

Two problems on homogeneous structures, revisited

Gregory Cherlin

ABSTRACT. We take up Peter Cameron’s problem of the classification of countably infinite graphs which are homogeneous as metric spaces in the graph metric [Cam98]. We give an explicit catalog of the known examples, together with results supporting the conjecture that the catalog may be complete, or nearly so.

We begin in Part I with a presentation of Fraïssé’s theory of amalgamation classes and the classification of homogeneous structures, with emphasis on the case of homogeneous metric spaces, from the discovery of the Urysohn space to the connection with topological dynamics developed in [KPT05]. We then turn to a discussion of the known metrically homogeneous graphs in Part II. This includes a 5-parameter family of homogeneous metric spaces whose connections with topological dynamics remain to be worked out. In the case of diameter 4, we find a variety of examples buried in the tables at the end of [Che98], which we decode and correlate with our catalog.

In the final Part we revisit an old chestnut from the theory of homogeneous structures, namely the problem of approximating the generic triangle free graph by finite graphs. Little is known about this, but we rephrase the problem more explicitly in terms of finite geometries. In that form it leads to questions that seem appropriate for design theorists, as well as some questions that involve structures small enough to be explored computationally. We also show, following a suggestion of Peter Cameron (1996), that while strongly regular graphs provide some interesting examples, one must look beyond this class in general for the desired approximations.

1. Introduction

The core of the present article is a presentation of the known *metrically homogeneous graphs*: these are the graphs which, when viewed as metric spaces in the graph metric, are homogeneous metric spaces.

1.1. Homogeneous Metric Spaces, Fraïssé Theory, and Classification.

A metric space M is said to be *homogeneous* if every isometry between finite subsets of M is induced by an isometry taking M onto itself. An interesting and early

2010 *Mathematics Subject Classification.* Primary 03C10; Secondary 03C13, 03C15, 05B25, 20B22, 20B27.

Key words and phrases. homogeneous structure, permutation group, finite model property, model theory, amalgamation, extreme amenability, Ramsey theory.

Supported by NSF Grant DMS-0600940.

example is the Urysohn space U [Ury25, Ury27] found in the summer of 1924, the last product of Urysohn’s short but intensely productive life. While the problem of Fréchet that prompted this construction concerned universality rather than homogeneity, Urysohn took particular notice of this homogeneity property in his initial letter to Hausdorff [Huš08], a point repeated in much the same terms in the posthumous announcement [Ury25]. We will discuss this in §2.

From the point of view of Fraïssé’s later theory of *amalgamation classes* [Fra54], the essential point in Urysohn’s construction is that finite metric spaces can be amalgamated: if M_1, M_2 are finite metric spaces whose metrics agree on their common part $M_0 = M_1 \cap M_2$, then there is a metric on $M_1 \cup M_2$ extending the given metrics; and more particularly, the same applies if we limit ourselves to metric spaces with a rational valued metric.

Fraïssé’s theory facilitates the construction of infinite homogeneous structures of all sorts, which are then universal in various categories, and the theory is often used to that effect. This gives a construction of Rado’s universal graph [Rad64], generalized by Henson to produce universal K_n -free graphs [Hen71], and uncountably many quite similar homogeneous directed graphs [Hen72]. A variant of the same construction also yields uncountably many homogeneous nilpotent groups and commutative rings [CSW93]. More subtly, Fraïssé’s of amalgamation classes can be used to *classify* homogeneous structures of various types: homogeneous graphs [LW80], homogeneous directed graphs [Sch79, Lac84, Che88, Che98], the finite and the imprimitive infinite graphs with two edge colors [Che99], the homogeneous partial orders with a vertex coloring by countably many colors [TrT08], and even homogeneous permutations [Cam03].

There is a remarkable 3-way connection involving the theory of amalgamation classes, structural Ramsey theory, and topological dynamics, developed in [KPT05]. In this setting the Urysohn space appears as one of the natural examples, but more familiar combinatorial structures come in on an equal footing. The Fraïssé theory, and its connection with topological dynamics, is the subject of §3.

In §4 we conclude Part I with a discussion of the use of Fraïssé’s theory to obtain classifications of all the homogeneous structures in some limited classes. Noteworthy here is the classification by Lachlan and Woodrow of the homogeneous graphs [LW80], which plays an important role in Part II.

1.2. Metrically Homogeneous Graphs. Part II is devoted to a classification problem for a particular class of homogeneous metric spaces singled out by Peter Cameron, an ambitious generalization of the case of homogeneous graphs treated by Lachlan and Woodrow. Every connected graph is a metric space in the graph metric, and Peter Cameron raised the question of the classification of the graphs for which the associated metric spaces are homogeneous [Cam98]. Such graphs are referred to as *metrically homogeneous* or *distance homogeneous*. This condition is much stronger than the condition of distance transitivity which is familiar in finite graph theory; the complete classification of the finite distance transitive graphs is much advanced and actively pursued. Cameron raised this issue in the context of his “census” of the very rich variety of countably infinite distance transitive graphs, and we develop Cameron’s “census” into what may reasonably be considered a “catalog.” The most striking feature of our catalog is a 4-parameter family of metrically homogeneous graphs determined by constraints on triangles.

For all of the known metrically homogeneous graphs, the next order of business would be to explore the Ramsey theoretic properties of their associated metric spaces, as well as the behavior of their automorphism groups in the setting of topological dynamics, in the spirit of §3.

We believe that our catalog may be complete, though we are very far from having a proof of that, or a clear strategy for one. But we will show that as far as certain natural classes of extreme examples are concerned, this catalog is complete. The catalog, and a statement of the main results about it, will be found in §5.

There are, a priori, three kinds of metrically homogeneous graphs which are sufficiently exceptional to merit separate investigation.

First, there are those for which the graph induced on the neighbors of a given vertex is exceptional. Here the distinction between the generic and exceptional cases is furnished, very conveniently, by the Lachlan/Woodrow classification. When a graph is metrically homogeneous, the induced graph on the neighborhood of any vertex is homogeneous as a graph, and its isomorphism type is independent of the vertex selected as base point. So this induced graph provides a convenient invariant which falls under the Lachlan/Woodrow classification. Those induced graphs which can occur within graphs in our known 4-parameter family are treated as non-exceptional, while the others are considered as exceptional: according to the Lachlan/Woodrow classification, these are the ones which do not contain an infinite independent set, and the imprimitive ones.

Next, there are the imprimitive ones, which carry a nontrivial equivalence relation invariant under the automorphism group.

A third class of metrically homogeneous graphs meriting separate consideration may be described in terms of the *minimal constraints* on the graph, which are the minimal finite integer-valued metric spaces which cannot be embedded isometrically into the graph with the graph metric. There is a profusion of metrically homogeneous graphs in which the minimal constraints all have order at most 3 (i.e., exactly 3, together with a possible bound on the diameter), and one would expect their explicit classification to be an important step in the construction of an appropriate catalog. Our catalog is complete in this third sense—it contains all the metrically homogeneous graphs whose minimal constraints have order at most 3. This last point proved elusive.

In the construction of our catalog, we began by determining those in the first class, that is the metrically homogeneous graphs whose neighborhood graphs are exceptional in the sense indicated above. These turn out to be of familiar types, namely the homogeneous graphs in the ordinary sense [LW80], the finite ones given in [Cam80], and the natural completion of the class of tree-like graphs considered in [Mph82].

Turning to the imprimitive case, we find that this cannot really be separated from the “generic” case. With few exceptions, a primitive metrically homogeneous graph is either bipartite or is of “antipodal” type, which in our context comes down to the following, after some analysis: the graph has finite diameter δ , and an “antipodality” relation $d(x, y) = \delta$ defines an involutory automorphism $y = \alpha(x)$ of the graph. In the bipartite case there is a reduction to an associated graph on each half of the bipartition. The antipodal case does not have a neat reduction: we will discuss the usual “folding” operation applied in such cases, and show that it does not preserve metric homogeneity. Our catalog predicts that the bipartite and

antipodal graphs occur largely within the generic family by specializing some of the numerical parameters to extreme values, with the proviso that the treatment of side conditions of Henson's type varies slightly in the antipodal case. We investigated some special classes of bipartite graphs using the standard reduction and found some examples that looked curious at the time but find their natural place in the present catalog as metrically homogeneous graphs of generic type.

The class of graphs determined by constraints of order 3 appears to be the key to the classification of the remaining metrically homogeneous graphs. Examples of such graphs are found in [Cam98], and we took as our own starting point the class of graphs $\Gamma_{K,C,\mathcal{S}}^\delta$ defined as follows. We introduce the class $\mathcal{A}_{K,C,\mathcal{S}}^\delta$ of all finite integer-valued metric spaces of diameter at most δ , in which there are no metric triangles with odd perimeter less than $2K + 1$, and there are no metric triangles with perimeter at least C . Here we take $1 \leq K \leq \delta$ or $K = \infty$, $2\delta + 1 \leq C \leq 3\delta + 1$. The extreme values of K correspond to no constraint or the bipartite case, while the extreme values of C correspond to the antipodal case or no constraint. Finally, \mathcal{S} is a set of $(1, \delta)$ -spaces, that is, spaces in which only the distances 1 and δ occurs (here $\delta \geq 3$), which we take as additional forbidden substructures. This is the natural extension of Henson's use of forbidden cliques to our setting. It turns out that this misses a significant source of examples determined by constraints of order 3, and when we came belatedly to test the catalog on this point we were led finally to a more complicated but similar family which will be denoted $\Gamma_{K_1, K_2; C_0, C_1; \mathcal{S}}^\delta$ in which the parameters δ, K_1, C_0, C_1 and the set \mathcal{S} are much as in the previous case, but the parameter K_2 is a little more exotic: it forbids the occurrence of certain triangles of odd perimeter P , but only those satisfying

$$P > 2K_2 + d(a, b)$$

for some pair of vertices a, b . Of course the parameters $\delta, K_1, K_2, C_0, C_1$ and the set \mathcal{S} must satisfy some auxiliary conditions to provide an amalgamation class. We will lay out the precise conditions on the parameters in detail. We will observe also that as far as the parameters $\delta, K_1, K_2, C_0, C_1$ are concerned, our conditions are expressible in Presburger arithmetic, and that this is to be expected, given that the set of constraints involved is itself definable in Presburger arithmetic from the 5 given numerical parameters. We also deal with an antipodal variation in which the set \mathcal{S} of side constraints is modified.

We will present our catalog in §5. The main point is to state the conditions on the auxiliary parameters in full. We will not give complete proofs of existence: this comes down to the amalgamation property for the classes we associate with admissible choices of the parameters. We take this up in §6, where we confine ourselves largely to a presentation of the amalgamation method. The essential point here is that one may determine amalgamations by determining one distance at a time, and in that case the range of values available for that distance (subject to the triangle inequality) is a non-empty interval. Then the various constraints associated with our parameters and the set \mathcal{S} may restrict the size and parity of the desired distance further, and one must show that some suitable value remains in all cases. We generally deal with the set \mathcal{S} by avoiding the values 1 and δ entirely. As the antipodal case is somewhat different from the rest, we give more detail in that case.

In §7 we discuss the imprimitive case, showing that with minor exceptions these are bipartite or antipodal (and in the latter case, of a very restricted type). We also classify the antipodal bipartite graphs of odd diameter.

For any metrically homogeneous bipartite graph Γ , there is an associated graph $B\Gamma$ on either half of the bipartition of Γ which is again metrically homogeneous. If Γ is an antipodal bipartite graph of odd diameter, our classification implies that when $B\Gamma$ is in our catalog, Γ is as well. When the diameter is even, the associated graph $B\Gamma$ is again antipodal, but not bipartite. Again one should aim to show that when $B\Gamma$ is in the catalog, then Γ is as well, but we did not look into that.

In §8 we turn to the proof that the exceptional metrically homogeneous graphs (the ones for which the induced graph on the neighbors of a fixed vertex either contains no infinite independent set, or is imprimitive) all lie in our catalog.

In §9 we look at another class of metrically homogeneous bipartite graphs, in terms of the structure of the associated graph $B\Gamma$. It is natural to take up the case in which $B\Gamma$ is itself exceptional. We show in this case that Γ is in the catalog. We note that there are some cases in which Γ falls under the “generic” case in our catalog and $B\Gamma$ is exceptional. These occur up through diameter 5, with $B\Gamma$ a homogeneous graph.

Lastly, we look back in §10 to some examples that can be found in tables in the Appendix to [Che98] in a very different form. There we listed all amalgamation classes corresponding to primitive infinite homogeneous structures with four orbits on pairs of distinct elements, all self-paired (in other words, four nontrivial 2-types, all of them symmetric). There are 27, and within that list one finds 17 which can be interpreted as metric spaces, some in more than one way, corresponding to 20 metrically homogeneous graphs. So we decode that list and recast it in terms which allow a direct comparison with our catalog, which does indeed contain all of these examples. We remark that with most of these 27 examples understood as metrically homogeneous graphs, one might take another look at finding a framework that accounts for the remaining ones.

The net result of our explorations has been to turn up nothing new on the sporadic side of the classification, but to broaden considerably our conception of the generic case. Given the structure of the resulting catalog, the natural way to think about a proof of its completeness is in the following terms, using the theory of amalgamation classes (which we will review in Part I). Let us use the term “generic type” for metrically homogeneous graphs which are not already in the catalog as exceptions of one kind or another.

- (1) Show that any amalgamation class of finite metric spaces associated with a metrically homogeneous graph of generic type involves exactly the same triangles (subspaces of order 3) as one of our generic classes determined by triangle constraints.
- (2) Show that any amalgamation class of finite metric spaces associated with a metrically homogeneous graph of generic type whose triangle constraints are the same as some catalogued graph of generic type, is in fact in the catalog.

In practice a proof may involve an elaborate induction in which the two sides of the issue become mixed together, but any step in the proof is likely to target just one of the two issues. An equivalent statement of the first point would be that for any amalgamation class \mathcal{A} of finite metric spaces associated with a metrically

homogeneous graph of generic type, the associated class \mathcal{A}' of finite metric spaces A such that every triangle in A belongs to \mathcal{A} is itself an amalgamation class. This is inherently plausible, but not something which one would aim to prove by a direct argument.

Our sense of these problems is that they are both difficult. We are also convinced of the correctness of the classification as far as diameter 3; the work of Amato and Macpherson [AMp10] covers the antipodal case and the case of triangle free metrically homogeneous graphs (in our notation, $K_1 > 1$). Identifying the compatible combinations of constraints on triangles is easy in this case, and the number of cases is reasonable. We have convinced ourselves that the full classification can be completed in diameter 3 by direct methods, and that the outcome is consistent with the catalog.

1.3. Is the generic triangle free graph pseudofinite? In Part III we turn our attention to another problem suggested by the study of homogeneous structures. In its general form, the problem is to find a testable criterion for a homogeneous structure to be *pseudofinite* (that is, a model of the theory of all finite structures). For example, the universal homogeneous triangle free graph will be pseudofinite if and only if for each n we can find a finite triangle free graph with the following two properties:

- (i) any maximal independent set of vertices has order at least n ; and
- (ii) for any set A of n independent vertices, and any subset B of A , there is a vertex adjacent to all vertices of B and to no vertices of $A \setminus B$.

