathcal{C} , each homomorphic image C^* of C is forbidden, in the sense that C^* has a subgraph isomorphic to one

of the graphs in \mathcal{C} . The case in which \mathcal{C} consists of all cycles of odd length up to a given order is a typical example. For a single graph, our formulation is equivalent to the statement as given by Füredi and Komjáth in the 2-connected case.

This result may suggest that there should be a completely explicit characterization of the sets \mathcal{C} of constraints for which the corresponding universal graph exists. This question is completely open, but at least it is well-posed. Stated rigorously, the question is this: is there an algorithm taking as input a finite set \mathcal{C} of connected graphs, and determining whether or not the corresponding class of \mathcal{C} -free graphs has a universal graph? The combinatorial content of this decision problem can be analyzed in detail, using the methods of model theory; cf. [3].

We will review enough of the combinatorial framework of [3] to keep the present paper self-contained, as need arises. It will be convenient for the purposes of our discussion in the remainder of this introduction to use the notation $\text{acl}_{\mathcal{C}}$ (read “algebraic closure” relative to \mathcal{C}) freely, without defining it; this is the closure operator alluded to above, and defined in the relevant instances in §4 below.

In the present paper, we focus on the case of a single constraint graph C , taken to be finite and connected, though we give some results in more general form. In the case of a single constraint graph, combining the known results together with some plausible conjectures yields a robust, though incomplete, picture of what we may expect. Call a graph *solid* if every induced subgraph which is 2-connected is complete (in other words, the blocks are complete). The main conjecture is the following.

Conjecture 1 (Solidity Conjecture) *If C is a finite connected graph for which there is a universal countable C -free graph, then C is solid.*

There is some theoretical support for this in terms of the closure operator $\text{acl} = \text{acl}_C$ associated to C ; indeed, in those terms we would strengthen the conjecture as follows. We call a closure operator cl *unary*, or *degenerate*, if we have $\text{cl}(A) = \bigcup\{\text{cl}(a) : a \in A\}$.

Conjecture 2 (Unarity Conjecture) *If \mathcal{C} is a finite set of finite connected graphs for which there is a universal countable \mathcal{C} -free graph, then $\text{acl}_{\mathcal{C}}$ is unary.*

In the Füredi-Komjáth construction, and its ancestors, one sees clearly how the failure of unarity is relevant. But in this general form the Unarity Conjecture is unequivocally ambitious. This conjecture should also be compared with the Reduction Conjecture given below.

For the case of a single constraint C , it can be shown that C is solid if and only if acl_C is unary. Thus the Unarity Conjecture reduces

to the Solidity Conjecture in the case of a single constraint. In the 2-connected case, this conjecture reduces to the known result of Füredi-Komjáth.

A more ambitious but ill-starred conjecture was given in [3]. This conjecture, which had reasonable empirical support, but not much of a theoretical basis, goes as follows.

Conjecture 3 (Monotonicity Conjecture) *If C is a finite connected graph for which there is countable universal C -free graph, and if C' is a connected induced subgraph of C , then there is a universal countable C' -free graph.*

This conjecture is *false*, and will be refuted in the present paper. In spite of this, it seems to represent a healthy point of view, barring really major surprises. It may be that the exceptions to the Monotonicity Conjecture are sporadic, and that the problem is to classify them. The idea behind this conjecture is as follows. The existence of a countable universal C -free graph is known to be controlled by the behavior of the associated algebraic closure operator acl_C , which should not be too complicated. The latter condition *seems* to require C to be “small”. More generally, the complexity of the algebraic closure operation seems to reflect the complexity of the constraint graph.

We do not know what notion of “smallness” or simplicity is actually involved here, or whether that intuition is at all accurate, and that is really the entire issue. The present paper shows that whatever notion of simplicity may be involved, it is one which is not inherited by induced subgraphs. At a combinatorial level, the analysis in §4 depends on an application of the Δ -system lemma in a case in which the “heart” of the Δ -system may turn out to be empty in some exceptional small cases. While we do not consider the Monotonicity Conjecture to be completely discredited, its failure certainly complicates matters in challenging ways.

The Solidity Conjecture is a special case of the Monotonicity Conjecture, but it is also a special case of the Unarity Conjecture. The latter seems more robust, as it involves some theoretical considerations.

Solid graphs are “tree-like” in that they arise from trees by blowing up vertices appropriately to complete graphs. For the case of tree constraints there is a longstanding conjecture due to the second author, whose proof is now complete.

Conjecture 4 (Tree Conjecture) *For C a tree of order n the following are equivalent.*

1. *There is a countable universal C -free graph.*
2. *There is a countable strongly universal (that is, universal with respect to embeddings as induced subgraphs) C -free graph*

3. C contains a path of order $n - 1$ (i.e., length $n - 2$).

We will find it convenient to refer to such a tree as a near-path. The first result of the present paper is the following, which confirms one direction of the Tree Conjecture.

Theorem 1 *If C is a finite set of finite connected graphs, one of which is a near-path, then there is a universal countable C -free graph.*

Corollary 1.2 *If C is a near-path then there is a universal countable C -free graph.*

The method used here is a mixture of [12] with [5]. Basically one extends the results of [12] to the context of graphs with a finite partition of the vertices (not a “coloring” in the graph theoretic sense, but an arbitrary partition). We do not actually proceed in that precise fashion, but by a more direct route that gives less explicit information about the universal graph. Model theorists interested in understanding the properties of the interesting first order theory involved might want to follow the route of [12] more closely.

The full tree conjecture is proved in [2], using our Corollary 1.2 for one direction, and providing the complementary non-existence results for all cases of forbidden trees other than near-paths (or rather, for the “critical” cases, according to an appropriate inductive scheme).

We remark that a consequence of the Tree Conjecture is that the Monotonicity Conjecture holds for trees. It seems likely that a restricted form of the Monotonicity Conjecture does hold, for pairs of connected graphs $C' \subseteq C$ such that every block of C' is a block of C . This formulation includes both the case of trees, and the Solidity Conjecture.

Staying within the class of solid graphs, but passing to the opposite extreme from general trees, we come to the *2-bouquets*, by which we mean graphs obtained by taking the union of two complete graphs with a pair of vertices identified, one vertex from each; in other words, their free join over a common vertex. If solid graphs are considered tree-like, then the 2-bouquets may be considered to be P_1 -like, where P_1 is a path of length one, or possibly P_2 -like; this depends on the notion of “blowing up vertices” which is used, and we will use a naive terminology here (and a less naive one in §4).

Komjáth [11] was the first to consider a 2-bouquet in our context: the *bowtie* $K_3 \dagger K_3$ is the 2-bouquet formed from two triangles, and Komjáth saw that there is countable universal bowtie-free graph, a result which was both highly unexpected and difficult at the time, via the Fraïssé amalgamation technique. This was one of the motivations for developing the combinatorial theory of [3], which makes it quite easy to establish the existence of the corresponding universal graph

while bypassing the amalgamation method; the price one pays for this is that one learns relatively little about the actual structure of the corresponding universal graph, or its theory, as a result, but this seems a reasonable trade-off as the combinatorial complexity of the classes involved increases.

At the same time, Komjáth also showed that for $n \geq 3$, an n -bouquet formed from n copies of K_m with $m \geq 3$ never gives rise to universal graphs. More general n -bouquets with blocks of variable size have not been examined. The case of 2-bouquets appears to be quite exceptional; in any case, it contains a number of examples allowing universal graphs. We prove the following.

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

1. *There is a countable universal C -free graph.*
2. $\min(m, n) \leq 5$ and $(m, n) \neq (5, 5)$.

In view of the clause $(m, n) \neq (5, 5)$, we have a blatant refutation of the Monotonicity Conjecture. One can see explicitly in §3 and more theoretically in §4 why the case of $K_5 \uplus K_5$ arises as an exception; actually all of the cases $K_4 \uplus K_4$, $K_5 \uplus K_5$, and even $K_6 \uplus K_5$ require special consideration, but only one of them provides an actual exception. As noted above, this involves exceptional Δ -system configurations.

At this point we have, conjecturally, a fairly robust picture of what we may expect, in spite of recent surprises. Even this conjectural picture is incomplete; what is lacking from that point of view is a closer analysis of path-like constraints. Our results on 2-bouquets mean that the boundary between the “tight” constraints, allowing countable universal graphs, and all the rest is more ragged than has been expected. The fundamental question is the following.

Problem 1 (Universality Decision Problem) *Is there an algorithm which will decide, given as input a finite connected constraint graph C , whether there is a countable universal C -free graph?*

This problem makes perfectly good sense for finite sets of constraints as well (with respect to any class of combinatorial structures). We think it should be decidable for single constraints. It is far less clear whether it should be decidable in general. In the negative direction, there are some encoding results in [6], showing that various generalized forms of the problem are all equivalent. On the other hand, the Unarity Conjecture would tend to suggest decidability of the full problem. The kind of analysis undertaken in §4 illustrates the combinatorial content of this decision problem.

At this point, we will enter into a fuller discussion of the relationship between the complexity of acl_C and the existence of a C -universal

object. This is again based on the point of view of [3], and is illustrated here concretely by the work in §4. The relevant notion is the following.

Definition 1.3 *A closure operator cl on a set V is locally finite if $\text{cl}(A)$ is finite for each finite set A , and is uniformly locally finite if $|\text{cl}(A)|$ is bounded by a function of $|A|$ for A finite.*

In our context it follows from König’s Lemma that local finiteness of acl_C is equivalent to uniform local finiteness.

Fact 1.4 *Let C be a finite set of finite connected graphs for which acl_C is locally finite. Then there is a countable universal C -free graph.*

While not trivial, this follows from very general considerations, as shown in [3].

The converse is false, and the case of near-paths provides a simple counterexample. Still, the converse appears to be very nearly true, and in the case of a single constraint graph the following conjecture is not unreasonable.

Conjecture 5 (Reduction Conjecture) *Let C be a finite connected graph. Then the following are equivalent.*

1. *There is a countable universal C -free graph.*
2. *Either acl_C is locally finite, or C is a near-path.*

Like the ill-fated Monotonicity Conjecture, this conjecture has no theoretical basis, and in particular it only concerns the case of one constraint. To see why one might believe it nonetheless, notice that the proofs of nonexistence of universal graphs in §3 involve “decorating” an infinite set of the form $\text{acl}(A)$ with A finite. Once one has an infinite set of the form $\text{acl}(A)$ with A finite, one needs just one degree of freedom to complete the proof of nonexistence of the universal graph. For the most part a failure of local finiteness tends to lead to an easy proof of the nonexistence of a universal graph, along the lines of §3. However, the basic problem of local finiteness is difficult to settle in interesting cases. The key combinatorial decision problem is really the following.

Problem 2 (Local Finiteness Decision Problem) *Is there an algorithm which will decide, given as input a finite connected constraint graph C , whether the associated operator acl_C is locally finite?*

As a point of methodology, this is the problem we initially studied in the cases treated in §§3–4, after which the solution of the corresponding universality problems required just a few more details. While we will not give all of this preliminary analysis below, enough of it survives into the actual proofs to show clearly how this allows the problem to be

approached in a systematic way. We think it is useful to bear in mind both the close connection between these two decision problems, and the distinction between them. It is the *local finiteness problem* which can be closely analyzed by a uniform, canonical, combinatorial process: it is simply the *halting problem* for a specific computation encoded by the constraint set \mathcal{C} , in a way which is reminiscent of various simple models of computation. At first glance, in fact, one could hope for a reduction to an ordinary finite state automaton, but this does not work out.

Is it realistic to aim at a complete solution of the universality problem for the case of one constraint? Evidently the Tree Conjecture and the Solidity Conjecture are fundamental, and it is possible that both can eventually be proved by elaborating on the methods already used in special cases. These two problems are certainly worth examining further. The Tree Conjecture can be generalized a little more: we would expect a constraint C for which there is a countable universal C -free graph to be near-path-like, and in fact path-like in almost all cases (compare Komjáth's result on m -bouquets, in the case $m = 3$). To round out the picture, we also need to solve the following problem.

Problem 3 (Path-like Problem) *Call a graph C path-like if it is solid and its underlying tree structure is a path, or in other words if C is a string of complete graphs. Determine the pathlike constraints C for which there is a countable universal C -free graph.*

This is the problem we have taken up here, but only in the basic case of P_1 -like graphs, where we arrive at an unexpected answer, along lines broadly compatible with our expectations, but considerably different in detail. Not only is the Monotonicity Conjecture wrong, but the cross-over point ($n = 5$) is higher than we would have expected, given the rarity of universal graphs with forbidden subgraphs to date. On the general principle that “sporadic phenomena turn up early” it would be reasonable to carry on in the same vein, and classify the P_2 -like constraints which allow a countable universal graph. One expects the final *result* not to be much more complicated than the one in Theorem 2 below, though the *analysis* involved could well be more difficult. At the same time, this is the direction to explore if one is looking to uncover phenomena which could lead to an *undecidability* result. So it is possible that further exploration of this particular case could be very illuminating.

Remark. Further analysis (ongoing) by the first author and Shelah suggest that the analysis of universality problems for star-like constraints and path-like constraints hold the key to the analysis of universality problems for an arbitrary finite connected forbidden constraint. In addition to the “sporadic” phenomena encountered in our present

study, one more surprising family of examples of P_2 -type has been identified, and we believe we begin to see the shape of the general case, and a strategy pinning it down rigorously; key ingredients for this strategy are the analysis of bouquets given here, and the techniques of [2], especially the inductive approach given there. The analysis of bouquets, and considerably more, must all be viewed as lying within the base of this induction.

We will comment on related problems, dealt with in the literature, that are not covered by our setup. One may consider universality problems also for *uncountable* graphs of fixed cardinality, a set theoretic topic which is completely different from ours in character, but also highly developed. On the other hand, staying in the countable case, infinite constraint graphs have been considered, and again the nature of the problem changes radically; our approach has no bearing at all on such cases.

Staying in the countable case, and allowing only finite constraint graphs, there are two more variations on the theme of universality which have been investigated sporadically. We may consider *disconnected* constraint graphs, and we may allow *infinite sets* of finite constraints. These generalizations have much in common with the more limited problems we deal with. Allowing disconnected constraints forces some change in viewpoint, or at least in terminology, and it is no longer appropriate to speak of universal *graphs*, but rather universal *sets of graphs*, as Komjáth has pointed out. This point is illustrated by our analysis in §2. The general theory of [3] has only been worked out for sets of *connected* constraints. In order to handle disconnected constraints, it would be convenient to change categories, replacing the category of graphs by the category of partitioned graphs—graphs equipped with a finite partition of their vertices. This is useful in practice.

On the other hand, if infinite sets of finite connected constraint graphs are considered, our theory does not apply in its present form. This extension includes some cases which are of clear interest, notably the case of bipartite graphs, and it leads to a number of intriguing examples. The theory given in [3] definitely does not apply here, but perhaps it can be extended. We may illustrate the difference by examples involving cycles as constraints. If the constraint set \mathcal{C} consists of *finitely many* cycles, then there is a universal \mathcal{C} -free graph if, and only if, the set \mathcal{C} contains all cycles of odd order up to some specified order. But if \mathcal{C} is the set of all cycles having length *at least* some specified order, then there is a universal \mathcal{C} -free graph, and there is no result, or even conjecture, as to what other sets of forbidden cycles may allow countable universal graphs. What happens in this particular case is that the algebraic closure operation is not unary, but the universal ob-

ject exists anyway. On the other hand the algebraic closure operator appears to be locally finite in these particular cases. So what we have here is not so much a breakdown in the theory of [3] as a combinatorial difference in the analysis of acl , which could be explored further. The initial axiomatizability result of [3] is simply false at this level of generality, but it is not clear whether that result is really needed.

2 Near-paths

Definition 2.1 *A near-path is a tree consisting of a path with at most one additional edge adjoined.*

Our goal is the following.

Proposition 2.2 *Let L be a finite near-path, and let \mathcal{C} be a finite set of finite connected graphs with $L \in \mathcal{C}$. Then there is a universal countable \mathcal{C} -free graph.*

This result can be obtained in a number of loosely related ways. We will prove a more general result which has other applications. We work in the broader category of graphs with a vertex partition, or coloring by a specified finite set of colors. This coloring is not assumed to stand in any particular relation to the edges of the graph, and may be completely arbitrary.

Definition 2.3 *A d -graph is a graph G together with a coloring of the set its vertices by d colors, that is, a function $c : V(G) \rightarrow \{1, \dots, d\}$.*

Before dealing with near-paths, we generalize the corresponding result for paths, which is found in [12], to the context of d -graphs. In the article [12], graphs omitting a single path are considered. In the context of d -graphs, there are many isomorphism types of path of a fixed length, and the condition we impose is that all of these paths are forbidden; we also allow some side constraints at the same time. Furthermore, departing from our usual practice, we do not insist here that our constraint graphs be connected. We will prove the following.