It is known that the universal homogeneous graph is pseudofinite, because the associated problem on finite graphs is easily solved using random graphs. But this problem remains open for the universal homogeneous triangle free graph—and the more general problem is likely to remain in the shade till this particular instance is settled, one way or another.

We will discuss what is known about this problem in the triangle free case for extremely small values of n , namely $n = 3$ or 4 . There is not a great deal known, but it is worth noticing that the problem can be rephrased in terms of finite geometries, and that some concrete problems emerge that design theorists may be able to make something of. So here I aim less at “model theoretic methods in combinatorics,” and more at the hope that combinatorial methods can shed more light on this problem arising in model theory.

Some nice examples are known for the case $n = 3$, notably the Higman-Sims graph (as observed by Simon Thomas), as well as an infinite family constructed by Michael Albert, again with $n = 3$ [SWS93, p. 447]. We still have no example with $n = 4$, but we will suggest that the case $n = 3$ is worth much closer scrutiny, and raises problems that seem relatively accessible and which have a design-theoretic flavor. We also point out a hierarchy of conditions between the cases $n = 3$ and $n = 4$ which seems to us to represent a steep climb at each level. In the case $n = 3$ we have many examples which are degenerate in a precise sense, and which can be varied quite freely, while the other examples known are subgraphs of the Higman-Sims graph.

Thus we have the following key problem.

PROBLEM 1. *Is there a finite triangle free graph with the following properties, which is not a subgraph of the Higman-Sims graph?*

- (i) any maximal independent set of vertices has order at least 3; and
- (ii_{3,2}) for any set A of 3 independent vertices, and any subset B of A , there are at least two vertices adjacent to all vertices of B and to no vertices of $A \setminus B$.

This may be phrased equivalently in terms of finite combinatorial geometries and the Higman-Sims geometry on 22 points, as we shall see. In §11 we give an interpretation of the general problem of approximating the generic triangle free graph by finite ones in terms of combinatorial geometries. In §12.1 we explore the connection with strongly regular graphs. Following a suggestion made privately in 1996 by Peter Cameron, we show that there is no strongly regular graph which meets our conditions for $n = 4$. My notes on the matter are long gone, but my impression now is that the case $n = 4$ is something of a squeaker. In coming back to this I found the explicit formulas in [Big09] helpful. To take the next step and bound the size of a strongly regular graph satisfying our conditions for $n = 3$ seems to involve the central problems of the field. But perhaps the experts can do something clever in that direction.

In §12.4, we show by elementary and direct analysis directly from the definitions that any graph satisfying our conditions for $n = 4$ will have minimal degree at least 66. One might expect this lower bound to translate into an impressive lower bound on the total number of vertices, but I don't see that.

The case $n = 3$ without any assumption of strong regularity is taken up in earnest in §13, in terms of geometries rather than graphs. Just as the Higman-Sims graph on 100 vertices is more readily seen in terms of the Higman-Sims geometry on 22 points (itself a 1-point extension of a projective plane over a field of order 4), one can describe nontrivial examples in terms of geometries on relatively few points. In particular, the smallest geometry which does not fall into the class we call "Albert Geometries" lives on a set of just 8 points. There is not much general theory to be seen in the present state of knowledge, but there is one very basic open question. In a combinatorial geometry we have a set of points and a set of blocks, the blocks being sets of points. In each of the known geometries associated with the case $n = 3$, with very few exceptions, there is a block of order 2. The main question is whether there is any geometry associated with the case $n = 3$ in which the minimum block size is greater than 2, other than geometries embedding into the Higman-Sims geometry.

In §13.2 we give a family of geometries satisfying our conditions for $n = 3$, having a unique block of order 2, and with the next smallest block size arbitrarily large. Getting rid of that last block of order 2 seems to put a wholly different complexion on the matter. One would expect design theorists to be able to say something substantial about this situation, one way or another.

Bonato has explored similar problems in the context of graphs and tournaments [Bon09, Bon10]. Here one knows by probabilistic arguments that finite graphs or tournaments with analogous properties exist, but in looking for the minimal size one lands in somewhat similar territory, and again certain aspects of design theory come into consideration, including strongly regular graphs, skew Hadamard matrices, and Paley graphs or tournaments.

The theory of homogeneous structures has many other aspects that we will not touch upon, many connected with the study of the automorphism groups of homogeneous structures, e.g. the small index property and reconstruction of structures from their automorphism groups [DNT86, HHLS93, KT01, Rub94, Tru92], group theoretic issues [Tru85, Tru03, Tru09], and the classification of reducts of homogeneous structures [Tho91, Tho96], which is tied up with structural Ramsey theory. We mention also the extensive and readable survey [Mph11] by Macpherson which covers a number of directions not touched on here.

There is also an elaborate theory due to Lachlan treating stable homogeneous relational structures systematically as limits of finite structures, and by the same token giving a very general analysis of the finite case [Lac86b].

The subject of homogeneity falls under the much broader heading of “oligomorphic permutation groups”, that is the study of infinite permutation groups having only finitely many orbits on n -tuples for each n . As the underlying set is infinite, this property has the flavor of a very strong transitivity condition, and leads to a very rich theory [Cam90, Cam97].

1.4. Acknowledgment. This article has benefited enormously from a careful reading by the referee, particularly in Part II. The first draft of Part II described work in progress at the time; I learned a good deal from writing it, but not enough to rewrite it immediately. The form of the catalog given here (specifically, the classes $\mathcal{A}_{K_1, K_2; C_0, C_1; S}^\delta$ mentioned above) was still a couple of months off at that point. Having a robust catalog of examples in Part II has certainly helped matters, but the referee’s response to a variation of that first draft was also very helpful, as I was finding my way to what should be a much clearer account than the original.

CONTENTS

1. Introduction

Part I. Homogeneous Structures and Amalgamation Classes

2. Urysohn’s space
3. Fraïssé Classes and the Ramsey Property
4. Classification

Part II. Metrically Homogeneous Graphs

5. Metrically homogeneous graphs: A catalog
6. Constructions of metrically homogeneous graphs
7. Imprimitive Graphs
8. Exceptional Metrically Homogeneous Graphs
9. Exceptional Bipartite Metrically Homogeneous Graphs
10. Graphs of small diameter

Part III. Extension Properties of Finite Triangle Free Graphs

11. Extension properties
 12. Strongly regular graphs and E_4 -geometries
 13. Some E_3 -geometries
 14. Appendix: Amalgamation Classes (Tables)
- References

Part I. Homogeneous Structures and Amalgamation Classes

2. Urysohn’s space

2.1. A little history. A special number of *Topology and its Applications* (vol. 155) contains the proceedings of a conference on the Urysohn space (Beer-Sheva, 2006). A detailed account of the circumstances surrounding the discovery of that space, shortly before a swimming accident took Urysohn’s life off the coast of Brittany, can be found in the first section of [Huš08], which we largely follow here. There is also an account of Urysohn’s last days in [GK09], which provides additional context.

Urysohn completed his habilitation in 1921, and his well known contributions to topology were carried out in the brief interval between that habilitation and his fatal accident on August 17th, 1924. Fréchet raised the question of the existence of a universal complete separable metric space (one into which any other should embed isometrically) in an article published in the *American Journal of Mathematics* in 1925 with a date of submission given as August 21, 1924. Fréchet had communicated his question to Aleksandrov and Urysohn some time before that, and an announcement of a solution is contained in a letter from Aleksandrov and Urysohn to Hausdorff dated August 3, 1924, a letter to which Hausdorff replied in detail on August 11. The letter from Aleksandrov and Urysohn is quoted in the original German in [Huš08]. In that letter, the announcement of the construction of a universal complete separable metric space is followed immediately by the remark: “. . . and in addition [it] satisfies a quite powerful condition of homogeneity: the latter being, that it is possible to map the whole space onto itself (isometrically) so as to carry an arbitrary finite set M into an equally arbitrary set M_1 , congruent to the set M .” The letter goes on to note that this pair of conditions, universality together with homogeneity, actually characterizes the space constructed up to isometry. This comment on the property of homogeneity is highlighted in very similar terms in the published announcement [Ury25].

Urysohn’s construction proceeds in two steps. He first constructs a space U_0 , now called the *rational Urysohn space*, which is universal in the category of countable metric spaces with rational-valued metric. This space is constructed as a limit of *finite* rational-valued metric spaces, and Urysohn takes its completion U as the solution to Fréchet’s problem.

It is the *rational* Urysohn space which fits neatly into the general framework later devised by Fraïssé [Fra54]. A countable structure is called *homogeneous* if any isomorphism between finitely generated substructures is induced by an automorphism of M . If we construe metric spaces as structures in which the metric defines a weighted complete graph, with the metric giving the weights, then finitely generated substructures are just finite subsets with the inherited metric, and isomorphism is isometry. Other examples of homogeneity arise naturally in algebra, such as vector spaces (which may carry forms—symplectic, orthogonal, or unitary), or algebraically closed fields. We will mainly be interested in relational systems, that is combinatorial structures in which “f.g. substructure” simply means “finite subset, with the induced structure.” But Fraïssé’s general theory, to which we turn in the next section, does allow for the presence of functions.

2.2. Topological dynamics and the Urysohn space. The Urysohn space, or rather its group of isometries, turns up in topological dynamics as an example of

the phenomenon of *extreme amenability*. A topological group is said to be extremely amenable if any continuous action on a compact space has a fixed point. The isometry group of Urysohn space is shown to be extremely amenable in [Pes02], and subsequently the general theory of [KPT05] showed that the isometry group of the *ordered rational Urysohn space* (defined in the next section) is also extremely amenable. The general theory of [KPT05] requires Fraïssé’s setup, but we quote one of the main results in advance:


THEOREM 1. [KPT05, Theorem 2] *The extremely amenable closed subgroups of the infinite symmetric group Sym_∞ are exactly the groups of the form $\text{Aut}(F)$, where F is the Fraïssé limit of a Fraïssé order class with the Ramsey property.*

We turn now to the Fraïssé theory.

3. Fraïssé Classes and the Ramsey Property

3.1. Amalgamation Classes. It is not hard to show that any two countable homogeneous structures of a given type will be isomorphic if and only if they have the same isomorphism types of f.g. substructures. This uses a “back-and-forth” construction, as in the usual proof that any two countable dense linear orders are isomorphic, which is indeed a particular instance. In view of this uniqueness, it is natural to look for a characterization of countable homogeneous structures directly in terms of the associated class $\text{Sub}(M)$ of f.g. structures embedding into M . Fraïssé identified the relevant properties:

- I. $\text{Sub}(M)$ is hereditary (closed downward, and under isomorphism): in other words, if A is in the class, then any f.g. structure isomorphic with a substructure of A is in the class;
- II. There are only countably many isomorphism types represented in $\text{Sub}(M)$;
- III. $\text{Sub}(M)$ has the *joint embedding* and *amalgamation* properties: if A_1, A_2 are f.g. substructures in the class, and A_0 embeds into A_1, A_2 isomorphically via embeddings f_1, f_2 , then there is a structure \hat{A} in the class with embeddings $A_1, A_2 \rightarrow \hat{A}$, completing the diagram



The joint embedding property is the case in which A_0 is empty, which should be treated as a distinct condition if one does not allow empty structures.

A key point is that amalgamation follows from homogeneity: taking A_0 to be a subset of A_1 , and embedded into A_2 , apply an automorphism of the ambient structure to move the image of A_0 in A_2 back to A_0 , and A_2 to some isomorphic structure A'_2 containing A_0 ; then the structure generated by $A_1 \cup A'_2$ will serve as an amalgam.

Conversely, if \mathcal{A} is a class of structures with properties (I – III), then there is a countable homogeneous structure M , unique up to isomorphism, for which

$\text{Sub}(M) = \mathcal{A}$. This homogeneous structure M is called the *Fraïssé limit* of the class \mathcal{A} . Thus with \mathcal{A} taken as the class of finite linear orders, the Fraïssé limit is isomorphic to the rational order; with \mathcal{A} the class of all finite graphs, the Fraïssé limit is Rado’s universal graph [Rad64]; and with \mathcal{A} the class of finite rational-valued metric spaces, after checking amalgamation, we take the Fraïssé limit to get the rational Urysohn space. The computation that checks amalgamation can be found in Urysohn’s own construction, though not phrased as such.

So we see that Fraïssé’s theory is at least a ready source of “new” homogeneous structures, and we now give a few more examples in the same vein. Starting with the class of all partial orders, we obtain the “generic” countable partial order (to call it merely “dense,” as in the linear case, would be to understate its properties). Or starting with the class of finite triangle free graphs, we get the “generic” triangle free graph, and similarly for any n the generic K_n -free graph will be obtained [Hen71]. The amalgamation procedure here is simply graph theoretical union, and the special role of the complete graphs here is due to their *indecomposability* with respect to this amalgamation procedure: a complete graph which embeds into the graph theoretical union of two graphs (with no additional edges permitted) must embed into one of the two. The generalization to the case of directed graphs is immediate: amalgamating via the graph theoretical union, the indecomposable directed graphs are the tournaments. So we associate with any set of tournaments \mathcal{T} the Fraïssé class of finite directed graphs *omitting* \mathcal{T} —i.e., with no directed subgraph isomorphic to one in \mathcal{T} . The corresponding Fraïssé limits are the generic \mathcal{T} -free graphs considered by Henson [Hen72]. The richness of the construction is confirmed by showing that 2^{\aleph_0} directed graphs arise in this way, because of the existence of an *infinite antichain* in the class of finite tournaments, that is an infinite set \mathcal{X} of finite tournaments, which are pairwise incomparable under embedding, so that each subset of \mathcal{X} gives rise to a different Fraïssé limit. A suitable construction of such an antichain is given by Henson [Hen72]. A parallel construction in the category of commutative rings provides, correspondingly, uncountably many homogeneous commutative rings [CSW93].

The structure of the infinite antichains of finite tournaments has been investigated further, but has not been fully elucidated. Any such antichain lies over one which is minimal in an appropriate sense, and after some close analysis by Latka [Lat94, Lat03, Lat02] a general finiteness theorem emerged [ChL00] to the effect that for any fixed k , there is a finite set of minimal antichains which will serve as universal witnesses for any collection of finite tournaments determined by k constraints (forbidden tournaments) which allows an infinite antichain. This means that whenever such an antichain is present, one of the given antichains is also present, up to a finite difference. But there is still no known *a priori* bound for the number of antichains required, as a function of k . Even the question as to whether the number of antichains needed is bounded by a computable function of k remains open.

In the terminology of [KPT05], the notion of Fraïssé class is taken to incorporate a further condition of *local finiteness*, meaning that all f.g. structures are finite. This may be viewed as a strengthening of condition (II). This convention is in force, in particular, in the statement of Theorem 1, which we now elucidate further.

3.2. Order Classes. The Fraïssé classes that occur in the theorem of Kechris-Pestov-Todorćević above are *order classes*: this means that the structures considered are equipped with a distinguished relation $<$ representing a linear order. Thus in this theorem nothing is said about graphs, directed graphs, or metric spaces, but rather their ordered counterparts: ordered graphs, ordered directed graphs, ordered metric spaces. In particular the *ordered rational Urysohn space* is, by definition, the homogeneous ordered rational valued metric space delivered by Fraïssé’s theory. As there is no connection between the order and the metric the necessary amalgamation may be carried out separately in both categories.