Proposition 2.4 *Let \mathcal{C} be a finite set of finite d -graphs. Suppose that for some n , every isomorphism type of d -colored path of length n is represented in \mathcal{C} . Then there is a finite set of countable connected \mathcal{C} -free d -graphs which is universal for countable connected \mathcal{C} -free d -graphs.*

The universality condition given here means that we have a finite set \mathcal{X} of connected \mathcal{C} -free d -graphs such that any countable connected \mathcal{C} -free d -graph is isomorphic to an induced subgraph of one of the d -graphs in \mathcal{X} .

Before entering into the proof, observe the following corollary.

Corollary 2.5 *Let \mathcal{C} be a finite set of connected finite d -graphs. Suppose that for some n , every isomorphism type of d -colored path of length n is represented in \mathcal{C} . Then there is a universal \mathcal{C} -free d -graph.*

Proof. Let \mathcal{X} be the family of connected \mathcal{C} -free d -graphs afforded by Proposition 2.4. As we are assuming the d -graphs in \mathcal{C} are connected, the disjoint union of countably many copies of each graph in \mathcal{X} is also \mathcal{C} -free, and will serve as the desired universal graph. \square

Proof of Proposition 2.4. We proceed by induction on n , writing $\mathcal{X} = \mathcal{X}_n$. By our induction hypothesis, we have a finite family \mathcal{X}_{n-1} with the following property: any \mathcal{C} -free d -graph which contains no d -colored path of length $n-1$ embeds as an induced subgraph in one of the d -graphs in \mathcal{X}_{n-1} . We wish to extend the family \mathcal{X}_{n-1} to a suitable family \mathcal{X}_n . So we need only consider the structure of a connected \mathcal{C} -free d -graph G which contains some path P of length $n-1$.

We associate to the pair (G, P) a d' -graph G_P , with $d' = d \cdot 2^{n-1}$, as follows. The vertex set of G_P is $V(G) \setminus V(P)$. The graph structure on G_P is induced by G . The coloring c' in $V(G_P)$ is defined by $c'(v) = (c(v), c_P(v))$, where we set

$$c_P(v) = \{a \in V(P) : a \sim v\}$$

writing \sim for the edge relation in G .

The important point is the following, as in [12]: Every connected component of G_P contains no path of length $n-1$. This follows since G is connected, P has length $n-1$, and G has no path of length n .

Let $k = \max\{|C| : C \in \mathcal{C}\}$. For X a connected component of G_P , let $\mathcal{G}_k(X)$ be the set of d' -graphs with vertex set contained in $\{1, \dots, k\}$ which do not embed in X . As $\mathcal{G}_k(X)$ contains all paths of length $n-1$, by induction there is a finite set \mathcal{X}_X of connected countable $\mathcal{G}_k(X)$ -free d' -graphs which is universal for this class. Extend G_P to \hat{G}_P by replacing each connected component X of G_P by a d' -graph in \mathcal{X}_X into which it embeds as an induced subgraph. Then $\mathcal{G}_k(\hat{G}_P) = \mathcal{G}_k(G_P)$, and the connected components of \hat{G}_P lie among finitely many isomorphism types, since the set $\mathcal{G}_k(X)$ varies over a finite set of possibilities. Hence as G and G_P vary, there are, all together, finitely many distinct d' -graphs occurring as connected components in the various d' -graphs \hat{G}_P .

We make a further modification of the d' -graph \hat{G}_P . For any connected component X of the d' -graph \hat{G}_P which occurs at least k times as a connected component of that d' -graph, we extend \hat{G}_P so that there are infinitely many connected components of \hat{G}_P isomorphic to X . After this modification, each connected component of \hat{G}_P occurs with a multiplicity which is either less than k , or equal to \aleph_0 , and there are, all told, only finitely many isomorphism types occurring among the d' -graphs \hat{G}_P themselves.

Now each d' -graph \hat{G}_P can be decoded into a d -graph \hat{G} on $V(P) \sqcup V(\hat{G}_P)$. Furthermore $\mathcal{G}_k(\hat{G}) = \mathcal{G}_k(G)$, so \hat{G} is \mathcal{C} -free. Taking \mathcal{X}_n to consist of \mathcal{X}_{n-1} together with the d -graphs \hat{G} just constructed, our claim follows. \square

Proof of Proposition 2.2. Let $n = |L| - 2$; L contains a path P of length n , on $n + 1$ vertices. We will aim at finding a countable set of connected countable \mathcal{C} -free graphs such that each connected countable \mathcal{C} -free graph is isomorphic with an induced subgraph of one such. Then one forms a single universal graph by taking the disjoint union of countably many copies of each of these connected graphs.

By Proposition 2.4, there is a countable universal graph for the class of \mathcal{C} -free graphs containing no path of length $2n$. We proceed to analyze the structure of connected \mathcal{C} -free graphs which do contain a path of length $2n$. Fix one such graph G , and let P be a maximal path of length at least $2n$ in G , which may even be infinite in one or two directions.

If P is 2-sided infinite, then as G is L -free and connected we have $G = P$. Putting this case aside, and writing $P = (a_0, a_1, \dots)$, consider the induced graph on $V(P)$.

If P is infinite, then for $i \geq 2n$ the only neighbors of a_i are the two vertices $a_{i \pm 1}$. Thus there are only countably many isomorphism types occurring as induced graphs on the vertices of such a path P . We may accordingly fix one such isomorphism type for further analysis.

Now let G_P be the associated 2^{2n} -graph on the vertices $V(G_P) = V(G) \setminus V(P)$, where the coloring is given by $c(v) = \{i < 2n : a_i \sim v\}$. Observe now that G_P contains no path of length $2n$. Indeed, assuming the contrary, one finds a path Q of length at least $n + 1$ disjoint from P except at one endpoint, where the endpoint in question is a_i for some $i \leq 2n$. By maximality of P , we have $i > 0$, and hence the edge (a_{i-1}, a_i) exists as an additional edge which is attached to Q . But this configuration now contains the forbidden near-path, and provides a contradiction.

Since G_P contains no path of length $2n$, arguing as above we may find a finite number of 2^{2n} -graphs \hat{G}_P with the following two properties:

- (a) each graph G_P embeds into one of the given ones;
- (b) after decoding \hat{G}_P as a graph \hat{G} containing G , \hat{G} again omits \mathcal{C} .

So in this way we obtain the desired countable family of connected countable \mathcal{C} -free graphs. \square

3 2-Bouquets without universal graphs

There are two notions of universality studied in the literature. While the constraint graphs function as forbidden subgraphs (rather than

forbidden induced graphs), it is reasonable to require universality in the strong form we gave: every appropriate graph embeds as an induced subgraph. If we only require an embedding as a *subgraph*, we speak of *weak universality*. While from our own point of view this is not very natural, it is often considered in the graph theoretic literature, so we prefer to prove *existence* results using the “strong” definition and *nonexistence* results using the “weak” definition, since these are the strongest forms in their respective cases. The present section deals with nonexistence theorems, so we formulate them accordingly. This whole section is based on the theory of [3] and some computations along the lines of those found in the next section. We omit that preliminary analysis, which is not needed for the proofs in this section, and the interested reader can look into the following section to see the sort of calculation that is involved in such cases. It is quite helpful for finding constructions like those that follow here.

Proposition 3.1 *Let $C = K_m + K_n$ with $m, n \geq 6$. Then there is no weakly universal countable C -free graph; that is, there is no countable C -free graph which contains an isomorphic copy of every countable C -free graph as a subgraph.*

Proof.

We may suppose $m \geq n$.

We begin with a construction of a family of C -free graphs. We take two disjoint vertex sets B_0, B_1 of cardinality $m - 4$ and $n - 4$ respectively, four additional vertices a_i for $i = 0, 1, 2, 3$, and an infinite set of vertices $U = \{u_i : i \in \mathbb{Z}\}$. Our vertex set V is $\{a_i : i = 0, 1, 2, 3\} \cup B_0 \cup B_1 \cup U$. This is best thought of as the set U extended by a finite number of vertices which will be used to impose some very rigid structure on U . We now define an edge relation on V .

For $i = 0, 1, 2, 3$, let A_i be the set $\{a_i\} \cup B_{i \bmod 2}$. Using \sim to denote the edge relation, we impose the following edges on V .

- i.* $u_i \sim u_j$ iff $|j - i| \leq 2$.
- ii.* A_0, A_1, A_2, A_3 are cliques.
- iii.* B_0 and B_1 are linked to every vertex of U .
- iv.* We take $a_i \sim u_j$ iff $i \not\equiv j \pmod{4}$.

Let H be the graph with vertex set V defined in this way. Then H has the following properties.

1. For any clique C , $|C \cap U| \leq 3$.
2. For any clique C there is an i such that $C \subseteq A_i \cup U$.
3. If C is a clique and $|C| \geq 6$ then for some i we have $|C \cap B_i| \geq 2$.
4. If C is a clique of order m then for some i we have $B_i \subseteq C$.

We may now check that H is C -free. If C_0, C_1 are cliques in H of order m and n respectively, with $|C_0 \cap C_1| = 1$, then there is some i for which $B_i \subseteq C_0$, and thus $|B_i \cap C_1| \leq 1$. On the other hand we have $|B_j \cap C_1| \geq 2$ for some j , and hence $j \neq i$, that is $\{i, j\} = \{0, 1\}$. It follows that $B_j \subseteq C_1$; if $m = n$ this is clear, and if $m > n$ then we have $i = 0$ and $j = 1$, and as $|B_j| = n - 4$ it follows that $B_j \subseteq C_1$. Hence there are $i^*, j^* \in \{0, 1, 2, 3\}$, distinct modulo 2, such that $C_0 \subseteq A_{i^*} \cup U$ and $C_1 \subseteq A_{j^*} \cup U$. In particular $C_0 \cap C_1 \subseteq U$.

But $C_0 = A_{i^*} \cup \{u_{i'+1}, u_{i'+2}, u_{i'+3}\}$ with $i' \equiv i^* \pmod{4}$, and similarly $C_1 = A_{j^*} \cup \{u_{j'+1}, u_{j'+2}, u_{j'+3}\}$ and $i' \not\equiv j' \pmod{2}$. Hence the sets $\{u_{i'+1}, u_{i'+2}, u_{i'+3}\}$ and $\{u_{j'+1}, u_{j'+2}, u_{j'+3}\}$ cannot intersect in just one element. So H is C -free.

Now suppose that $f : H \rightarrow G$ is an embedding of H into a C -free graph G , that is an isomorphism of H with a subgraph of G . Let $A = \{a_0, a_1, a_2, a_3\} \cup B_0 \cup B_1$. We claim that the function f is completely determined by its restriction to the finite set $X = A \cup \{u_{-1}, u_0, u_1\}$. Proceeding inductively, and taking account of symmetry, it suffices to check that the values of $f(u_2)$ and $f(u_3)$ are determined by the restriction of f to X . This is the key ‘‘rigidity’’ property that the construction was intended to achieve.

For $i \geq 0$, let C_i be the clique $A_{i \bmod 4} \cup \{u_{i+1}, u_{i+2}, u_{i+3}\}$. Set $Q_i = C_{2i}$ and $Q'_i = C_{2i+1}$. Then the Q_i are cliques of order m , and the Q'_i are cliques of order n .

For any i , the pair (u_i, u_{i+4}) does not lie on an edge of G , as otherwise the set $\tilde{C}_i = A_{i \bmod 4} \cup \{u_i, u_{i+2}, u_{i+3}\}$ is a clique and $C_{i-3} \cup \tilde{C}_i \cong C$.

Consider a clique Q in G of order m whose intersection with $f[X]$ is $f[A_0 \cup \{u_1\}]$. One such clique is $f[Q_0]$, and we claim that it is unique. Note that $Q \cap f[Q'_{-1}]$ contains $f(u_1)$ and no other vertex except possibly $f(u_2)$. Since G is C -free, Q contains $f(u_2)$. So $Q \cap f[Q'_0]$ contains $f(u_2)$ and no other vertex except possibly $f(u_3)$. So $Q = f[Q_0]$. Thus the unordered pair $\{f(u_2), f(u_3)\}$ is determined by the restriction of f to X .

Now consider the neighbors of a_2 in $\{f(u_2), f(u_3)\}$. Certainly $f(u_3)$ is one such; and we claim that $f(u_2)$ is not. Given this, it will be clear that the ordered pair $(f(u_2), f(u_3))$ is uniquely determined by $f \upharpoonright X$.

So suppose $(f(a_2), f(u_2))$ is an edge in G , or, what amounts to the same, consider the effect of making (a_2, u_2) into an edge in H . Then $B_0 \cup \{a_2, u_0, u_1, u_2\}$ becomes a clique of order m , meeting Q'_0 in a single vertex, and the resulting graph is not C -free.

So for any embedding f of H into a C -free graph G , the restriction of f to the finite set X determines f on all of U .

Now we modify the construction of H slightly. We associate to any bit string $\epsilon = (\epsilon_0, \epsilon_1, \dots)$ of 0's and 1's the graph H_ϵ , which differs from H only by the addition of edges (u_{8i}, u_{8i+4}) when $\epsilon_i = 1$, and

by the addition of an auxiliary graph *obstructing* such an edge when $\epsilon_i = 0$: we may take as the obstructing graph C' the graph derived from $K_m \uplus K_n$ by omitting a single edge between two vertices of the clique K_n (but not the vertex shared with K_m). So when $\epsilon_i = 0$, we amalgamate C' freely with H , identifying the two vertices of C' on the deleted edge with the two vertices u_{8i}, u_{8i+4} of H .

In this way we obtain 2^{\aleph_0} different graphs H_ϵ , each of which contains a sequence (u_{2i}) which is rigid over the finite set X , and such that for two bit strings $\epsilon \neq \epsilon'$, the induced graphs on the vertices (u_{2i}) are incompatible (that is, any graph containing both must have a subgraph isomorphic with C).

It follows that there is no weakly universal countable C -free graph G . Otherwise, the graph G would contain all the H_ϵ as induced subgraphs, and then two of these graphs would involve the same embedding of the finite set X into G , and hence by rigidity would have the same embedding of the sequence (u_i) into G , resulting in a copy of C wherever $\epsilon_i \neq \epsilon'_i$. \square

The analysis that leads to the particular construction used in our next result is elaborate. See the end of the next section for a similar analysis worked out. The configuration is of an unusual type, in fact unique as far as known examples are concerned.


Proposition 3.2 *Let $C = K_5 \uplus K_5$. Then there is no universal C -free graph.*

Proof.

We follow the line of the previous argument. First we introduce a particular C -free graph H with a very rigid structure.

As our vertex set $V(H)$ we take vertices a_i, a'_i for $i \in \mathbb{Z}$. For the edge relation \sim on H we take the following.

1. $a_i \sim a_j$ iff $|j - i| = 2$;
2. $a'_i \sim a'_j$ iff $|j - i| = 1$ or 2 ;
3. $a_i \sim a'_j$ iff $|j - i| \leq 2$.



Let C_i be the clique $\{a_{i-1}, a_{i+1}, a'_{i-1}, a'_i, a'_{i+1}\}$. We wish to view this clique as “marking” the sequence a'_{i-1}, a'_i, a'_{i+1} in a certain sense, and for this purpose it is necessary to analyze the cliques of H , and more generally the cliques in any C -free graph G containing H as a subgraph and meeting H .

So let G be a C -free graph containing H . We claim that the only cliques of order 5 in G which meet H are the cliques C_i .

Let Q be such a clique. We argue first that Q contains some vertex a'_i . Otherwise, if i is maximal so that $a_i \in Q$, we would find that the induced graph on $Q \cup C_{i+1}$ is a copy of $K_5 \dagger K_5$.

So suppose that $a'_i \in Q$. We claim then that Q is the clique C_j with $j \in \{i - 1, i, i + 1\}$.

As G omits $K_5 \dagger K_5$, Q meets the cliques C_{i-1} , C_i , and C_{i+1} in at least two vertices each. Consider the pair of disjoint sets $C_{i-3} \cup C_{i-2}$ and $C_{i+2} \cup C_{i+3}$. If Q meets both of them, then as these sets are disjoint and do not contain a'_i , we find that Q intersects each one in exactly two vertices. Then consideration of the intersections of Q with C_{i-5}, C_{i-4} shows that $a'_{i-1}, a_{i-1} \in Q$; similarly, $a'_{i+1}, a_{i+1} \in Q$. So $Q = C_i$ in this case.