In the most straightforward, and most common, applications of the Fraïssé theory there is often some notion of “free amalgamation” in use. In the case of order classes amalgamation cannot be entirely canonical, as some “symmetry breaking” is inevitable. But there are also finite homogeneous structures—such as the pentagon graph, or 5-cycle—for which the theory of amalgamation classes is not illuminating, and the amalgamation procedure consists largely of the forced identification of points.

When one passes from the *construction* of examples to their systematic *classification*, there is typically some separation between the determination of more or less sporadic examples, and the remaining cases described naturally by the Fraïssé theory. Such a classification has only been carried out in a few cases, and perhaps a more nuanced picture will appear eventually. But this simple picture continues to guide our expectations for the classification of metrically homogeneous graphs, considered in the next Part, and so far everything we have seen is consistent with that picture in this case.

Below we will say something more about how the theory of [KPT05] applies in the absence of order, but first we complete the interpretation of Theorem 1 by discussing the second key property required.

3.3. The Ramsey Property. The ordinary Ramsey theorem is expressed in Hungarian notation by the symbolism:

$$\forall k, m, n \exists N : N \rightarrow (n)_k^m$$

meaning that for given k, m, n , there is N so that: for any coloring of increasing m -tuples from $A = \{1, \dots, N\}$ by k colors, there is a subset B of cardinality n which is monochromatic with respect to the coloring.

Structural Ramsey theory deals with a locally finite hereditary class \mathcal{A} of finite structures of fixed type, which on specialization to the case of the class \mathcal{L} of finite linear orders will degenerate to the usual Ramsey theory. In general, given two structures A, B in \mathcal{A} , write $\binom{B}{A}$ for the class of induced substructures of B isomorphic to A . This gives $\binom{N}{n}$ an appropriate meaning if $\mathcal{A} = \mathcal{L}$, namely increasing sequences of length n from an ordered set of size N .

We may then use the Hungarian notation

$$M \rightarrow (B)_k^A$$

to mean that whenever we have a coloring of $\binom{M}{A}$ by k colors, there is a copy of B inside M which is monochromatic with respect to the induced coloring of $\binom{B}{A}$. And the *Ramsey property* will be:

$$\forall A, B \in \mathcal{A} \forall k \exists M \in \mathcal{A} : M \rightarrow (B)_k^A$$

So the Ramsey property for \mathcal{L} is the usual finite Ramsey theorem.

Ramsey theory for Fraïssé classes is a subtle matter, but a highly developed one. In [HN05] it is shown that the Ramsey property implies the amalgamation property, by a direct argument. What one would really like to classify are the Fraïssé classes with the Ramsey property, but according to [HN05] the most promising route toward that is via classification of amalgamation classes first, and then the identification of the Ramsey classes.

For unordered graphs, the only instances of the Ramsey property that hold are those for which the subgraphs being colored are complete graphs K_n , or their complements [NR75b]. But the collection of finite ordered graphs does have the Ramsey property [NR77a, NR77b, AH78].

To illustrate the need for an ordering, consider colorings of the graph $A = K_1 + K_2$, a graph on 3 vertices with one edge, and let $B = 2 \cdot K_2$ be the disjoint sum of two complete graphs of order 2. If we order B in any way, we may color the copies of A in B by three colors according to the relative position of the isolated vertex of A , with respect to the other two vertices, namely before, after, or between them. Then B , with this coloring, cannot be monochromatic. Thus we can never have a graph G satisfying $G \rightarrow (B)_3^A$, since given such a graph G we would first order G , then define a coloring of copies of A in G as above, using this order, and there could be no monochromatic copy of B .

The topological significance of the Ramsey property for the ordered rational Urysohn space only emerged in [KPT05], and the appropriate structural Ramsey theorem was proved “on demand” by Nešetřil [Neš07].

At this point, we have collected all the notions needed for Theorem 1. Before we leave this subject, we note that the theory of [KPT05] also exhibits a direct connection between topological dynamics and the more classical examples of the Fraïssé theory (lacking a built-in order). The following is a fragment of Theorem 5 of [KPT05].

THEOREM 2. *Let G be the automorphism group of one of the following countable structures M :*

- (1) *The random graph;*
- (2) *The generic K_n -free graph, $n \geq 2$;*
- (3) *The rational Urysohn space.*

Let L be the space of all linear orderings of M , with its compact topology as a closed subset of $2^{M \times M}$. Then under the natural action of G on L , L is the universal minimal compact flow for G .

The *minimality* here means that there is no proper closed invariant subspace; and the *universality* means that this is the largest such minimal flow (projecting on to any other). Again, Theorem 2 has an abstract formulation in terms of Fraïssé theory [KPT05]. The following is a special case.

THEOREM 3 ([KPT05, Theorem 4]). *Let \mathcal{A} be a (locally finite) Fraïssé class and let \mathcal{A}^+ be the class of ordered structures $(K, <)$ with $K \in \mathcal{A}$. Suppose that \mathcal{A}^+ is a Fraïssé class with the order property and the Ramsey property. Let M be the Fraïssé limit of \mathcal{A} , $G = \text{Aut}(M)$, and \mathcal{L} the space of linear orderings of M , equipped with the compact topology inherited by inclusion into $2^{M \times M}$. Then under the natural action of G on \mathcal{L} , the space \mathcal{L} is the universal minimal compact flow for G .*

Here one has in mind the case in which amalgamation in \mathcal{A} does not require any identification of vertices (strong amalgamation); then \mathcal{A}^+ is certainly an amalgamation class. The *order property* is the following additional condition: given $A \in \mathcal{A}$, there is $B \in \mathcal{A}$ such that under any ordering on A , and any ordering on B , there is some order preserving isomorphic embedding of A into B . This is again a property which must be verified when needed, and is known in the cases cited. To see an example where the order property does not hold, consider the class \mathcal{A} of finite equivalence relations. Any equivalence relation B may be ordered so that its classes are intervals; thus the order property fails.

We turn next to the problem of classifying homogeneous structures of particular types. Here again the Fraïssé theory provides the key.

4. Classification

The homogeneous structures of certain types have been completely classified, notably homogeneous graphs [LW80], homogeneous tournaments [Lac84], homogeneous tournaments with a coloring by finitely many colors and homogeneous directed graphs [Che88], homogeneous partial orders with a coloring by countably many colors [TrT08], and homogeneous permutations [Cam03]: this last is a less familiar notion, that we will enlarge upon. There is also work on the classification of homogeneous 3-hypergraphs [LT95, AL95], and on graphs with two colors of edges [Lac86a, Che99], the latter covering only the finite and imprimitive cases: this uncovers some sporadic examples, but the main problem remains untouched in this class.

4.1. Homogeneous Permutations. Cameron observed that permutations have a natural interpretation as structures, and that when one adopts that point of view the model theoretic notion of embedding is the appropriate one. A finite permutation may naturally be viewed as a finite structure consisting of two linear orderings. This is equivalent to a pair of bijections between the structure and the set $\{1, \dots, n\}$, n being the cardinality, and thus to a permutation. In this setting, an embedding of one permutation into another is an occurrence in the second of a *permutation pattern* corresponding to the first, so that this formalism meshes nicely with the very active subject of permutations omitting specified patterns (“pattern classes”).

By a direct analysis, Cameron showed that there are just 6 homogeneous permutations, in this sense, up to isomorphism: the trivial permutation of order 1, the identity permutation of \mathbb{Q} or its reversal, the class corresponding to the lexicographic order on $\mathbb{Q} \times \mathbb{Q}$, where the second order agrees with the first in one coordinate and reverses the first in the other coordinate, and the generic permutation (corresponding to the class of all finite permutations). The existence of the generic permutation is immediate by the Fraïssé theory.

As amalgamations of linear orders are tightly constrained, the classification of the amalgamation classes of permutations is quite direct. Cameron also observes that it would be natural to generalize from structures with two linear orders to an arbitrary finite number, but I do not know of any further progress on this interesting question.

4.2. Homogeneous Graphs. This is the case that really launched the classification project. The classification of homogeneous graphs by Lachlan and Woodrow

involves an ingenious inductive setup couched directly in terms of amalgamation classes of finite graphs. We will need the results of that classification later, when discussing metrically homogeneous graphs. Indeed, homogeneous graphs are just the diameter two (or less) case of metrically homogeneous graphs. Furthermore, in any metrically homogeneous graph, the graph induced on the neighbors of a point is a homogeneous graph, and we will find it useful to consider the possibilities individually, or more exactly to distinguish the “exceptional” and “generic” cases, and to treat the exceptional ones on an ad hoc basis.

In our catalog of the homogeneous graphs we will use the following notation. The graph K_n is a complete graph of order n , allowing $n = \infty$, which stands for \aleph_0 in this context. We write I_n for the complement of K_n , that is an independent set of vertices of order n , and $m \cdot K_n$ for the disjoint sum of m copies of K_n , again allowing m and n to become infinite. Bearing in mind that the complement of a homogeneous graph is again a homogeneous graph, we arrange the list as follows.

- I. *Degenerate cases, K_n or I_n* ; these are actually homogeneous structures for a simpler language (containing just the equality symbol).
- II. *Imprimitive homogeneous graphs, $m \cdot K_n$ and their complements*, where $m, n \geq 2$. The complement of $m \cdot K_n$ is complete n -partite with parts of constant size.
- III. *Primitive, nondegenerate, homogeneous finite graphs* (highly exceptional): the pentagon or 5-cycle C_5 , and a graph on 9 points which may be described as the line graph of the complete bipartite graph $K_{3,3}$, or the graph theoretic square of K_3 . These graphs are isomorphic with their complements.
- IV. *Primitive, nondegenerate, infinite homogeneous graphs*, with which the classification is primarily concerned: Henson’s generic graphs omitting K_n , and their complements, generic omitting I_n , as well as the generic or random graph Γ_∞ (Rado’s graph) corresponding to the class of all finite graphs. Rado’s graph is isomorphic with its complement.

In this setting there is no difficulty identifying the degenerate and imprimitive examples, and little difficulty in identifying the remaining finite ones by an inductive analysis. Since the class of homogeneous graphs is closed under complementation, the whole classification comes down to the following result.

THEOREM 4 ([LW80, Theorem 2’, paraphrased]). *Let Γ be a homogeneous non-degenerate primitive graph containing an infinite independent set, as well as the complete graph K_n . Then Γ contains every finite graph not containing a copy of K_{n+1} .*

Let us see that this completes the classification in the infinite, primitive, nondegenerate case. As the graph Γ under consideration is infinite, by Ramsey’s theorem it contains either K_∞ or I_∞ , and passing to the complement if necessary, we may suppose the latter. So if K_n embeds in Γ and K_{n+1} does not, then Theorem 4 says that Γ is the corresponding Henson graph, while if K_n embeds in Γ for all n , Theorem 4 then says that it is the Rado graph.

The method of proof is by induction on the order N of the finite graph G which we wish to embed in Γ . The difficulty is that on cursory inspection, Theorem 4 does not at all lend itself to such an inductive proof. Lachlan and Woodrow show that as sometimes happens in such cases, a stronger statement may be proved

by induction. Their strengthening is on the extravagant side, and involves some additional technicalities, but it arises naturally from the failure of the first try at an inductive argument. So let us first see what difficulties appear in a direct attack.

Let G be a graph of order N , not containing K_{n+1} , and let Γ be the homogeneous graph under consideration. We aim to show that G embeds in Γ , proceeding by induction on N . Pick a vertex v of G . If v is isolated, or if v is adjacent to the remaining vertices of G , we will need some special argument (even more so later, once we strengthen our inductive claim). For example, if v is adjacent to the remaining vertices of G , then we have an easy case: we identify v with any vertex of Γ , we consider the graph Γ_1 on the vertices adjacent to v in Γ , and after verifying that Γ_1 inherits all hypotheses on Γ (with n replaced by $n - 1$) we can conclude directly by induction on n . If the vertex v is isolated, the argument will be less immediate, but still quite manageable.

Turning now to the main case, when the vertex v has both a neighbor and a non-neighbor in G , matters are considerably less simple. Let G_0 be the graph induced on the other vertices of G , and let a, b be vertices of G_0 chosen with a adjacent to v , and b not. At this point, we must build an amalgamation diagram which forces a copy of G into Γ , and we hope to get the factors of the diagram by induction on N , which of course does not quite work. It goes like this.

Let A be the graph obtained from G by deleting v and a , and let B be the graph obtained from G by deleting v and b . Let H_0 be the disjoint sum $A + B$ of A and B , and form two graphs $H_1 = H_0 \cup \{u\}$ and $H_2 = H_0 \cup \{c\}$, with edge relations as follows. The vertex u plays the role of v , and is therefore related to A and B as v is in G . The vertex c plays a more ambiguous role, as a or b , and is related to A as a is, and to B as b is.

Suppose for the moment that copies of H_1, H_2 occur in Γ . Then so does an amalgam $H_1 \cup H_2$ over H_0 , and in that amalgam either u is joined to c , which may then play the role of a , with the help of A , or else u is not joined to c , and then c may play the role of b , with the help of B . In either case, a copy of G is forced into Γ .

What may be said about the structure of H_1 and H_2 ? These are certainly too large to be embedded into Γ *by induction*, but they have a simple structure: H_2 is the free amalgam of $A \cup \{a\}$ with $B \cup \{b\}$ under the identification of a with b , and H_1 is constructed similarly, over u . Here each factor (e.g., $A \cup \{a\}$, $A \cup \{b\}$ in the case of H_2) embeds into Γ by induction, but we need also the sum of the two factors over a common vertex. This leads to the following definitions.

DEFINITION 4.1.

- (1) A *pointed graph* (G, v) is a graph G with a distinguished vertex v .
- (2) The *pointed sum* of two pointed graphs (G, v) and (H, w) is the graph obtained from the disjoint sum $G + H$ by identifying the base points.
- (3) Let $\mathcal{A}(n)$ be the set of finite graphs belonging to every amalgamation class which contains K_n, I_∞ , the path of order 3, and its complement (the last two eliminate imprimitive cases).
- (4) Let $\mathcal{A}^*(n)$ be the set of finite graphs G such that for any vertex v of G , and any pointed graph (H, w) with $H \in \mathcal{A}(n)$, the pointed sum $(G, v) + (H, w)$ belongs to $\mathcal{A}(n)$.

Notice that $\mathcal{A}^*(n)$ is contained in $\mathcal{A}(n)$ for trivial reasons, simply taking for (H, w) the pointed graph of order 1. Now we can state the desired strengthening of Theorem 2'.

THEOREM 5 ([LW80, Lemma 6]). *For any n , if G is a finite graph omitting K_{n+1} , then G belongs to $\mathcal{A}^*(n)$.*

With this definition, the desired inductive proof actually goes through. Admittedly the special cases encountered in our first run above become more substantial the second time around. As a result, this version of the main theorem will be preceded by 5 other preparatory lemmas required to support the final induction. However the process of chasing one's tail comes to an end at this point.

4.3. Homogeneous Tournaments. In Lachlan's classification of the homogeneous tournaments [Lac84] two new ideas occur, which later turned out to be sufficient to carry out the full classification of the homogeneous directed graphs [Che98], with suitable orchestration. A byproduct of that later work was a more efficient organization of the case of tournaments, given in [Che88]. The main idea introduced at this stage was a certain use of Ramsey's theorem that we will describe in full. The second idea arises naturally at a later stage as one works through the implications of the first; it involves an enlargement of the setting beyond tournaments, where much as in the case of the Lachlan/Woodrow argument, the point is to find an inductive framework large enough to carry through an argument that leads somewhat beyond the initial context of homogeneous tournaments.

It turns out that there are only 5 homogeneous tournaments, four of them of a special type which are easily classified, and the last one fully generic. The whole difficulty comes in the characterization of this last tournament as the only homogeneous tournament of general type, in fact the only one containing a specific tournament of order 4 called $[T_1, C_3]$. In this notation, T_1 is the tournament of order 1, C_3 is a 3-cycle, and $[T_1, C_3]$ is the tournament consisting of a vertex (T_1) dominating a copy of C_3 . So the analog of Theorem 2' of [LW80] is the following.

THEOREM 6 ([Lac84]). *Let T be a countable homogeneous tournament containing a tournament isomorphic with $[T_1, C_3]$. Then every finite tournament embeds into T .*

We now give the classification of the homogeneous tournaments explicitly, and indicate the reduction of that classification to Theorem 6.

A *local order* is a tournament with the property that for any vertex v , the tournaments induced on the sets $v^+ = \{u : v \rightarrow u\}$ and $v^- = \{u : u \rightarrow v\}$ are both transitive (i.e., given by linear orders). Equivalently, these are the tournaments not embedding $[T_1, C_3]$ or its dual $[C_3, T_1]$. There is a simple structure theory for the local orders, which we will not go into here. But the result is that there are exactly four homogeneous local orders, two of them finite: the trivial one of order 1, and the 3-cycle C_3 . The infinite homogeneous local orders are the rational order $(\mathbb{Q}, <)$ and a very similar generic local order, which can be realized equally concretely.

Now a tournament T which does not contain a copy of $[T_1, C_3]$ can easily be shown to be of the form $[S, L]$ where S is a local order whose vertices all dominate a linear order L ; here S or L may be empty. Indeed, if T is homogeneous, one of the two must be empty, and in particular T is a local order. Thus if the homogeneous tournament T contains a copy of $[T_1, C_3]$ then it contains a copy of $[C_3, T_1]$ as well, and it remains only to prove Theorem 6.

At this point, the following interesting technical notion comes into the picture. If \mathcal{A} is an amalgamation class, let \mathcal{A}^* be the set of finite tournaments T such that every tournament T^* of the following form belongs to \mathcal{A} : $T^* = T \cup L$, L is linear, and every pattern of edges between T and L is permitted. Theorem 6 is equivalent to the following rocco variation.

THEOREM 7. *If \mathcal{A} is an amalgamation class containing $[T_1, C_3]$, then \mathcal{A}^* is an amalgamation class containing $[T_1, C_3]$.*

That \mathcal{A}^* is an amalgamation class is simply an exercise in the definitions, but worth working through to see why the definition of \mathcal{A}^* takes the particular form that it does (because linear orders have strong amalgamation). The deduction of Theorem 6 from Theorem 7 is also immediate, as we will now verify.

Assuming Theorem 7, we argue by induction on $N = |A|$ that any finite tournament A belongs to any amalgamation class \mathcal{A} containing $[T_1, C_3]$. Take any vertex v of A and let A_0 be the tournament induced on the remaining vertices. By induction, A_0 belongs to every amalgamation class containing $[T_1, C_3]$; in particular $A_0 \in \mathcal{A}^*$. Since A is the extension of A_0 by a single vertex, and since a single vertex constitutes a linear tournament, then from $A_0 \in \mathcal{A}^*$ we derive $A \in \mathcal{A}$, and we are done.

Note the progress which has been made. In Theorem 6 we consider arbitrary tournaments; in Theorem 7 we consider only linear extensions of $[T_1, C_3]$. Now a further reduction comes in, and eventually the statement to be proved reduces to a finite number of specific instances of Theorem 6 which can be proved individually. But we have not yet encountered the leading idea of the argument, which comes in at the next step.

4.4. The Ramsey Argument. We introduce another class closely connected with \mathcal{A}^* .

DEFINITION 4.2.

1. For tournaments A, B we define the *composition* $A[B]$ as usual as the tournament derived from A by replacing each vertex of A by a copy of B , with edges determined within each copy of B as in B , and between each copy of B , as in A . The composition of two tournaments is a tournament.

2. If T is any tournament, a *stack* of copies of T is a composition $L[T]$ with L linear.

3. If \mathcal{A} is an amalgamation class of finite tournaments, let \mathcal{A}^{**} be the set of tournaments T such that every tournament T^* of the following form belongs to \mathcal{A} : $T^* = L[T] \cup \{v\}$ is an extension of some stack of copies of T by one more vertex.

The crucial point here is the following.

FACT 4.3. *Let \mathcal{A} be an amalgamation class of finite tournaments, and T a finite tournament in \mathcal{A}^{**} . Then T belongs to \mathcal{A}^* .*

We will not give the argument here. It is a direct application of the Ramsey theorem, given explicitly in [Lac84] and again in [Che88, Che98]. The idea is that one may amalgamate a large number of one point extensions of a long stack of copies of T so that in any amalgam, the additional points contain a long linear tournament, and one of the copies of T occurring in the stack will hook up with that linear tournament in any previously prescribed fashion desired.

This leads to our third, and nearly final, version of the main theorem.

THEOREM 8. *Let \mathcal{A} be an amalgamation class of finite tournaments containing $[T_1, C_3]$. Then C_3 belongs to \mathcal{A}^{**} .*

Notice that a stack of copies of $[T_1, C_3]$ embeds in a longer stack of copies of C_3 , so that Theorem 8 immediately implies the same result for $[T_1, C_3]$. Since we already saw that $\mathcal{A}^{**} \subseteq \mathcal{A}^*$, Theorem 8 implies Theorem 7. In view of the very simple structure of a stack of copies of T , we are almost ready to prove Theorem 8 by induction on the length of the stack. Unfortunately the additional vertex v occurring in $T^* = L[T] \cup \{v\}$ complicates matters, and leads to a further reformulation of the statement.

At this point, it is convenient to return from the language of amalgamation classes to the language of structures. So let the given amalgamation class correspond to the homogeneous tournament \mathbb{T} , and let $T = L[C_3] \cup \{v\}$ be a 1-point extension of a finite stack of 3-cycles. Theorem 8 says that T embeds into \mathbb{T} . It will be simpler to strengthen the statement slightly, as follows.

Let a be an arbitrary vertex in \mathbb{T} , and consider

$$\mathbb{T}_1 = a^+ = \{v : a \rightarrow v\} \text{ and } \mathbb{T}_2 = a^- = \{v : v \rightarrow a\}$$

separately. We claim that T embeds into \mathbb{T} with $L[C_3]$ embedding into \mathbb{T}_1 , and with the vertex v going into \mathbb{T}_2 . This now sets us up for an inductive argument in which we consider a single 3-cycle C in \mathbb{T}_1 , and the parts \mathbb{T}'_1 and \mathbb{T}'_2 defined relative to C as follows: \mathbb{T}'_1 consists of the vertices of \mathbb{T}_1 dominated by the three vertices of C , and \mathbb{T}'_2 consists of the vertices v' of \mathbb{T}_2 which relate to C as the specified vertex v does. What remains at this point is to clarify what we know, initially, about \mathbb{T}_1 and \mathbb{T}_2 , and to show that these properties are inherited by \mathbb{T}'_1 and \mathbb{T}'_2 (in particular, \mathbb{T}'_2 should be nonempty!). This then allows an inductive argument to run smoothly.

At this point we have traded in the tournament \mathbb{T} for a richer structure $(\mathbb{T}_1, \mathbb{T}_2)$ consisting of a tournament with a distinguished partition into two sets. The homogeneity of \mathbb{T} will give us the homogeneity of $(\mathbb{T}_1, \mathbb{T}_2)$ in its expanded language. Such structures will be called 2-tournaments, and the particular class of 2-tournaments arising here will be called *ample tournaments*. The main inductive step in the proof of Theorem 8 will be the claim that an ample tournament $(\mathbb{T}_1, \mathbb{T}_2)$ gives rise to an ample tournament $(\mathbb{T}'_1, \mathbb{T}'_2)$ if we fix a 3-cycle in \mathbb{T}_1 and pass to the subsets considered above.

We will not dwell on this last part. The main steps in the proof are the reduction to Theorem 8, and then the realization that we should step beyond the class of homogeneous tournaments to the class of homogeneous 2-tournaments, to find a setting which is appropriately closed under the construction corresponding to the inductive step of the argument. This then leaves us concerned only about the base of the induction, which reduces to a small number of specific claims about tournaments of order not exceeding 6. Once the problem is finitized, it can be settled by explicit amalgamation arguments.

Lachlan's Ramsey theoretic argument functions much the same way in the context of directed graphs as it does for tournaments, and comes more into its own there, as it is not a foregone conclusion that Ramsey's theorem will necessarily produce a linear order; but it will produce something, and modifying the definition of \mathcal{A}^* to allow for this additional element of vagueness, things proceed much as they did before.

In [Che98] there is also a treatment of the case of homogeneous graphs using the ideas of [Lac84] in place of the methods of [LW80]. This cannot be said to be a simplification, having roughly the complexity of the original proof, but it is a viable alternative, and the proof of the classification of homogeneous directed graphs is more or less a combination of the ideas of the tournament classification with the ideas which appear in a treatment of homogeneous graphs by this second method.

This ends our general survey of the general theory of amalgamation classes and its application to classification results. It is not clear how much further those ideas can be taken. The proofs are long and ultimately computational even when the final classifications have a reasonably simple form, and at this level of generality one has good methods but no very general theory. In the finite case, one has an excellent theory due to Lachlan making good use of model theoretic ideas. These were Lachlan’s words in [Lac86b]:

The situation can be summarized as follows: Finite homogeneous structures are well understood. Stable homogeneous structures turn out to be just the unions of chains of finite ones. Thus, understanding stable homogeneous structures goes hand in hand with understanding finite ones. Beyond this, some special cases have been investigated successfully, but almost no general results have been obtained.

That assessment stands today as far as the theory of homogeneous structures is concerned. A generalization of the theory of finite homogeneous structures beyond the homogeneous framework was also envisioned by Lachlan, and came to fruition, based on a combination of permutation group theory and model theory [KLM89, ChH03].

It seems to be impossible to say at this point how much farther one can go with the methods of classification for homogeneous structures currently available. While we have few general results, we also have no known limitations on the method. The existence of 2^{\aleph_0} homogeneous directed graph was once taken as such a limitation, a point of view I shared till seeing the classification of homogeneous tournaments. At that point, there was a tension between the existence of 2^{\aleph_0} known examples of homogeneous digraphs and the fact that there was no clear obstruction to the use of Lachlan’s methods in this case. This conflict was resolved in favor of Lachlan’s methods [Che98]. So we know less now than we thought we did originally. It still seems doubtful to me that the methods of Lachlan and Woodrow can be pushed very far beyond their current range, but on the other hand we have not actually found any concrete evidence of their limitations, or indeed any homogeneous relational structures that are not readily accounted for as either occurring in nature or coming naturally from the Fraïssé theory.

We believe that the classification of the metrically homogeneous graphs will provide another case in which some sporadic examples can be accounted for as growing naturally out of one or another special phenomenon, and the remainder fall neatly into the theory of amalgamation classes of “generic” type. The evidence for this is admittedly thin—the catalog of known types reached its present form after the first draft of the present paper was complete, and a coherent plan for a proof of its completeness still does not exist. But the catalog strikes me as satisfactorily robust now, and whenever one has a catalog with a clear division of

exceptional and generic cases, one has some reason to expect the existing theory to be adequate to a proof of its completeness, with the proviso that there is often a striking disparity between the complexity of the catalog and the complexity of the resulting proof.

With all this in mind—or out of mind, as the case may be—we will turn in the next part to a catalog of the known metrically homogeneous graphs, and a full discussion of those which are “non-generic” in one of a number of senses. As we mentioned earlier, we will not give full existence proofs for the metrically homogeneous graphs in our catalog, though we will give some of the leading ideas and full proofs in special cases. The main class we present depends on four numerical parameters satisfying some simple linear inequalities and parity constraints, in more than one possible combination. We can account for those inequalities and parity constraints on abstract grounds: they are connected heuristically with quantifier elimination in Presburger arithmetic (§5.4).

In another direction, we think the generalization of Cameron’s classification of homogeneous permutations to the case of structures equipped with k linear orders (also called k -dimensional permutations [Wat07, §5.9]) is another attractive classification problem. Here one can easily make a catalog of the “natural” examples but it is unclear whether one should expect that catalog to be complete. This problem seems to us to have something to do with the existing theory of weakly o -minimal structures. But we will not explore the matter here.

Part II. Metrically Homogeneous Graphs

5. Metrically homogeneous graphs: A catalog

5.1. The Classification Problem. Any connected graph may be considered as a metric space under the graph metric, and if the associated metric space is homogeneous then the graph is said to be *metrically homogeneous*¹ [Cam98]). Cameron asked whether this class of graphs can be *completely classified*, and gave some examples of constructions via the Fraïssé theory of amalgamation classes.

We believe that such a classification can be given. As a first step, we will give a catalog of all the known metrically homogeneous graphs, with the expectation that this catalog is complete or nearly so. That catalog is the focus of the present part. It consists of a few graphs of exceptional types, and two “generic” families which are best understood in terms of the Fraïssé theory of amalgamation classes.

Since the main examples in the catalog are presented in terms of amalgamation classes, it is necessary to check the amalgamation property for the classes we define. This is not trivial, and as there are a number of distinct cases to consider, it will not be covered in detail. The main point of §6 will be to lay out explicitly the amalgamation procedure followed. We will leave for another occasion a detailed proof that this procedure succeeds for all the classes considered. We will go into more detail in the discussion of a variant appropriate to the so-called antipodal case, as this is distinct perturbation of the general case.

The second order of business is to show that this catalog is reasonably complete. This is largely a byproduct of the way the catalog was constructed. There are two natural notions of “exceptional” metrically homogeneous graph. In addition, there

¹A considerably weaker notion occurs in the geometry literature under the name *metrically homogeneous set*

is a natural Fraïssé style construction whose main ingredient is an amalgamation class determined entirely by constraints on triangles. The most difficult point to work through was the determination of this last class of examples. This caught us by surprise; the analogous step in previous problems of this type has been straightforward. In the present case, the conditions on an amalgamation class determined by constraints of order 3 depend on five numerical parameters (one of which is the diameter δ , which we always have lurking in the background). We denote the resulting classes by $\mathcal{A}_{K_1, K_2; C_0, C_1}^\delta$ where the numerical parameters K_1, K_2 are used to specify which metric triangles of small odd perimeter are forbidden, and the parameters C_0, C_1 are used to specify which triangles of large perimeter, even or odd respectively, are forbidden. When we come to the details it will be seen that the parameter K_2 is used in a more subtle way than the other three parameters. We more or less guessed the role of K_1, C_0, C_1 at the outset, except that we expected $C_1 = C_0 \pm 1$, and hence we worked with just two parameters, $K = K_1$ and $C = \min(C_0, C_1)$. The role of K_2 came as a surprise and we took the first examples found to be sporadic.