So we may suppose that Q is disjoint from $C_{i-3} \cup C_{i-2}$. Then $Q \cap C_{i-1} = \{a_i, a'_i\}$, and looking at $Q \cap C_i$ we see that Q also meets $\{a_{i+1}, a'_{i+1}\}$. On the other hand, the pair (a_i, a_{i+1}) does not lie on an edge, as otherwise $Q' = \{a_i, a_{i+1}, a'_{i-1}, a'_i, a'_{i+1}\}$ would be a clique of order 5 meeting C_{i-2} in a unique vertex, giving a contradiction. So $a_{i+1} \notin Q$ and Q contains $\{a_i, a'_i, a'_{i+1}\}$. It follows quickly that $Q = C_{i+1}$: easily Q must be disjoint from C_{i+4} , and then looking at the intersection with C_{i+2} , followed by C_{i+3} , the result follows.

Now with G a C -free graph containing H as a subgraph, consider the set V_0 of vertices lying in three distinct cliques of order 5, and let

G^* be the graph induced on V_0 by edges whose endpoints lie in two distinct cliques. Then the vertices $\{a'_i\}$ form a connected component of G^* and the edges induced on these vertices by G^* are those induced by H . It follows that any embedding of H into G is uniquely determined by the images of a'_0, a'_1 in G . With this rigidity result, we can conclude much as in the previous argument. The details are as follows.

For any bit string $\epsilon = (\epsilon_i : i \in \mathbb{Z})$, let H_ϵ be the extension of H with the same vertex set, and with additional edges $(a_{4i} \sim a_{4(i+1)})$ when $\epsilon_i = 1$, while for $\epsilon_i = 0$ a “contradictory configuration” is attached: take a copy of K_5 with one edge deleted, and where the endpoints of the deleted edge are to be identified with $(a_{4i}, a_{4(i+1)})$. Then the graphs H_ϵ are C -free, and there are 2^{\aleph_0} of them.

Given a weakly universal C -free graph G , we would have embeddings of all the graphs H_ϵ into G as subgraphs, and then at least two of these embeddings would have to agree on the vertices a_0, a'_0 , and hence on all of $V(H)$. But then wherever $\epsilon(i) \neq \epsilon'(i)$, we would have a new clique meeting H , contradicting our analysis above. \square

4 2-Bouquets with Universal Graphs

In this final section we will show that some fairly large 2-bouquets, taken as forbidden subgraphs, allow a countable universal graph. We only know one way to do this, namely by showing that the closure operator analyzed in [3] is locally finite. In model theoretic terms, this implies that the class of existentially complete C -free graphs has a unique countable model, which provides the strong form of universality at once.

Fortunately, it is not necessary to work out the closure operator in any detail. Some general considerations, primarily relating to the Δ -system lemma, will give us a very rough bound on the complexity of this operator. In this section we will however need to work with an explicit definition of this operator.

Our 2-bouquets $K_m + K_n$ are *solid* graphs in the following sense.

Definition 4.1 *A graph C is called solid if all of its blocks (2-connected components) are complete.*

We will next review the definition of the operator cl^* given in [3]. This is the operator whose local finiteness is in question.

First, let G be a C -free graph, C_0 a subgraph of G which is isomorphic to a subgraph of the graph C , and let $B \subseteq V(C_0)$ be a subset of its set of vertices. Let $(C_0^\infty)_B$ be the free amalgam of an infinite sequence of copies of C_0 over B ; in other words, form the disjoint union of copies C_0^i of C_0 for $i \in \mathbb{N}$, then identify the various copies of B .

We say that G is *free* over C_0 if the graph $(C_0^\infty)_B$ embeds in G over B . We say that B is a *base* for C_0 in G over the vertex $v \in B$ if C_0 is free over B in G , and B is minimal subject to these requirements: that is, C_0 is not free in G over any set B' properly contained in B for which $v \in B'$.

With this notation, we define the following ‘‘closure operators’’. Let G be a C -free graph, and C_0 a subgraph of C with a distinguished basepoint v_0 . For v a vertex of G , let $\text{cl}_{C_0}(v)$ be the union of all bases B over v for subgraphs of G isomorphic to C_0 , with v corresponding to the basepoint of C_0 . For A a set of vertices in G , let $\text{cl}_{C_0}(A) = \bigcup_{v \in A} \text{cl}_{C_0}(v)$. Similarly, if \mathcal{F} is a family of subgraphs of C , then $\text{cl}_{\mathcal{F}}$ is the union of cl_{C_0} as C_0 ranges over \mathcal{F} .

These are not true closure operators, as they are not idempotent: we do not have $\text{cl}(\text{cl}(A)) = \text{cl}(A)$. But these partial closure operators generate the closure operators that interest us. For any of these partial closure operators cl , let cl^* denote the corresponding closure operator. In detail, setting $\text{cl}^0 = \text{cl}$ and $\text{cl}^{n+1} = \text{cl} \circ \text{cl}^n$, we let $\text{cl}^* = \bigcup_n \text{cl}^n$.

Fact 4.2 ([3])

1. For G a C -free graph, \mathcal{F} a family of subgraphs of C , and $v \in G$, the set $\text{cl}_{\mathcal{F}}(v)$ is finite.
2. For \mathcal{F} the family of all subgraphs of C , if $\text{cl}_{\mathcal{F}}^*(v)$ is finite for all vertices v in all C -free graphs G , then there is a universal C -free graph (and even a canonical one).

The second clause above is not sufficiently precise for our purposes, as we will need to make computations involving $\text{cl}_{\mathcal{F}}$, and the family \mathcal{F} defined here is broader, and the associated closure operator $\text{cl}_{\mathcal{F}}$ more complex, than either needs to be. So we will now sketch the definition of a narrower class \mathcal{F}_0 associated with a 2-bouquet C . The main point of this definition is that it yields a more tightly controlled partial closure operator $\text{cl}_{\mathcal{F}_0}$ with the property that after iteration it yields the full closure operator: $\text{cl}_{\mathcal{F}_0}^* = \text{cl}_{\mathcal{F}}^*$. The definition of \mathcal{F}_0 is based on the representation of a graph as a tree of blocks (2-connected components with linking edges). We will now review this.

Let C be a graph. We define an equivalence relation on the edges of C as follows. For $e, e' \in E(C)$, write $C(e, e')$ to mean that e, e' belong to a cycle. Let C^* be the equivalence relation generated by C . A block is the graph associated to a single equivalence class for the relation C^* by taking the edges in this class together with their vertices. The blocks of C partition the edges and are either 2-connected or trivial (a single edge). Two blocks intersect in at most one vertex. Let V_0 be the set of blocks in C , and V_1 the set of vertices lying in more than one block. We put a bipartite graph structure on $V_0 \cup V_1$ by linking $v_0 \in V_0$ with $v_1 \in V_1$ if $v_1 \in v_0$. This gives a forest; if C is connected, it gives a tree T . For what follows, we take C to be connected and

choose a root in V_1 for T ; we treat T as a partial order with the root at the bottom.

For $t \in T$ we let T^t be $\{s \in T : s > t\}$, and let C^t be the corresponding part of C . We set

$$\mathcal{F}_0 = \{(H, t) : t \in V_1, t \in H, H \setminus \{t\} \text{ is a connected component of } C^t\}$$

Here the pair (H, t) represents a graph H with a basepoint t .

Fact 4.3 [3] *Let \mathcal{C} be a finite set of finite solid graphs. Let \mathcal{F} be the set of all finite subgraphs of graphs in \mathcal{C} , and let \mathcal{F}_0 be the set of pairs (C^t, t) introduced above, with C varying over \mathcal{C} . Then for $G \in \mathcal{E}_{\mathcal{C}}$ and $A \subseteq V(G)$, we have $\text{cl}_{\mathcal{F}_0}^*(A) = \text{cl}_{\mathcal{F}}^*(A)$.*

We now specialize to the case at hand, $\mathcal{C} = \{C\}$ with $C = K_m \uplus K_n$ and $m \geq n$. The relevant closure operator is $\text{cl}_{(H,a)}$ with $H \cong K_m$ and $a \in H$. We will work inside an existentially complete C -free graph G : the one significant graph theoretic consequence of this assumption is that for any subgraph $H \cong K_m$ in G , and any subset $A \subseteq H$, if H is not free over A in G then there is a clique Q of order 5 with $|Q \cap A| = 1$ (in other words, there is a reason why H is not free over A). Any C -free graph embeds as a subgraph in a C -free graph with this additional property, and if one does not wish to invoke existential completeness, it suffices to check this fact directly.