In the next subsection we will lay out our notions of “exceptional” metrically homogeneous graph and “generic” metrically homogeneous graph explicitly, and indicate the way the catalog was devised, before actually giving the catalog. One of our two notions of “exceptional” metrically homogeneous graph (the imprimitive case) turns out to lie mainly on the generic side in the catalog as it now stands.

Cameron made a number of fundamental observations on the classification problem in [Cam98]. He noted that the Lachlan/Woodrow classification is the diameter 2 case. He pointed out that Fraïssé constructions give graphs of any fixed diameter whose associated metric spaces are analogous to the Urysohn space, but with bounded integral distance, and that there is a bipartite variant of this construction. He also noted related work by Komjáth, Mekler and Pach which turns out to be very closely connected with the construction of graphs of generic type [KMP88], specifically with the classes $\mathcal{A}_{K_1, K_2; C_0, C_1}^\delta$. Cameron also observed that one may forbid cliques in the manner of Henson. When this is generalized a little more one gets the classes $\mathcal{A}_{K_1, K_2; C_0, C_1; \mathcal{S}}^\delta$ where \mathcal{S} is a set of side constraints slightly more general than Henson’s cliques.

Also noteworthy in this regard is the classification by Macpherson of the infinite, locally finite distance transitive graphs [Mph82], which occupies a privileged position among the exceptional entries in the catalog (slightly generalized). The finite case was dealt with in [Cam80],

Ongoing work by Amato and Macpherson [AMp10] sheds considerable light on the case of diameter 3, and similar methods appear to suffice to confirm the completeness of the catalog in diameter 3. In particular, with this small diameter the potentially problematic antipodal case can be handled directly and is covered in [AMp10]. The role of the generalized Henson constraints is very clear in that work, and is the main focus of that article. In the form I have seen, that article treats the antipodal and triangle free cases.

Some examples of metrically homogeneous graphs can be extracted from tables given at the end of [Che98]. These tables present all the primitive metrically homogeneous graphs of diameter 3 or 4 which can be defined by forbidding a set of triangles, excluding those of diameter 3 in which none of the forbidden triangles involve the distance 2 (where there is a notion of free amalgamation which seemed

not very interesting). At the time I produced those tables, it never crossed my mind that a significant number of them could be construed as metric spaces. But the 27 examples originally listed with 4 nontrivial self-paired orbits on pairs (i.e., 4 nontrivial symmetric 2-types) give rise to 20 distinct metric spaces, each derived from a metrically homogeneous graph by taking distance 1 as the edge relation. Some of the examples listed have no interpretation as metric spaces, while others can be interpreted as metric spaces in two distinct ways. Some of these examples lie outside the catalog based on the families $\mathcal{A}_{K,C,S}^\delta$ but are consistent with the catalog as it now stands, as indeed they must be, since we have proved the completeness of our catalog in its present form for examples determined entirely by constraints of order 3. We will tabulate the relevant examples and give a translation from the notation of [Che98] to the notation of our catalog.

There are good methods, but there is no clear strategy, for an eventual proof of completeness of the catalog. In Part I we have indicated some of the methods which have been applied to similar problems, notably in the classification of the homogeneous connected graphs by Lachlan and Woodrow. Since the homogeneous connected graphs are the metrically homogeneous connected graphs of diameter at most 2, the Lachlan/Woodrow classification appears as the point of departure for Cameron's problem. In the case of diameter 2 the fact that the triangle inequality is vacuous is helpful when applying Lachlan's method. With higher diameters the emphasis shifts toward the use of the triangle inequality, but Lachlan's Ramsey argument (discussed in Part I) retains considerable power.

After presenting our notions of exceptional and generic type metrically homogeneous graphs, we will state the main facts known to us regarding the completeness of the catalog, and then present the catalog itself. The remainder of this Part will then deal with proofs of our results in the exceptional case, as well as some discussion of the generic case, in considerably less detail.

Most of the metrically homogeneous graphs considered are taken to be connected, a point occasionally mentioned. In an inductive analysis, disconnected metrically homogeneous subgraphs may come into play. In such cases we focus mainly but not exclusively on their connected components.

5.2. Exceptional Cases. Our basic strategy in designing the catalog presented below was to try to ensure that the following three types of connected metrically homogeneous graph Γ were all adequately covered.

Exceptional: The induced graph on the neighbors of a fixed vertex of Γ is exceptional (in a sense specified below).

Imprimitive: The graph Γ is imprimitive; that is, it carries a nontrivial equivalence relation invariant under $\text{Aut}(\Gamma)$.

3-constrained: The class of finite metric spaces which embed isometrically into Γ can be specified in terms of forbidden substructures of order 3.

Let us take up these three possibilities separately. We will mention various results about them along the way, but leave the precise statements to follow the presentation of the catalog.

First, the class of connected metrically homogeneous graphs which we officially declare to be "exceptional" is defined as follows. Let Γ be a metrically homogeneous graph, and $v \in \Gamma$ a fixed vertex. Let $\Gamma_1 = \Gamma_1(v)$ be the graph induced on the neighbors of v . Then Γ_1 is a homogeneous graph, and its isomorphism type is

independent of the choice of base vertex v . As Γ_1 is homogeneous, it can be found in the list of Lachlan/Woodrow presented in §4.2, which we should keep in mind throughout.

The cases in which the induced graph Γ_1 is an independent set, a Henson graph, or the Rado graph are all associated with natural constructions of Fraïssé type and do not belong on the exceptional side. We will put all other possibilities for Γ_1 on the exceptional side: these are the imprimitive or finite cases, and the complements of the Henson graphs. More abstractly, they are the cases in which Γ_1 is either imprimitive or contains no infinite independent set. We give the classification of the metrically homogeneous graphs Γ with Γ_1 exceptional as Theorem 10. These turn out to be of familiar types.

The second class of graphs we looked into were the imprimitive graphs. Here the situation is at first quite close to what is known in the finite case for distance transitive graphs, where it goes under the name of Smith's Theorem [AH06, Smi71], though eventually the analysis diverges, losing some ground in the infinite case, but with the much stronger hypothesis of metric homogeneity providing considerable compensation. It is easy to show, as in the finite case, that the imprimitive metrically homogeneous graphs are either bipartite, or *antipodal*, or possibly both. Here a graph is antipodal if its diameter δ is finite, and the relation $d(x, y) = \delta$ (or 0) defines an equivalence relation. For metrically homogeneous graphs it turns out that with minor exceptions the equivalence classes in the antipodal case have order 2, and that the pairing defined by the relation $d(x, y) = \delta$ defines an involutory automorphism.

This analysis is useful but does not lead to a complete classification for reasons that will become quite clear in a moment. There is a general reduction from the bipartite case to the nonbipartite (but possibly antipodal) case. One considers the graph $B\Gamma$ induced on one half of the bipartition by taking as the edge relation $d(x, y) = 2$. This is then a metrically homogeneous graph in its own right, and we may consider the case in which $B\Gamma$ is exceptional. In fact, we give a complete classification of the bipartite metrically homogeneous graphs for which $B\Gamma_1$ is not the Rado graph as Theorem 13. We also characterize the antipodal bipartite metrically homogeneous graphs of odd diameter in Theorem 12. It would also be reasonable at this point to try to complete the reduction of the bipartite case to the nonbipartite (possibly antipodal) case by proving the existence and uniqueness of Γ with $B\Gamma$ falling under the remaining cases; or, if necessary, at least to prove that result for $B\Gamma$ falling within the catalog of known examples.

The third class of metrically homogeneous graphs requiring special attention leads us to a conjectured description of the graphs of generic type. Informally, the graphs of generic type are those which come from amalgamation classes using natural methods of amalgamation approximating some notion of free amalgamation. In our case we are dealing with classes of metric spaces with an integer valued metric, typically with a bound δ on the diameter, and for which any geodesic of total length at most δ is allowed to occur. In that case, if the class \mathcal{A} in question has the amalgamation property, the associated homogeneous metric space carries the graph structure given by the edge relation $d(x, y) = 1$, and the metric coincides with the graph metric (a point made in [Cam98]). As the class \mathcal{A} must be hereditary (downward closed), it may be specified by giving a set of forbidden subspaces; that is, we specify a collection \mathcal{C} of finite metric spaces, and then let \mathcal{A} be the class of

finite metric spaces X such that no space in \mathcal{C} embeds isometrically into X . We single out for attention the case in which all constraints in \mathcal{C} have order 3. We note that the triangle inequality is already a set of constraints of this type. Furthermore, imprimitivity is itself an example of a constraint given by a set of forbidden metric triangles.

Our first idea was to consider the classes $\mathcal{A}_{K;C}^\delta$ given by constraints of the following type: there are no triangles of odd perimeter less than $2K + 1$, and no triangles of perimeter C or more. Here the perimeter is the sum of the three distances between pairs of points (and the term “triangle” refers to metric triangles, that is, to metric spaces with three points). This idea is suggested by known examples in the theory of universal graphs described by Konjath, Mekler, and Pach [KMP88], and a proof of the amalgamation property in such cases is straightforward. We may combine this construction with the idea of omitting cliques as follows. With δ the diameter, let a $(1, \delta)$ -space be any finite metric space in which only the distances 1 and δ occur. We will be concerned only with the case in which $\delta \geq 3$, in which case such a space consists of equivalence classes which are cliques with respect to the edge relation $d(x, y) = 1$, with distinct classes separated by the maximal distance δ . It turns out that one may usually avoid both the minimal and the maximal values of the distance in completing amalgamation diagrams, and thus we can generalize the Henson construction to get amalgamation classes of the form $\mathcal{A}_{K;C;\mathcal{S}}$ with \mathcal{S} a set of $(1, \delta)$ -spaces. Here one must pay a little attention to the interaction of K , C , and \mathcal{S} to ensure that the amalgamation property holds.

As a test of our original catalog, we set out to prove that the pattern of triangles occurring in any metrically homogeneous graph would be that of one of the classes $\mathcal{A}_{K;C}^\delta$. Once we saw that this was false it was not immediately clear whether the exceptions were sporadic. In the end the class with five parameters $\mathcal{A}_{K_1, K_2; C_0, C_1}^\delta$ emerged along with its Hensonian variations $\mathcal{A}_{K_1, K_2; C_0, C_1; \mathcal{S}}^\delta$, and a further antipodal variant, by solving a much simpler problem: identify all amalgamation classes determined by constraints of order 3. It still needs to be shown that the collection of triangles not embedding in an arbitrary metrically homogeneous graph agrees with the collection associated with some graph in our extended catalog. We would view this as a significant step toward a proof of completeness of the catalog.

The case of imprimitive graphs has been analyzed to the point at which the remaining cases should fall under the generic case, corresponding to extreme values of the parameters: in the bipartite case we forbid all triangles of odd perimeter, and in the antipodal case we forbid all triangles of perimeter greater than 2δ (as we shall see). But the antipodal case involves another variation on the Henson construction, so it appears separately in the catalog.

5.3. The Catalog.

NOTATION 5.1. If Γ is metrically homogeneous then for $v \in \Gamma$ we denote by $\Gamma_i(v)$ the set of vertices at distance i from v , with the induced metric; this is a homogeneous metric space, but it does not necessarily come from a graph metric, and in fact the distance 1 may not even be represented in $\Gamma_i(v)$. Since the isomorphism type of $\Gamma_i(v)$ is independent of the choice of v , we often write Γ_i rather than $\Gamma_i(v)$.

If the distance 1 is represented in Γ_i and Γ_i is connected, then the metric on Γ_i is the graph metric (see [Cam98]).

Our catalog uses specialized notations and constructions which will be explained in §5.4.

CATALOG

- I. $\delta \leq 2$ (Cf. §4.2)
 - (a) Finite primitive: $C_5, L[K_{3,3}]$
 - (b) Defective or imprimitive: $m \cdot K_n, K_m[I_m]$.
 - (c) Infinite primitive, not defective: G_n, G_n^c, G_∞ .
- II. $\delta \geq 3, \Gamma_1$ finite or imprimitive.
 - (a) An n -gon with $n \geq 6$.
 - (b) Antipodal double cover of one of the graphs $C_5, L[K_{3,3}]$, or a finite independent set.
 - (c) A tree-like graph $T_{r,s}$ as described by Macpherson in [Mph82], where $2 \leq r, s \leq \infty$, and if $s = \infty$ then $r \geq 3$.
- III. Γ_1 infinite and primitive
 - (a) $T_{2,\infty}$, the infinitely branching regular tree.
 - (b) The generic antipodal graph omitting K_n : $\Gamma_{a,n}^\delta$, where either $\delta \geq 4$, or $\delta = 3$ and $n = 3$ or ∞ .
 - (c) The generic graph $\Gamma_{K_1, K_2; C_0, C_1; \mathcal{S}}^\delta$ associated with the class $\mathcal{A}_{K_1, K_2; C_0, C_1; \mathcal{S}}^\delta$, for an admissible choice of parameters $K_1, K_2; C_0, C_1; \mathcal{S}$.

The next order of business is to explain the following notions referred to in the catalog, after which we will state the main results relating to completeness of the catalog.

- (1) The treelike graphs $T_{r,s}$ of [Mph82] (generalized to allow r or s to be infinite);
- (2) The antipodal double cover;
- (3) The notation

$$\Gamma_{K_1, K_2; C_0, C_1; \mathcal{S}}^\delta$$

for graphs constructed via the Fraïssé theory from a suitable amalgamation class, and the antipodal variations denoted by $\Gamma_{a,n}^\delta$, as well as the precise conditions for admissibility of the parameters.

5.4. Three Constructions.

DEFINITION 5.2 (The treelike graphs $T_{r,s}$). For $2 \leq r, s \leq \infty$, we may construct an r -tree of s -cliques $T_{r,s}$ as follows. Take a tree $T(r, s)$ partitioned into two sets of vertices A, B , so that each vertex of A has r neighbors, all in B , and each vertex of B has s neighbors, all in A . Consider the graph induced on A with edge relation given by “ $d(u, v) = 2$ ”. This is $T_{r,s}$ (and the corresponding graph on B is $T_{s,r}$).

LEMMA 5.3. *For any r, s the tree $T(r, s)$ is homogeneous as a metric space with a fixed partition into two sets, and the graph $T_{r,s}$ is metrically homogeneous.*

PROOF. For any finite subset A of a tree T , one can see that the metric structure on A induced by T determines the structure of the convex closure of A , the smallest subtree of T containing A . Given that, a map between two finite subsets of $T(r, s)$ that respects the partition will extend first to the convex closures and then to the whole of $T(r, s)$.

This applies in particular to the two halves of $T(r, s)$. □

Now we turn to some “doubling” constructions.

DEFINITION 5.4.

- (1) The *double cover* $\Gamma = 2 * G$ of G is the graph on $V(G) \times \mathbb{Z}_2$ with edges given by $(u, i) \sim (v, j)$ iff

$$\begin{cases} u \sim v & \text{if } i = j \\ u \not\sim v \text{ and } u \neq v & \text{if } i \neq j \end{cases}$$

- (2) The *antipodal double cover* $\Gamma = \hat{G}$ of G is the double cover of the graph G^* obtained from G by adding one additional vertex $*$ adjacent to all vertices of G .
- (3) Let Γ be a graph of diameter δ . The *bipartite double cover* of Γ is the graph with vertex set $V(\Gamma) \times \mathbb{Z}_2$ and edge relation \sim given by $(u, i) \sim (v, j)$ iff:

$$d(u, v) = \delta \text{ and } i \neq j$$

Finally, we give the explicit definition of the classes $\mathcal{A}_{K_1, K_2; C_0, C_1}^\delta$.

We will write \mathcal{M}^δ for the class of all finite, integer valued metric spaces in which all lengths are bounded by δ . A triangle is a metric space containing three points. The type of a triangle is the triple (i, j, k) of distances realized in the triangle, taken in any order. The perimeter of a triangle of type i, j, k is the sum $i + j + k$.

DEFINITION 5.5. For $1 \leq K_1 \leq K_2 \leq \delta$ (or $K_1 = \infty, K_2 = 0$) and for $2\delta + 1 \leq C_0, C_1 \leq 3\delta + 2$, we define

- (1) $\mathcal{A}_{K_1, K_2}^\delta$ is the subclass of \mathcal{M}^δ with forbidden triangles of types (i, j, k) with $P = i + j + k$ odd and either
- $P < 2K_1 + 1$; or
 - $P > 2K_2 + \min(i, j, k)$
- (2) $\mathcal{B}_{C_0, C_1}^\delta$ is the subclass of \mathcal{M}^δ with forbidden triangles of types (i, j, k) where $P = i + j + k$ satisfies

$$P \geq C_\ell, \quad P \equiv \ell \pmod{2}$$

- (3) $\mathcal{A}_{K_1, K_2; C_0, C_1}^\delta = \mathcal{A}_{K_1, K_2}^\delta \cap \mathcal{B}_{C_0, C_1}^\delta$.

DEFINITION 5.6. A choice of parameters $\delta, K_1, K_2, C_0, C_1, \mathcal{S}$ is *admissible* if the following conditions are satisfied, where we write $C = \min(C_0, C_1)$ and $C' = \max(C_0, C_1)$.

- $\delta \geq 3$
- Either $1 \leq K_1 \leq K_2 \leq \delta$ or $K_1 = \infty$ and $K_2 = 0$
- $2\delta + 1 \leq C < C' \leq 3\delta + 2$, with one of C, C' even and the other odd.
- \mathcal{S} is a set of finite $(1, \delta)$ -spaces of order at least 3, and one of the following combinations of conditions holds:

- (1) $K_1 = \infty$:

$$K_2 = 0, C_1 = 2\delta + 1, \text{ and}$$

$$\mathcal{S} \text{ is } \begin{cases} \text{empty} & \text{if } \delta \text{ is odd, or } C_0 \leq 3\delta \\ \text{a set of } \delta\text{-cliques} & \text{if } \delta \text{ is even, } C_0 = 3\delta + 2 \end{cases}$$

- (2) $K_1 < \infty$ and $C \leq 2\delta + K_1$:

$$C = 2K_1 + 2K_2 + 1, K_1 + K_2 \geq \delta, \text{ and } K_1 + 2K_2 \leq 2\delta - 1$$

If $C' > C + 1$ then $K_1 = K_2$ and $3K_2 = 2\delta - 1$.
 If $K_1 = 1$ then \mathcal{S} is empty.

(3) $K_1 < \infty$, and $C > 2\delta + K_1$:

$K_1 + 2K_2 \geq 2\delta - 1$ and $3K_2 \geq 2\delta$.
 If $K_1 + 2K_2 = 2\delta - 1$ then $C \geq 2\delta + K_1 + 2$.
 If $C' > C + 1$ then $C \geq 2\delta + K_2$.
 If $K_2 = \delta$ then \mathcal{S} cannot contain a triangle of type $(1, \delta, \delta)$.
 If $K_1 = \delta$ then \mathcal{S} is empty.
 If $C = 2\delta + 2$, then \mathcal{S} is empty.

Note that if $\delta = \infty$ then C_0, C_1 should be omitted and the $(1, \delta)$ -spaces are just cliques.

We claim of course that these admissibility conditions are precisely the conditions required on our parameters to ensure that the corresponding class is an amalgamation class. We will not prove this here, though we will present the amalgamation procedure which works when the parameters meet these conditions.

We also have some antipodal variations on the classes $\mathcal{A}_{K_1, K_2; C_0, C_1; \mathcal{S}}^\delta$ to present. The antipodal case falls under our formalism for appropriate values of the numerical parameters, but the additional constraint set \mathcal{S} is somewhat different. On one hand, the $(1, \delta)$ -spaces of order at least 3 with the distance δ present are already excluded by the antipodality condition, so we are only concerned with cliques. On the other hand, if K_n does not embed into an antipodal graph, then none of the graphs obtained by replacing k vertices of K_n by antipodal vertices can embed. So we make the following definition, which takes into account some further constraints when $\delta = 3$.

DEFINITION 5.7. Let $\delta \geq 4$ be finite and $2 \leq n \leq \infty$, or $\delta = 3$ and $n = \infty$. Then

- (1) $\mathcal{A}_a^\delta = \mathcal{A}_{1, \delta-1; 2\delta+2, 2\delta+1; \emptyset}^\delta$ is the set of finite integral metric spaces in which no triangle has perimeter greater than 2δ .
- (2) $\mathcal{A}_{a, n}^\delta$ is the subset of \mathcal{A}_a^δ containing no subspace of the form $I_2^{\delta-1}[K_k, K_\ell]$ with $k + \ell = n$; here $I_2^{\delta-1}$ denotes a pair of vertices at distance $\delta - 1$ and $I_2^{\delta-1}[K_k, K_\ell]$ stands for the corresponding composition, namely a graph of the form $K_k \cup K_\ell$ with K_k, K_ℓ cliques (at distance 1), and $d(x, y) = \delta - 1$ for $x \in K_k, y \in K_\ell$. In particular, with $k = n, \ell = 0$, this means K_n does not occur.

One may make a general observation about the form of the admissibility conditions in Definition 5.6. Take $\mathcal{S} = \emptyset$, so that we consider $\mathcal{A}_{K_1, K_2; C_0, C_1; \emptyset}^\delta$ as a 5-parameter family of classes of finite metric structures. Looking over the conditions on the five parameters $\delta, K_1, K_2, C_0, C_1$ given above, we observe the following.

The condition (*) " $\mathcal{A}_{K_1, K_2; C_0, C_1}^\delta$ is an amalgamation class"
 is expressible in Presburger arithmetic.

If we knew this fact on *a priori* grounds we would then interpret the admissibility conditions as the result of expressing the property (*) in quantifier free terms in a language suitable for quantifier elimination in Presburger arithmetic, namely a language permitting the formation of linear combinations of variables with integer

coefficients, and with predicates for congruence modulo a fixed integer (above, parity will suffice).

To cast some light on this heuristically, consider the following three conditions, whose precise meaning will require some elucidation.

- (C1) The family of constraints defining $\mathcal{A}_{K_1, K_2; C_0, C_1; \emptyset}^\delta$ is uniformly definable in Presburger arithmetic, in the parameters $\delta, K_1, K_2, C_0, C_1$.
- (C2) For any fixed k , the condition “The family $\mathcal{A}_{K_1, K_2; C_0, C_1; \emptyset}^\delta$ has the k -amalgamation property” is definable in Presburger arithmetic as a property of $\delta, K_1, K_2, C_0, C_1$.
- (C3) The condition “ $\mathcal{A}_{K_1, K_2; C_0, C_1}^\delta$ is an amalgamation class” is expressible in Presburger arithmetic as a property of $\delta, K_1, K_2, C_0, C_1$.

Let us clarify the meaning of these statements before considering their relationship. Since we vary δ , the language here is a binary language with predicates $R_i(x, y)$ for $i \in \mathbb{Z}$, $i \geq 0$. We consider only structures in which every ordered pair of points satisfies exactly one of the relations R_i (with R_0 the equality relation); we may restrict our attention to symmetric relations, but this is not essential.

The family of constraints defining $\mathcal{A}_{K_1, K_2; C_0, C_1; \emptyset}^\delta$ is the set of minimal finite structures not in the class, which in the case at hand are structures of order at most 3, and include all pairs x, y violating the bound δ , all failures of symmetry, and all triples violating the triangle inequality, as well as the more specific constraints associated with K_1, K_2, C_0, C_1 . We may identify a structure A whose universe is $\{1, \dots, s\}$ with the s^2 -tuple

$$(d(i, j) : 1 \leq i, j \leq s)$$

where we use the metric notation $d(i, j)$ to denote the unique subscript d such that $R_d(i, j)$ holds in A . With $n = s^2$ we arrange this n -tuple in a definite order. Then the constraints of order s associated with K_1, K_2, C_0, C_1 become a subset of \mathbb{N}^n and with s bounded we are dealing with a finite number of such sets. So we may say that the set of constraints defining $\mathcal{A}_{K_1, K_2; C_0, C_1; \emptyset}^\delta$ is uniformly definable in Presburger arithmetic from the parameters if the encoded constraint sets in \mathbb{N}^n are so definable (for $s = 2, 3$). This gives our first statement a definite sense.

In condition (C2) we consider the k -amalgamation property, which is the amalgamation condition for pairs of structures A_1, A_2 of orders $k' < k$, over a base of order $k' - 1$; in other words, we require the ability to complete amalgamation diagrams in which the relation of one pair of points a_1, a_2 remains to be determined, and the total number of points involved is at most k . In our second condition, we require that the set of parameter 5-tuples $(\delta, K_1, K_2, C_0, C_1)$ for which the associated class has the k -amalgamation property should be definable in Presburger arithmetic, for each fixed value of k .

The third condition is similar, but refers to the full amalgamation property.

These three conditions (C1)–(C3) are all satisfied by our family $\mathcal{A}_{K_1, K_2; C_0, C_1; \emptyset}^\delta$, which has the additional property that it contains every amalgamation class of finite metric spaces which is determined by a set of constraints of order at most 3 (and 3 seems to be something of a magic number in terms of capturing a good deal of the sporadic side in classifications of homogeneous structures for finite binary languages). The following questions arise: is this kind of definability to be expected, not only in the present instance, but more generally; and could our analysis be materially simplified using automated techniques to perform an appropriate quantifier

elimination in Presburger arithmetic? We do not have answers to these questions, but we point out the following: the first property is interesting, but evident in the case at hand; the first property implies the second; and the second makes the third seem highly probable. We add a word on each of these three points.

The uniform definability in Presburger arithmetic of the constraint sets associated with the classes $\mathcal{A}_{K_1, K_2; C_0, C_1; \emptyset}^\delta$ is evident by inspection of the definition, which is already in that form. In particular, the triangle inequality is expressed by a formula of Presburger arithmetic. We had expected to define an appropriate family of amalgamation classes by imposing conditions of triangles involving only the size and parity of their perimeters, but this was inadequate. Our first try, a 3-parameter family $\mathcal{A}_{K, C}^\delta$, had that simple form, but did not cover all relevant amalgamation classes. The modification required in going to the 5-parameter family brought an additional term of Presburger arithmetic into the definition. We do *not* have any heuristic which explains why we are able to define a single family of classes uniformly in Presburger arithmetic, in terms of a bounded number of numerical parameters, which captures all amalgamation classes defined by constraints of order at most 3.

Now let us consider the implication $(C1) \implies (C2)$. Consider the amalgamation property for amalgams of order *precisely* k in which the desired structure is completely determined apart from the relations holding between one pair of elements a, b . This amalgamation property is directly expressed by a formula of the following form, with $n = k^2$:

$$\forall x_1 \forall x_2 \cdots \forall x_{n-2} \exists y_1 \exists y_2 \phi(x_1, \dots, x_{n-2}, y_1, y_2, \delta, K_1, K_2, C_0, C_1)$$

where ϕ simply states that the corresponding structure satisfies the required constraints, which will be expressible in Presburger arithmetic if condition $(C1)$ holds. Of course in symmetric structures we will require $y_1 = y_2$ and we may dispense with one existential quantifier (and several universal ones).

The displayed formula is very far from a quantifier free condition, and is not really useful for our purposes until it is expressed directly in terms of the parameters $\delta, K_1, K_2, C_0, C_1$. The admissibility conditions on the parameters given in Definition 5.6 express the 5-amalgamation property for the case of our particular family $\mathcal{A}_{K_1, K_2; C_0, C_1; \emptyset}^\delta$. Indeed, in deriving these conditions we considered only certain key amalgamation diagrams involving at most 5 elements; and eventually we found that we had enough conditions to derive the full amalgamation property. Thus the analysis that produced the admissibility conditions showed that they are equivalent to k -amalgamation for all $k \geq 5$. So at that point the obvious condition $(C1)$ explained the general form of our conditions, corresponding to $k = 5$ in condition $(C2)$, and we understood that our work to that point could be construed as carrying through an elimination procedure for a specific formula of Presburger arithmetic.

Of course, to pass from $(C2)$ to $(C3)$ in general requires finding a value of k_0 such that the quantifier-free form of condition $(C2)$ is independent of k for all $k \geq k_0$. It seems intrinsically reasonable that when we have a bound on the sizes of the constraints involved, there is also a bound k such that k -amalgamation implies full amalgamation. So there is some expectation that $(C2)$ will lead to $(C3)$, and thus that $(C1)$ will lead to $(C3)$.

Bearing in mind that Presburger arithmetic is decidable while diophantine problems over \mathbb{Z} are in general undecidable, we may wonder whether the general

classification problem for homogeneous structures in a finite relational language involves conditions which are largely expressible in Presburger arithmetic, or perhaps conditions of a more general diophantine type. One approach to this question would be to look more carefully at the examples tabulated in §14, giving 27 amalgamation classes determined by constraints on triangles whose Fraïssé limit is primitive, of which an impressive proportion (17) can be interpreted as metrically homogeneous graphs. We have just described a natural embedding of those 17 metrically homogeneous graphs into a 5-parameter family of graphs characterized by conditions expressible in Presburger arithmetic. One might ask for a natural extension of this 5-parameter family to another one given by similar conditions, in which all 27 examples for the case of 4 symmetric 2-types are found. Observe however that the collection of homogeneous structures for a given language is invariant under permutations of the language, while a useful definable relation in Presburger arithmetic should not be invariant under arbitrary permutations of the base set! So there is a prior issue of “symmetry breaking.” Even the 17 examples which can be interpreted as metrically homogeneous graphs give rise to 20 examples after breaking the symmetry (§10.1).

5.5. Statement of Results. Our first two theorems can be taken as a summary of the way the catalog was constructed.

THEOREM 9. *Let $\delta \geq 3$, $1 \leq K_1 \leq K_2 \leq \delta$ or $K_1 = \infty$, $2\delta + 1 \leq C_0, C_1 \leq 3\delta + 2$ with C_0 even and C_1 odd, and \mathcal{S} a set of $(1, \delta)$ -spaces occurring in $\mathcal{A}_{K_1, K_2; C_0, C_1}^\delta$. Then the class*

$$\mathcal{A}_{K_1, K_2; C_0, C_1; \mathcal{S}}^\delta$$

is an amalgamation class if and only if the parameters $\delta, K_1, K_2, C_0, C_1, \mathcal{S}$ are admissible, or $K_1 = 1$, $K_2 = \delta - 1$, $C_0 = 2\delta + 2$, and $C_1 = 2\delta + 1$, and $\mathcal{S} = \emptyset$.