After these lengthy preliminaries, we can proceed. Our objective is to show that in certain cases the closure operator $\text{cl}_{\mathcal{F}_0}$ defined above is locally finite. If this fails, then the infinite iteration which defines cl^* really is infinite, in other words $\text{cl}^n \neq \text{cl}^{n+1}$ for all n . Using the compactness theorem of logic or König's theorem, it follows that in some C -free graph there is an infinite sequence of vertices a_n such that $a_{n+1} \in \text{cl}(a_n) \setminus \text{cl}(a_0, \dots, a_{n-1})$ for all n . One thing to keep in mind is that according to our definitions, this forces the pair $\{a_n, a_{n+1}\}$ to be inside some m -clique, and in view of the hypothesis that the ambient graph G be C -free, it should be clear that this already introduces a certain amount of tension into the situation. The constructions of the previous section give a good indication of what such sequences actually look like in practice. What we need to show below is that for $n \leq 4$, or for $n = 5$ and $m > 5$, no such sequence can exist. This is done by a fairly close analysis, after some preliminary considerations involving Δ -systems.

Proposition 4.4 *Let $C = K_m \uplus K_n$ with $m \geq n$, and $n \leq 4$. Then $\text{cl}_{\mathcal{F}_0}^*$ is locally finite, and hence there is a universal countable C -free graph.*

Proof. We write cl^* for $\text{cl}_{\mathcal{F}_0}^*$, and cl for $\text{cl}_{\mathcal{F}_0}$, and we begin by running through the analysis sketched above.

Suppose the proposition fails. Then as the partial operator cl is locally finite, for each i there must be some C -free graph G and some vertex $a \in V(G)$ for which $\text{cl}^{i+1}(a) \neq \text{cl}^i(a)$. What we actually want to work with is an infinite sequence a_i of vertices in a single C -free graph G , such that $a_{i+1} \in \text{cl}(a_i)$, $a_{i+1} \notin \text{cl}(a_j)$ for $j < i$. By the Compactness Theorem of logic (or a brute force application of König's Lemma) it suffices to find arbitrarily long sequences (a_i) of this type. One finds them by reverse induction. Begin with a vertex a for which $\text{cl}^n(a) \neq \text{cl}^{n-1}(a)$, with n large, and choose a_i inductively, starting with a_n and working downward, as follows. First, fix $a_n \in \text{cl}^n(a) \setminus \text{cl}^{n-1}(a)$. Then, given a_i , one has $a_i \in \text{cl}(b)$ for some $b \in \text{cl}^{i-1}(a)$, and one takes $a_{i-1} = b$; it is easy to verify that the desired conditions are met.

Now we use the explicit definition of cl given above. We will write $\text{cl}(v, H)$ for the portion of $\text{cl}_{(K_m, a)}(v)$ associated with a particular m -clique H containing v : in other words, the union of the bases for H over v in G . We have $a_i \in \text{cl}(a_{i-1}, H_i)$ for some m -clique H_i , and thus there is a subset $B_i \subseteq H_i$ with $a_{i-1} \in B_i$, H_i free over B_i , such that B_i is minimal subject to these conditions, and with $a_i \in B_i$. We study the whole configuration that results.

In view of the freeness of H_i over B_i , we may suppose that $H_i \cap H_j \subseteq B_i \cap B_j$ for all i, j . In particular for $i < j - 1$, the a_j are not in B_i , by the choice of the sequence, and hence $a_j \notin H_i$.

Now we apply a Δ -system argument to the sequence (H_i) . This sequence is an infinite sequence of sets all of fixed size. Hence there is an infinite subsequence $(H_i : i \in I)$ which forms a " Δ -system with heart B ," that is we have $H_i \cap H_j = B$ for all $i, j \in I$ distinct, with B a fixed set. This is the starting point for our analysis. The main case division is between the cases $B \neq \emptyset$ and $B = \emptyset$. The case in which B is empty is one that only arises when m is small, and in fact is illustrated by the construction in the previous section corresponding to $m = n = 5$. Given that example, it is not surprising that that possibility calls for some close analysis here.

Note that the Δ -system property shows that the cliques H_i are free over B for $i \in I$.

The case $n = 2$ is straightforward, and is also included in a case treated at the end of [3] by more direct methods, so we will assume here that $n \geq 3$.

The case $n = 3$

For $n = 3$, the situation is so simple that we do not even need a Δ -system. We can simply consider $H_i \cap H_j$ with $i = j - 2$. Then $a_{j-1}, a_j \notin H_i$, so $|H_j \setminus H_i| \geq 2$, and if $H_i \cap H_j \neq \emptyset$ then we have an embedding of $K_m \dagger K_3$ into G with H_i representing K_m and with a_{j-1}, a_j representing the additional vertices of K_3 . This is a contra-

diction, so $H_{j-2} \cap H_j = \emptyset$.

On the other hand, $a_{j-1} \in H_{j-1} \cap H_j$, so this intersection is definitely nonempty, and hence by the same argument $|H_{j-1} \cap H_j| = m-1$, and similarly $|H_{j-2} \cap H_{j-1}| = m-1$, so $|H_{j-2} \cap H_j| \geq m-2$. Since $m \geq n = 3$ we see that $H_{j-2} \cap H_j \neq \emptyset$ after all, and we have the desired contradiction. This, in miniature, is how the Δ -system argument will permit us to argue in less straightforward cases.

The case $n = 4$

With $n = 4$, suppose first that the ‘‘heart’’ B of the Δ -system is nonempty. Taking $i, j \in I$ with $i < j-1$ we see that $|B| = |H_i \cap H_j| \leq m-2$. If $|B| \leq m-3$ then $H_i \dagger K_4$ embeds into $H_i \cup H_j$, a contradiction. So $|B| = m-2$ in this case, and $B = H_i \setminus \{a_{i-1}, a_i\}$.

We have $B \subseteq H_i$ for $i \in I$, and we will show that $B \subseteq H_i$ for all i . Supposing the contrary, we have some choice of i for which $B \subseteq H_i$ and $B \not\subseteq H_{i-1}$. On the other hand, $H_{i-1} \cap H_i \neq \emptyset$, and as G is C -free this implies $|H_{i-1} \cap H_i| \geq m-2$. Hence $|H_{i-1} \cap B| \geq m-3$ and as $m \geq 4$ it follows that H_{i-1} meets B . Taking $j \in I$ with $j > i$, we have $H_{i-1} \cap H_j \neq \emptyset$ and $a_{j-1}, a_j \notin H_i$. Hence $H_{i-1} \cap H_j \subseteq B$, and as G is C -free we must have $|H_{i-1} \cap H_j| \geq m-2$. Hence $B \subseteq H_{i-1}$ after all.

Using our Δ -system, we have more or less pinned down the whole configuration. At this point since $B \subseteq B_i$ for each i , and $a_{i-1}, a_i \in B_i$, we have $B_i = H_i$.

Take $i \in I$. By the minimality of the base B_i , H_i is not free over $H_i \setminus \{a_i\}$. Hence by our assumption on G we have the following ‘‘obstruction’’.

There is some $Q \cong K_4$ in G such that $|Q \cap (H_i \setminus \{a_i\})| = 1$

As G is C -free we find $|H_i \cap Q| = 2$, and $a_i \in Q$. So $|Q \cap B| \leq 1$, and as H_i is free over B , if we had $|Q \cap B| = 1$ we would get an embedding of $Q \dagger H_i$ into G for some i , a contradiction. This is a key Δ -system argument.

Our conclusion is that $Q \cap B = \emptyset$ and hence $Q \cap H_i = \{a_{i-1}, a_i\}$. Since $H_{i+1} = B \cup \{a_i, a_{i+1}\}$ it follows that $Q \cap H_{i+1} \subseteq \{a_i, a_{i+1}\}$ and hence $Q \cap H_{i+1} = \{a_i, a_{i+1}\}$. Continuing inductively, $a_j \in Q$ for infinitely many j , a contradiction. Thus

B is empty

Now choose i and j so that $H_i \cap H_j = \emptyset$, $i < j$, and $j-i$ is *minimal*. Evidently $j \geq i+2$. Let $A = H_{j-1} \setminus \{a_{j-2}, a_{j-1}\}$.

If $j > i+2$ then $H_i \cap H_{j-1} = H_i \cap A$, and as G is C -free we have $|H_i \cap H_{j-1}| \geq m-2$, hence $A \subseteq H_i$. As $H_i \cap H_j = \emptyset$ we have $H_{j-1} \cap H_j \subseteq \{a_{j-1}, a_j\}$; but $a_j \notin H_{j-1}$, so $|H_{j-1} \cap H_j| = 1$. As G is C -free this is a contradiction. Thus $j = i+2$.

Suppose that $m > 4$. Then as $|H_i \cap H_{i+1}| \geq m - 2$ and $|H_{i+1} \cap H_{i+2}| \geq m - 2$, we have $|H_i \cap H_{i+2}| \geq m - 4$, so $H_i \cap H_{i+2} \neq \emptyset$, a contradiction. So we have $m = 4$ and $j = i + 2$. Furthermore the preceding argument gives $|H_i \cap H_{i+1}| = |H_{i+1} \cap H_{i+2}| = 2$.

Now H_{i+1} is not free over $H_i \cap H_{i+1}$, and thus there is a clique $Q \cong K_4$ in G so that $|Q \cap (H_i \cap H_{i+1})| = 1$. Accordingly Q meets $H_i \setminus H_{i+1}$ and $H_{i+1} \setminus H_i$; let u, v belong to $Q \cap (H_i \setminus H_{i+1})$ and $Q \cap (H_{i+1} \setminus H_i)$, respectively. Set $Q' = \{u, v\} \cup (H_i \cap H_{i+1})$. Then $Q' \cong K_4$ and $Q' \cap H_{i+2} = \{v\}$, so $Q' \cup H_{i+2} \cong C$, a contradiction.

This completes the proof. \square

We note that the penultimate line of the foregoing proof is particularly important in terms of eliminating the more plausible sort of construction, similar to the constructions which are actually seen in §3.

Now we can carry out the same line of argument at the next level, with $n = 5$. Our analysis then leads more clearly in the direction of the examples that actually do exist, given in the previous section. The analysis becomes more elaborate. Since all the ideas needed are visible in the proof of the previous argument, we will assume that this style of argument is familiar at this point.

Proposition 4.5 *Let $C = K_m + K_5$ with $m \geq 6$. Then $\text{cl}_{\mathcal{F}_0}^*$ is locally finite, and hence there is a universal countable C -free graph.*

Proof. Write cl^* for $\text{cl}_{\mathcal{F}_0}^*$. Supposing this is not locally finite, work in an existentially complete C -free graph G , and extract sequences a_i, H_i, B_i witnessing the failure of local finiteness, as in the previous case. Again, apply the Δ -system lemma to get a Δ -system $(H_i : i \in I)$ with heart B . As G is C -free, one has one of the following: $|B| = m - 3$, $|B| = m - 2$, or B is empty.

While the first two cases are similar, each one has to be considered on its own. In some sense the third case, though not requiring much analysis, is really the main one. In particular the distinction between $m = 5$ and $m > 5$ becomes visible in that case.

(Case I) The case $|B| = m - 3$

If $|B| = m - 3$ one shows easily that $B \subseteq H_i$ for all i , and not just for $i \in I$.

A critical point is that

$$(1) \quad \bigcup_i (H_i \setminus B) \text{ contains no clique } Q_0 \text{ of order 4.}$$

Otherwise, one takes i with $H_i \cap Q_0 = \emptyset$, and one extends Q_0 by a point of B to get a contradiction.

This is generally applied in the following form: for any clique Q of order 5 in G , the intersection $Q \cap \bigcup_i (H_i \setminus B)$ has order at most 3. Since $|Q \cap H_i| \geq 2$ whenever the intersection is nontrivial, this restricts possibilities substantially.

We will show that $a_i \notin H_j$ for $i \neq j-1, j$. This is immediate if $j < i$. If $i \leq j-2$ but $a_i \in H_j$, one first maximizes i with respect to these conditions. Then consider $B' \subseteq B_j$ minimal so that $a_i \in B'$ and H_j is free over B' . By definition the elements of B' belong to $\text{cl}(a_i)$ and therefore by the choice of the sequence (a_n) we have $a_j \notin B'$, and hence also $a_{j-1} \notin B'$. As $a_i \notin B$, we find $B' \setminus \{a_i\} \subseteq B \subseteq H_i$. So $B' \subseteq H_i$. But then H_i is free over B' , since H_j is. But also $B \subseteq B_i$ and hence $B' \subseteq B_i$, so by minimality of B_i we have $a_{i+1} \in B'$, and hence $i+1 \leq j-2$ as well, which violates the choice of i .

Let $H_i = B \cup \{a_{i-1}, a_i, u_i\}$. It follows that $u_i \neq a_j$ for all i, j . The possibility that $u_i = u_j$ for some $i \neq j$ (notably, $j = i+2$) is one of the main issues, and brings to mind the first construction in the previous section.

At this point the analysis divides into two short-lived subcases.

(Case I.1)

For some i , there is $Q \cong K_5$ such that $Q \cap H_i = \{a_{i-1}, a_i\}$.

Fix such a Q and i throughout this part of the analysis.

(2) $u_{i-1} = u_{i+1}$

Take $Q \cong K_5$ with $Q \cap H_i = \{a_{i-1}, a_i\}$. Then $Q \cap B = \emptyset$. Applying (1) to Q , It follows quickly that $Q \cap \bigcup_i H_i = \{u_{i-1}, a_{i-1}, a_i\}$, and in particular $Q \cap H_{i+1} = \{u_{i-1}, a_i\}$, from which claim (2) follows.

Let $u = u_{i-1} = u_{i+1}$. One may check that $u \in H_j$ only for $j = i \pm 1$ (using Q). Now to the main step:

(3) There is $Q' \cong K_5$ with $Q' \cap H_{i+1} = \{a_i, a_{i+1}\}$

As H_{i+2} is not free over $\{a_{i+1}\} \cup B$ and G is existentially complete, we have some $Q' \cong K_5$ with $Q' \cap |\{a_{i+1}\} \cup B| = 1$, hence $a_{i+1} \in Q'$ and $Q' \cap B = \emptyset$. Recall also that $|Q' \cap \bigcup_i (H_i \setminus B)| \leq 3$.

As $u \neq u_{i+3}$ we cannot have $Q' \cap H_{i+2} = \{a_{i+1}, a_{i+2}\}$, by the analysis which produced claim (2). So we have $u_{i+2} \in Q'$.

If $Q' \cap H_{i+2} = \{a_{i+1}, u_{i+2}, a_{i+2}\}$ then we find $|Q \cap H_{i+1}| = 1$, a contradiction. So $Q' \cap H_{i+2} = \{a_{i+1}, u_{i+2}\}$. As $u_{i-1} = u_{i+1}$ one checks easily that Q' satisfies claim (3).

To complete the analysis of Case I.1, observe that we can repeat the proof of (3) to get some $Q'' \cong K_5$ with $Q'' \cap H_{i+2} = \{a_{i+1}, a_{i+2}\}$, and then as in (2) we have $u = u_{i+3}$, a contradiction.

(Case I.2) There is no i and $Q \cong K_5$ such that $Q \cap H_i = \{a_{i-1}, a_i\}$.

(4) There is no i and $Q \cong K_5$ such that $Q \cap H_i = H_i \setminus B$

Suppose the contrary. One finds easily that $u_{i-1} = u_i = u_{i+1}$. As H_{i-1} is not free over $\{a_{i-2}\} \cup B$, one can find $Q' \cong K_5$ with $|Q' \cap (\{a_{i-2}\} \cup B)| = 1$. Hence $a_{i-2} \in Q'$ and $Q' \cap B = \emptyset$. Using the case assumption, one finds $u_{i-1} \in Q'$. Hence $Q' \cap \{a_i, a_{i+1}\} \neq \emptyset$ and we have found three elements in $Q' \cap \bigcup_i H_i$. So $Q' \cap H_{i-2} \subseteq \{a_{i-2}, u_{i-1}\}$, forcing $u_{i-1} = u_{i-2}$. But then $|Q' \cap H_{i-2}| = 1$, a contradiction. This proves (4).

Now choose some i so that $u_{i-1} \neq u_i$. As H_i is not free over $\{a_{i-1}\} \cup B$, we have some $Q \cong K_5$ with $|Q \cap (\{a_{i-1}\} \cup B)| = 1$, and at this point this forces $Q \cap H_i = \{a_{i-1}, u_i\}$ and $Q \cap H_{i-1} = \{a_{i-1}, u_{i-1}\}$. Thus $Q \cap \bigcup_j H_j = \{u_{i-1}, a_{i-1}, u_i\}$ and looking at $Q \cap H_{i+1}$ it follows that $u_i \neq u_{i+1}$. So we can choose $Q' \cong K_5$ with $Q' \cap \bigcup_j H_j = \{u_i, a_i, u_{i+1}\}$. Let $H = \{u_i, a_i, u_{i+1}\} \cup B$; then $H \cong K_m$ and $H \cap Q = \{u_i\}$, so $H \cup Q \cong K_m \dagger K_5$, a contradiction.