There is a point of notation that needs to be explained here, as it is possible for two sets of parameters to define the same class, with one set admissible and the other not. We make canonical choices of parameters, taking for example K_1 to be the least k for which there is an odd triangle of perimeter $2k + 1$, and K_2 the greatest, with values $K_1 = \infty$ and $K_2 = 0$ if there is none. In particular if $C_1 = 2\delta + 1$ this forces $K_2 < \delta$. One defines C_0, C_1 similarly as the least values strictly above 2δ such that perimeters of appropriate parity above the given bound are forbidden. Thus if there are no odd triangles, then $C_1 = 2\delta + 1$. And of course, we take only minimal constraints in \mathcal{S} which are not consequences of the other constraints.

Note that when $K_1 = 1$, $K_2 = \delta - 1$, $C_0 = 2\delta + 2$, and $C_1 = 2\delta + 1$, we have the condition $\mathcal{S} = \emptyset$ here, but we also have the amalgamation classes $\mathcal{A}_{a, n}^\delta$ with the same parameters K_1, K_2, C_0, C_1 and a modified constraint set \mathcal{S} .

THEOREM 10. *Let Γ be a connected metrically homogeneous graph and suppose that Γ_1 is either imprimitive or contains no infinite independent set. Then Γ_1 is in our catalog under case I or II.*

We will prove Theorem 10 in §8, but our discussion of Theorem 9 in §6 is confined primarily to a detailed description of the amalgamation procedure. It takes some additional calculation to see that this procedure produces a metric meeting the required constraints in all cases.

Now let us turn to the case in which Γ is imprimitive. The first point is a very general fact which is quite familiar in the finite case but does not depend on finiteness, and which is also discussed in [AMp10].

FACT (cf. [AH06, Theorem 2.2]). *Let Γ be an imprimitive connected distance transitive graph of degree at least 3. Then Γ is either bipartite or antipodal (possibly both).*

The next result goes some distance toward bringing the antipodal case under control. It is proved in §7.2.

THEOREM 11. *Let Γ be a connected metrically homogeneous and antipodal graph, of diameter $\delta \geq 3$. Then for each vertex $u \in \Gamma$, there is a unique vertex $u' \in \Gamma$ at distance δ from u , and we have the “antipodal law”*

$$d(u, v) = \delta - d(u', v) \text{ for } u, v \in \Gamma$$

In particular, the map $u \mapsto u'$ is an automorphism of Γ .

In the study of finite distance transitive graphs, there are good reductions in both of the imprimitive cases—bipartite as well as antipodal—back (eventually) to the primitive case. The reduction in the bipartite case is straightforward in our context as well, but we will see that there is no simple reduction in the antipodal case, in the category of metrically homogeneous graphs. Still Theorem 11 suggests that the situation can be viewed as a modest elaboration beyond the primitive case.

We made some further explorations of the bipartite case, emphasizing extreme behavior, so as not to overlook anything obvious that would belong in the catalog. This led to the following results.

THEOREM 12. *Let Γ be a connected metrically homogeneous graph of odd diameter $\delta = 2\delta' + 1$ which is both antipodal and bipartite. Then $B\Gamma$ is connected, and Γ is the bipartite double cover of $B\Gamma$. The graph $B\Gamma$ is a metrically homogeneous graph with the following properties:*

- (1) $B\Gamma$ has diameter δ' ;
- (2) No triangle in $B\Gamma$ has perimeter greater than $2\delta' + 1$;
- (3) $B\Gamma$ is not antipodal.

Conversely, for any metrically homogeneous graph G with the three stated properties, there is a unique antipodal bipartite graph Γ of diameter $2\delta' + 1$ such that $B\Gamma \cong G$.

By inspection we find the following.

COROLLARY 12.1. *Let Γ be a connected metrically homogeneous graph of odd diameter $\delta = 2\delta' + 1$ which is both antipodal and bipartite. Suppose that $B\Gamma$ is in the catalog. Then Γ is in the catalog, and the relevant pairs $(B\Gamma, \Gamma)$ are as follows.*

- (1) $B\Gamma$ complete of order n , $3 \leq n \leq \infty$, and Γ the bipartite double of an independent set of order $n - 1$ (the bipartite complement of a perfect matching between two sets of order n).
- (2) (C_n, C_{2n}) with C_n an n -cycle, and n odd.
- (3) $(G_3^c, \Gamma_{\infty, 0; 12, 13; \emptyset}^5)$.
- (4) $(\Gamma_{1, \delta'; 2\delta'+2, 2\delta'+3; \emptyset}^{\delta'}, \Gamma_{\infty, 0; 2\delta'+2, 2\delta'+1; \emptyset}^{\delta'})$ with $\delta' \geq 3$.

In particular, for Γ antipodal and bipartite of odd diameter $\delta \leq 5$, since the Lachlan/Woodrow classification is complete through diameter 2, the corollary applies (Corollary 12.2, §7.2).

Another natural case in the bipartite context is the following. The structure of Γ_1 tells us nothing in the case of a bipartite graph when Γ_1 is infinite, but we may single out for special attention the metrically homogeneous bipartite graphs for which Γ_1 is infinite while $B\Gamma$ is itself exceptional. Here too everything conforms to the catalog.

THEOREM 13. *Let Γ be a connected, bipartite, and metrically homogeneous graph, of diameter at least 3, and degree at least 3, and with Γ_1 infinite. Then either $B\Gamma_1$ is isomorphic to the Rado graph, or $B\Gamma$ and Γ are in the catalog under one of the following headings.*

- (1) $B\Gamma \cong T_{\infty, \infty}$, and Γ is an infinitely branching tree $T_{2, \infty}$.
- (2) $B\Gamma \cong K_{\infty}$, and Γ has diameter 3, with Γ either the complement of a perfect matching, or the generic bipartite graph $\Gamma_{\infty, 0; 10, 7; \emptyset}^3$.
- (3) $B\Gamma \cong K_{\infty}[I_2]$, Γ has diameter 4, and $\Gamma \cong \Gamma_{\infty, 0; 10, 9; \emptyset}^4$ is the generic antipodal bipartite graph of diameter 4.
- (4) $B\Gamma \cong G_n^c$, the complement of a Henson graph, for some $n \geq 3$, and $\Gamma \cong \Gamma_{\infty, 0; 14, 9; \{I_n^{(4)}\}}^4$ the generic bipartite graph in which there is no set of n vertices which are pairwise at distance 4.
- (5) $B\Gamma \cong G_3^c$, $\Gamma \cong \Gamma_{\infty, 0; 12, 11; \emptyset}^5$ antipodal bipartite of diameter 5 as in Corollary 12.1.

One could explore the imprimitive case more, and perhaps find something else of a similar sort that needs to be added to the catalog. But after that the question becomes how to approach a proof of its completeness. This is governed by the structure of the catalog. To show that the catalog is complete, we would need to show on the one hand that any example not falling under one of the exceptional headings has the same constraints on triangles as one of our standard amalgamation classes, and then that for each metrically homogeneous graphs with one of the specified patterns of forbidden triangles, any further minimal forbidden configurations must be $(1, \delta)$ -spaces (or the analogous constraints appropriate to the antipodal case). Some instances of the latter type of analysis were given in [AMP10].

It is no doubt useful to divide the classification problem into the two parts just mentioned, but both parts appear challenging, and they are likely to become intertwined in an inductive treatment of the problem.

In proving the completeness of the catalog one first takes δ finite. There is something more to be checked when δ is infinite. Once one has a degree of control on the triangles involved it becomes easy to reduce from the case of infinite diameter to finite diameter. And that degree of control may well follow quickly from the classification in finite diameter, but we have not looked closely at that.

The rest of this Part is devoted to the proofs of Theorems 10 and 11–13. In the case of Theorem 9 we will confine ourselves to an indication of the amalgamation procedure used when the parameters are admissible.

6. Constructions of metrically homogeneous graphs

We are concerned here with the use of Fraïssé constructions to produce metrically homogeneous graphs of generic type, including some bipartite and antipodal cases.

6.1. The Main Construction. Recall that in the definition of the class $\mathcal{A}_{K_1, K_2; C_0, C_1; \mathcal{S}}^\delta$ of finite integral metric spaces of diameter δ with constraints K_1 , K_2 , C_0 , C_1 , and \mathcal{S} , the parameter K_1 controls the presence of triangles with small odd perimeter, the parameters C_0, C_1 control those of large perimeter, \mathcal{S} is a set of $(1, \delta)$ -spaces, and the parameter K_2 is used to further control the presence of odd triangles, excluding all those whose perimeter P satisfies

$$P > 2K_2 + d(a, b)$$

for some choice of a, b in the triangle. If $K_2 = \delta$ then this condition is vacuous. The admissible combinations of these parameters are described in Definition 5.6. In that presentation the parity of C_0 and C_1 plays less of a role than their relative sizes, so we introduce the notation $C = \min(C_0, C_1)$ and $C' = \max(C_0, C_1)$. The simplest cases are those in which $C' = C + 1$: we exclude all triangles of perimeter C or greater. We use the values $3\delta + 1$ or $3\delta + 2$ (with appropriate parity) to indicate that the constraints corresponding to C_0 or C_1 are absent.

The main result concerning amalgamation for classes of type $\mathcal{A}_{K_1, K_2; C_0, C_1; \mathcal{S}}^\delta$ can then be phrased as follows.

THEOREM (9). *Let $\delta \geq 3$, $1 \leq K_1 \leq K_2 \leq \delta$ or $K_1 = \infty$, $2\delta + 1 \leq C_0, C_1 \leq 3\delta + 2$ with C_0 even and C_1 odd, and \mathcal{S} a set of $(1, \delta)$ -spaces occurring in $\mathcal{A}_{K_1, K_2; C_0, C_1}^\delta$. Then the class*

$$\mathcal{A}_{K_1, K_2; C_0, C_1; \mathcal{S}}^\delta$$

is an amalgamation class if and only if the parameters $\delta, K_1, K_2, C_0, C_1, \mathcal{S}$ are admissible.

There is also a variation associated with the antipodal case and a modified constraint set \mathcal{S} (Theorem 14, below).

The proof that the admissibility constraint is *necessary* (for classes specified in this way) requires many constructions of specific amalgamation diagrams, mostly of order 5, which have no completion when the parameters lie outside the admissible range. The proof that in the admissible cases amalgamation can be carried out involves the specification of a particular amalgamation procedure varying in detail according to the different cases occurring within the definition of admissibility, and depending further on some details of the amalgamation problem itself, followed by additional computation to verify that the proposed procedure actually works. We will describe the procedure itself in detail and leave the rest for another occasion.

For the proof of amalgamation, it suffices to consider amalgamation diagrams of the special form $A_1 = A_0 \cup \{a_1\}$, $A_2 = A_0 \cup \{a_2\}$, that is with only one distance $d(a_1, a_2)$ needing to be determined; we call these *2-point* amalgamations. Furthermore, any distance lying between

$$d^-(a_1, a_2) = \max_{x \in A_0} (|d(a_1, x) - d(a_2, x)|)$$

and

$$d^+(a_1, a_2) = \min_{x \in A_0} (d(a_1, x) + d(a_2, x))$$

will give at least a pseudometric (and if $d^-(a_1, a_2) = 0$, we will identify a_1 and a_2).

When $C' = C + 1$ we have the requirement that all triangles have perimeter at most $C - 1$ and therefore we consider a third value

$$\tilde{d}(a_1, a_2) = \min_{x \in A_0} (C - 1 - [d(a_1, x) + d(a_2, x)])$$

In this case the distance $i = d(a_1, a_2)$ must satisfy:

$$d^-(a_1, a_2) \leq i \leq \min(d^+(a_1, a_2), \tilde{d}(a_1, a_2))$$

Similarly when $K_1 = \infty$ then as there are no odd triangles we modify the definition of \tilde{d} as follows:

$$\tilde{d}(a_1, a_2) = \min_{x \in A_0} (C_0 - 2 - [d(a_1, x) + d(a_2, x)])$$

The amalgamation procedure is then given by the following rules for completing a 2-point amalgamation diagram $A_i = A_0 \cup \{a_i\}$ ($i = 1, 2$). We will assume throughout that $d^-(a_1, a_2) > 0$, as otherwise we may simply identify a_1 with a_2 . We will also write:

$$i^- = d^-(a_1, a_2), i^+ = d^+(a_1, a_2), \tilde{i} = \tilde{d}(a_1, a_2)$$

and we seek a suitable value for $i = d(a_1, a_2)$.

(1) If $K_1 = \infty$:

Then the parity of $d(a_1, x) + d(a_2, x)$ is independent of the choice of $x \in A_0$.

If \mathcal{S} is empty then any value i with $i^- \leq i \leq i^+$ and of the correct parity will do.

If \mathcal{S} is nonempty (and irredundant) then \mathcal{S} consists of a δ -clique, and δ is even. In particular $\delta \geq 4$. In this case take $d(a_1, a_2) = i$ with $1 < i < \delta$ and with i of the correct parity. We may take $i = i^+$ if this is less than δ , $i = i^-$ if this is greater than 1, and otherwise, when $i^- \leq 1$ and $i^+ \geq \delta$, take $i = 2$ or 3 of the correct parity.

(2) If $K_1 < \infty$ and $C \leq 2\delta + K_1$:

(a) If $C' = C + 1$ then:

(i) If $\min(i^+, \tilde{i}) \leq K_2$ let $d(a_1, a_2) = \min(i^+, \tilde{i})$. Otherwise:

(ii) If $i^- \geq K_1$ let $d(a_1, a_2) = i^-$. Otherwise:

(iii) Let $d(a_1, a_2) = K_2$.

(b) If $C' > C + 1$ then:

(i) If $i^+ < K_2$ let $d(a_1, a_2) = i^+$. Otherwise:

(ii) If $d^- > K_2$ let $d(a_1, a_2) = i^-$. Otherwise:

(iii) Take $d(a_1, a_2) = K_2$ unless there is $x \in A_0$ with $d(a_1, x) = d(a_2, x) = \delta$, in which case take $d(a_1, a_2) = K_2 - 1$.

(3) If $K_1 < \infty$ and $C > 2\delta + K_1$:

(a) If $i^- > K_1$, let $d(a_1, a_2) = i^-$.

- (b) Otherwise:
- (i) If $C' = C + 1$:
 - (A) If $i^+ \leq K_1$ let $d(a_1, a_2) = \min(i^+, \tilde{i})$. Otherwise:
 - (B) Let $d(a_1, a_2) = K_1$ unless we have one of the following:
 - There is $x \in A_0$ with $d(a_1, x) = d(a_2, x)$,*
 - and $K_1 + 2K_2 = 2\delta - 1$; or $K_1 = 1$.*
 In these cases, take $d(a_1, a_2) = K_1 + 1$.
 - (ii) If $C' > C + 1$:
 - If $i^+ < K_2$ let $d(a_1, a_2) = i^+$.
 - Otherwise, let $d(a_1, a_2) = \min(K_2, C - 2\delta - 1)$.

We note some extreme cases. With $K_1 = \infty$ we are dealing with bipartite graphs; with $C = 2\delta + 1$ we are dealing with antipodal graphs. With $K_1 > 1$ we have the case $\Gamma_1 \cong I_\infty$; with $K_1 = 1$ and $\mathcal{S} = \{K_n\}$, we have the case of Γ_1 the generic K_n -free graph.