(Case II) The case $|B| = m - 2$

For any j such that $H_j \cap B \neq \emptyset$, we have $|H_j \cap B| \geq m - 3$, as otherwise we can find some H_i containing B such that $|H_i \cap H_j| \leq m - 4$, and so embed C in G . It then follows that $H_{j \pm 1} \cap B \neq \emptyset$, and hence by induction we have $|H_j \cap B| \geq m - 3$ for all j .

Fix $i \in I$. Then $H_i = B \cup \{a_{i-1}, a_i\}$.

As H_i is not free over $\{a_{i-1}\} \cup B$ there is $Q \cong K_5$ such that $|Q \cap (\{a_{i-1}\} \cup B)| = 1$, and hence $Q \cap H_i = \{a_{i-1}, a_i\}$. Fix vertices u, v in $Q \cap (H_{i-1} \setminus \{a_{i-1}\})$ and $Q \cap (H_{i+1} \setminus \{a_i\})$, respectively.

Now $|H_{i \pm 1} \cap B| \geq m - 3$ and hence $H_{i-1} \cap B \cap H_{i+1} \neq \emptyset$. Take $w \in H_{i-1} \cap B \cap H_{i+1}$. Set $Q' = \{w\} \cup [Q \cap (H_{i-1} \cup H_{i+1})]$. Then Q' is a clique, and $|Q' \cap B| = 1$, so $|Q'| \leq 4$ and $|Q \cap (H_{i-1} \cup H_{i+1})| \leq 3$. It follows that $v = u$ or $v = a_{i-1}$. But $v \neq a_{i-1}$ since otherwise $a_{i-1} \in H_{i+1}$ and H_{i+1} is free over some $B' \subseteq H_{i+1}$ with $B' \subseteq H_{i+1} \setminus \{a_i, a_{i-1}\} \subseteq H_i$, which contradicts the minimality of B_i . So $v = u$. By a similar argument $v \neq a_{i-2}$.

At this point we have enough information to make a close analysis of the structure of $H^* = H_{i-1} \cup H_i \cup H_{i+1}$. As we have seen already, exploiting the vertex w , there is no clique of order 4 in $H^* \setminus B$. On the other hand, $u \in H_{i-1} \cap H_{i+1}$ and hence $a_{i-2} \not\sim a_i$, $a_{i-1} \not\sim a_{i+1}$ (using “ \sim ” for the edge relation).

We can take $Q' \cong K_5$ with $|Q' \cap (H_{i+1} \setminus \{a_{i+1}\})| = 1$, and then $a_{i+1} \in Q'$, $Q' \cap B = \emptyset$, and $|Q' \cap \{a_i, u\}| = 1$.

As $a_{i+1} \in Q'$, it follows that $a_{i-1} \notin Q'$, and hence $Q' \cap H_i = \emptyset$.

Now $u \notin H_{i-2}$, as otherwise $|Q \cap (H_{i-2} \cup H_{i-1})| \geq 4$ and hence $(H_{i-2} \cap H_i) \cap B = \emptyset$, which is false.

It then follows that $|Q' \cap (H_{i-2} \cup H_{i+1})| \geq 4$. Hence $H_{i-2} \cap H_{i+1} \cap B = \emptyset$, but $|B| \geq 4$ and $|B \setminus H_j| \leq 1$ for all j , so this is impossible.

(Case III) The case in which B is empty

We choose a pair i, j with $i < j$ such that $H_i \cap H_j = \emptyset$ and $j - i$ is minimal. Then $|H_i \cap H_{i+1}|, |H_{i+1} \cap H_j| \geq m - 3$, and it follows that $n = 6$.

If $j > i + 2$, consider $A_k = H_i \cap H_{i+k}$ and $A'_k = H_j \cap H_{i+k}$ for $k = 1, 2$. Then (A_k, A'_k) is a partition of H_{i+k} of type $(3, 3)$. As $A'_1 \cup A'_2 \subseteq H_j \setminus \{a_j\}$, there is a vertex $u \in A'_1 \cap A'_2$. If $A_1 \neq A_2$ we can find a clique Q of order 5 contained in $A_1 \cup A_2 \cup \{u\}$ such that $H_j \cap Q' = \{u\}$, a contradiction.

So $A_1 = A_2$, and therefore $A'_1 \neq A'_2$, and we can run the same argument with H_i in place of H_j . Hence

$$j = i + 2$$

Next we claim

$$H_k \cap H_{k+2} = \emptyset \text{ for all } k$$

Proceeding inductively, it is enough to consider $H_{i-1} \cap H_{i+1}$. Suppose this intersection is nonempty, and take k minimal so that $H_k \cap H_{i+1} \neq \emptyset$; if no such k exists, we can form a Δ -system with nonempty heart, reducing to a case already treated. Then easily H_k is partitioned by H_i, H_{i+2} and hence H_{k-1} meets $H_i \cup H_{i+2}$, while $a_{i-1}, a_{i+2} \notin H_{k-1}$, forcing $|H_{k-1} \cap H_i| \leq 2, |H_{k-1} \cap H_{i+2}| \leq 2$. These intersections are too small to be nontrivial, so $H_{k-1} \cap H_i = H_{k-1} \cap H_{i+2} = \emptyset$, a contradiction.

It now follows easily that all H_i, H_j are disjoint for $|j - i| \geq 2$.

Finally, since H_i is not free over $H_i \cap H_{i-1}$, there is $Q \cong K_5$ meeting $H_i \cap H_{i-1}$ in one vertex. Taking $u \in Q \cap (H_i \setminus H_{i-1})$ and $v \in Q \cap (H_{i-1} \setminus H_i)$, we let $Q' = \{u, v\} \cup (H_i \cap H_{i-1})$ and we find $H_{i-2} \cup Q' \cong C$, a contradiction. This ends the analysis. \square

The same analysis, in the case $C = K_5 \uplus K_5$, shows that we must have a Δ -system with $B = \emptyset$, embedded in a configuration similar to the one given in the previous section. The most significant difference between the cases $m = 5$ and $m > 5$ occurs precisely at the end of the analysis, in the last paragraph of the proof above.

References

- [1] G. Cherlin and P. Komjáth. There is no universal countable pentagon-free graph, *J. Graph Theory*, **18** (1994), 337–341.
- [2] G. Cherlin and S. Shelah. The tree conjecture for universal graphs, *J. Comb. Theory Ser. B*, **97** (2007), 293–333.
- [3] G. Cherlin, S. Shelah, and N. Shi. Universal graphs with forbidden subgraphs and algebraic closure, *Advances in Applied Mathematics* **22** (1999), 454–491.
- [4] G. Cherlin and N. Shi. Graphs omitting a finite set of cycles, *J. Graph Theory*, **21** (1996), 351–355.
- [5] G. Cherlin and N. Shi. Graphs omitting sums of complete graphs, *J. Graph Theory*, **24** (1997), 237–247.
- [6] G. Cherlin and N. Shi. Forbidden subgraphs and forbidden substructures. *J. Symbolic Logic*, **66** (2001), 1342–1352.
- [7] G. Cherlin, N. Shi, and L. Tallgren. Graphs omitting a bushy tree. *J. Graph Theory*, **26** (1997), 203–210.
- [8] Z. Füredi and P. Komjáth. On the existence of countable universal graphs. *J. Graph Theory*, **25** (1997), 53–58.
- [9] Z. Füredi and P. Komjáth. Nonexistence of universal graphs without some trees, *Combinatorica*, **17** (1997), 163–171.
- [10] M. Goldstern and M. Kojman. Universal arrow free graphs, *Acta Math. Hungary* **73** (1996), 319–326.
- [11] P. Komjáth. Some remarks on universal graphs. *Discrete Math.* **199** (1999), 259–265.
- [12] P. Komjáth, A. Mekler and J. Pach. Some universal graphs, *Israel J. Math*, **64** (1988), 158–168.
- [13] P. Komjáth and J. Pach. Universal graphs without large bipartite subgraphs. *Mathematika*, **31** (1984), 282–290.
- [14] P. Komjáth and J. Pach. Universal elements and the complexity of certain classes of infinite graphs. *Discrete Math.* **95** (1991), 255–270.
- [15] J. Pach. A problem of Ulam on planar graphs. *Eur. J. Comb.* **2** (1981), 357–361.
- [16] R. Rado. Universal graphs and universal functions. *Acta Arith.* **9** (1964), 331–340.