LEMMA 6.1. *Let Γ be a metrically homogeneous graph of diameter δ . Then Γ is antipodal if and only if no triangle has perimeter greater than 2δ .*

PROOF. The bound on perimeter immediately implies antipodality.

For the converse, let (a, b, c) be a triangle with $d(a, b) = i$, $d(a, c) = j$, $d(b, c) = k$, and let a' be the antipodal point to a . Then the triangle (a', b, c) has distances $\delta - i$, $\delta - j$, and k , and the triangle inequality yields $i + j + k \leq 2\delta$, as claimed. \square

6.2. An Antipodal Variation. We consider modifications of our definitions which allow us to include constraints on cliques when the associated graph is antipodal and $K_1 = 1$. In this case the associated amalgamation class may require constraints which are neither triangles nor $(1, \delta)$ -spaces.

DEFINITION 6.2. Let $\delta \geq 4$ be finite and $2 \leq n \leq \infty$. Then

- (1) $\mathcal{A}_a^\delta = \mathcal{A}_{1, \delta-1; 2\delta+2, 2\delta+1; \emptyset}^\delta$ is the set of finite metric spaces in which no triangle has perimeter greater than 2δ .
- (2) $\mathcal{A}_{a,n}^\delta$ is the subset of \mathcal{A}_a^δ containing no subspace of the form $K_k \cup K_\ell$ with K_k, K_ℓ cliques, $k + \ell = n$, and $d(x, y) = \delta - 1$ for $x \in K_k, y \in K_\ell$. In particular, K_n does not occur.

THEOREM 14. *If $\delta \geq 4$ is finite and $2 \leq n \leq \infty$, then $\mathcal{A}_{a,n}^\delta$ is an amalgamation class. If $n \geq 3$ then the associated Fraïssé limit is a connected antipodal metrically homogeneous graph which is said to be generic for the specified constraints.*

Here the parameter n stands in place of the set \mathcal{S} ; since there are no triangles of perimeter greater than 2δ , the only relevant $(1, \delta)$ -spaces are 1-cliques. In these graphs Γ_1 is the generic graph omitting K_n .

We will give the proof of amalgamation for these particular classes in detail. The following lemma is helpful.

LEMMA 6.3. *Let δ be fixed, and let A be a finite metric space with no triangle of perimeter greater than δ . Then there is a unique “antipodal” extension \hat{A} of A , up to isometry, to a metric space satisfying the same condition, in which every vertex is paired with an antipodal vertex at distance δ , and every vertex not in A is antipodal to one in A .*

If A is in $\mathcal{A}_{a,n}^\delta$, then \hat{A} is in the same class.

PROOF. The uniqueness is clear: let

$$B = \{a \in A : \text{There is no } a' \in A \text{ with } d(a, a') = \delta\}$$

and introduce a set of new vertices $B' = \{b' : b \in B\}$. Let $\hat{A} = A \cup B'$ as a set. Then there is a unique symmetric function on \hat{A} extending the metric on A , with $d(x, b') = \delta - d(x, b)$ for $x \in \hat{A}$, $b \in B$.

So the issue is one of existence, and for that we may consider the problem of extending A one vertex at a time, that is to $A \cup \{b'\}$ with $b \in B$, as the rest follows by induction.

We need to show that the canonical extension of the metric on A to a function d on $A \cup \{b'\}$ is in fact a metric, satisfies the antipodal law for δ , and also satisfies the constraints corresponding to n .

The triangle inequality for triples (b', a, c) or (a, b', c) corresponds to the ordinary triangle inequality for (b, a, c) or the bound on perimeter for (a, b, c) respectively, and the bound on perimeter for triangles (a, b', c) follows from the triangle inequality for (a, b, c) .

Now suppose $n < \infty$ and b' belongs to a configuration $K_k \cup K_\ell$ with $k + \ell = n$ and $d(x, y) = \delta - 1$ for $x \in K_k$, $y \in K_\ell$. We may suppose that $b' \in K_k$; then $K_k \setminus \{b'\} \cup (K_\ell \cup \{b\})$ provides a copy of $K_{k-1} \cup K_{\ell+1}$ of forbidden type. \square

LEMMA 6.4. *If $\delta \geq 4$ is finite, $2 \leq n \leq \infty$, then $\mathcal{A}_{a,n}^\delta$ is an amalgamation class.*

PROOF. We consider a two-point amalgam with $A_i = A_0 \cup \{a_i\}$, $i = 1, 2$. If $d(a_2, x) = \delta$ for some $x \in A_0$ then there is a canonical amalgam $A_1 \cup A_2$ embedded in \hat{A}_1 . So we will suppose $d(a_i, x) < \delta$ for $i = 1, 2$ and $x \in A_0$. Then applying Lemma 6.3 to A_1 and A_2 , we may suppose that for every vertex $v \in A_0$ there is an antipodal vertex $v' \in A_0$.

We claim that any metric d on $A_1 \cup A_2$ extending the given metrics d^i on A_i will satisfy the antipodal law for δ . So with d such a metric, consider a triangle of the form (a_0, a_1, a_2) with $a_0 \in A_0$. By the triangle inequality for (a_1, a'_0, a_2) we have

$$d(a_1, a_2) \leq 2\delta - [d(a_1, a_0) + d(a_2, a_0)]$$

and this is the desired bound on perimeter.

We know by our general analysis that any value r for $d(a_1, a_2)$ with

$$d^-(a_1, a_2) \leq r \leq d^+(a_1, a_2)$$

will give us a metric, and in the bipartite case we will want r to have the same parity as $d^-(a_1, a_2)$ (or equivalently, $d^+(a_1, a_2)$).

To deal with the the constraints involving the parameter n , it is sufficient to avoid the values $r = 1$ and $r = \delta - 1$. But $d^+(a_1, a_2) > 1$, and $d^-(a_1, a_2) < \delta - 1$, so we may take r equal to one of these two values unless we have

$$d^-(a_1, a_2) = 1, d^+(a_1, a_2) \geq \delta - 1$$

In this case, we take r to be some intermediate value, and as $\delta > 3$, there is one. \square

7. Imprimitve Graphs

7.1. Smith's Theorem. We now turn to Smith's Theorem, which provides a general analysis of the imprimitive case, following [AH06] (cf. [BCN89, Smi71]). This result applies to imprimitive distance transitive graphs (that is, the homogeneity condition is assumed to hold for pairs of vertices), and even more generally

in the finite case. There are three points to this theory in the finite case: (1) the imprimitive graphs are of two extreme types, bipartite or antipodal; (2) associated with each type there is a reduction (folding or halving) to a potentially simpler graph; (3) with few exceptions, the reduced graph is primitive. Among the exceptions that need to be examined are the graphs which are both antipodal and bipartite. As our hypothesis of metric homogeneity is *not* preserved by the folding operation in general, we lose a good deal of the force of (2) and a corresponding part of (3). On the other hand, we will see that metric homogeneity implies that with trivial exceptions, in antipodal graphs the antipodal equivalence classes have order two. In that sense, antipodal graphs are not far from primitive, but there is no simple reduction to the primitive case, and the full classification is best thought of as a variation on the primitive case, to be handled by similar methods.

We will first take up the explicit form of Smith's Theorem given in [AH06], restricting ourselves to the distance transitive case. If Γ is a distance transitive graph, then any binary relation R invariant under $\text{Aut}(\Gamma)$ is a union of relations R_i defined by $d(x, y) = i$, $R = \bigcup_{i \in I} R_i$. We denote by $\langle t \rangle$ the union $\bigcup_{t|i} R_i$ taken over the multiples of t . The first point is the following.

FACT 7.1 (cf. [AH06, Theorem 2.2]). *Let Γ be a connected distance transitive graph of diameter δ , and let E be a congruence of Γ .*

- (1) $E = \langle t \rangle$ for some t .
- (2) If $2 < t < \delta$, then Γ has degree 2.
- (3) If $t = 2$ then either Γ is bipartite, or Γ is a complete regular multipartite graph, of diameter 2.

In particular, if the degree of Γ is at least 3, then Γ is either bipartite or antipodal (and possibly both).

Of course, the exceptional case of diameter 2 has already been noticed within the Lachlan/Woodrow classification, where it occurs as the complement of $m \cdot K_n$, with $m, n \leq \infty$.

If Γ is a connected distance transitive bipartite graph, we write $B\Gamma$ for the graph induced on either of the two equivalence classes for the congruence $\langle 2 \rangle$; these are isomorphic, with respect to the edge relation R_2 : $d(x, y) = 2$. This is called a *halved graph* for Γ , and Γ is a *doubling* of $B\Gamma$. If Γ is a connected distance transitive antipodal graph of diameter δ (necessarily finite), then $A\Gamma$ denotes the graph induced on the quotient Γ/R_δ by the edge relation: C_1 is adjacent to C_2 iff there are $u_i \in C_i$ with (u_1, u_2) an edge of Γ . This is called a *folding* of Γ , and Γ is called an *antipodal cover* of $A\Gamma$. In our context, the halving construction is more useful than the folding construction.

FACT 7.2 (cf. [AH06, Theorem 2.3]). *Let Γ be a connected metrically homogeneous bipartite graph. Then $B\Gamma$ is metrically homogeneous.*

PROOF. Since $\text{Aut}(\Gamma)$ preserves the equivalence relation whose classes are the two halves of Γ , the homogeneity condition is inherited by each half. \square

Some insight into the folding construction is afforded by the following.

LEMMA 7.3. *Let Γ be a connected distance transitive antipodal graph of diameter δ , and let C_1, C_2 be two equivalence classes for the antipodality relation R_δ . Then the set of distances $d(u, v)$ for $u \in C_1, v \in C_2$ is a pair of the form $\{i, \delta - i\}$ (which is actually a singleton if $i = \delta/2$ with δ even).*

PROOF. Since there are geodesics (u, v, u') with $d(u, v) = i$, $d(v, u') = \delta - i$, and $d(u, u') = \delta$, whenever we have $d(u, v) = i$ we also have $d(u', v) = \delta - i$ for some u' antipodal to u , by distance transitivity.

We claim that for $u \in C_1$, $v, w \in C_2$, with $d(u, v) = i$, $d(u, w) = j$ and $i, j \leq \delta/2$, we have $i = j$. If $i < j$, then $d(v, w) \leq i + j < \delta$ and hence $v = w$, a contradiction.

Thus for $u \in C_1$ the set of distances $d(u, v)$ with v in C_2 has the form $\{i, \delta - i\}$, and for v in C_2 the same applies with respect to C_1 , with the same pair of values. This implies our claim. \square

COROLLARY 7.3.1 (cf. [AH06, Proposition 2.4]). *Let Γ be a connected distance transitive antipodal graph of diameter δ , and consider the graph $A\Gamma$ as a metric space. If $u, v \in \Gamma$ with $d(u, v) = i$, then in $A\Gamma$ the corresponding points \bar{u}, \bar{v} lie at distance $\min(i, \delta - i)$.*

PROOF. Replacing v by v' with $d(u, v) = \delta - i$, we may suppose $i = \min(i, \delta - i)$. We have $d(\bar{u}, \bar{v}) \leq d(u, v)$.

Let $j = d(\bar{u}, \bar{v})$ and lift a path of length j from \bar{u} to \bar{v} to a walk (u, \dots, v^*) in Γ . If $v^* = v$ then $d(u, v) \leq j$ and we are done. Otherwise, $\delta = d(v, v^*) \leq d(v^*, u) + d(u, v) \leq j + i$, and as $i, j \leq \delta/2$, we find $i = j = \delta/2$. \square

This implies in particular that the folding of an antipodal metrically homogeneous graph of diameter at most 3 is complete, and that the folding of any connected distance transitive antipodal graph is distance transitive. But now consider the ‘‘generic’’ antipodal graph $\Gamma_{1, \delta-1; 2\delta+2, 2\delta+1; \emptyset}^\delta$ of diameter $\delta \geq 4$. Let $P = (u_0, \dots, u_\delta)$ be a geodesic in Γ , and let $C = (v_0, \dots, v_\delta, v_0)$ be an isometrically embedded cycle. On P we have $d(u_i, u_j) = |i - j|$ and on C we have $d(v_i, v_j) = \min(|i - j|, \delta - |i - j|)$, so the images of these graphs in $A\Gamma$ are isometric. We claim however that there is no automorphism of $A\Gamma$ taking one to the other.

Let $\delta_1 = \lfloor \delta/2 \rfloor$ and $\delta' = \lceil \delta/2 \rceil$. Then there is a vertex w_C in Γ at distance precisely δ' from every vertex of C . Hence the same applies in the folded graph $A\Gamma$ to the image of C . We claim that this does not hold for the image of P . Supposing the contrary, there would be a vertex w_P in Γ whose distance from each vertex of P is either δ' or $\delta - \delta'$. On the other hand, the distances $d(w_P, u_i)$, $d(w_P, u_{i+1})$ can differ by at most 1, and $\delta - \delta' > \delta' + 1$ since $\delta \geq 4$, so the distance $d(w_P, u_i)$ must be independent of i . However $d(u_0, w_P) = \delta - d(u_\delta, w_P)$, which would mean $\delta = 2\delta'$, while in fact $\delta > 2\delta'$.

Still, we can get a decent grasp of the antipodal case in another way.

7.2. The Antipodal Case. All graphs considered under this heading are connected and of finite diameter. Our main goal is the following.

THEOREM (11). *Let Γ be a connected metrically homogeneous and antipodal graph, of diameter $\delta \geq 3$. Then for each vertex $u \in \Gamma$, there is a unique vertex $u' \in \Gamma$ at distance δ from u , and we have the law*

$$d(u, v) = \delta - d(u', v)$$

for $u, v \in \Gamma$. In particular, the map $u \mapsto u'$ is an automorphism of Γ .

For $v \in \Gamma$, $\Gamma_i(v)$ denotes the graph induced on the vertices at distance i from v , and since the isomorphism type is independent of v , this will sometimes be denoted

simply by Γ_i , when the choice of v is immaterial. In particular, $\Gamma_\delta \cong I_n^{(\delta)}$, a set of n vertices mutually at distance δ . Our claim is that $n = 1$.

We begin with a variation of Lemma 7.3.

LEMMA 7.4. *Let Γ be metrically homogeneous and antipodal, of diameter δ . Suppose $u, u' \in \Gamma$, and $d(u, u') = \delta$. Then for $i < \delta/2$, the relation R_δ defines a bijection between $\Gamma_i(u)$ and $\Gamma_i(u')$, while $\Gamma_{\delta/2}(u) = \Gamma_{\delta/2}(u')$.*

PROOF. First suppose $i < \delta/2$, and $v \in \Gamma_i(u)$. We work with the equivalence classes C_1, C_2 of u and v respectively, with respect to the relation R_δ . As $d(u, v) = i$, $d(u, u') = \delta$, $i \leq \delta/2$, and $d(v, u') \in \{i, \delta - i\}$, we have $d(v, u') = \delta - i$.

Now (v, u') extends to a geodesic (v, u', v') with $d(v, v') = \delta$, $d(u', v') = i$, and we claim that v' is unique. If (v, u', v'') is a second such geodesic then we have $d(v, v') = d(v, v'') = \delta$ and $d(v', v'') \leq 2(\delta - i) < \delta$, so $v' = v''$.

Thus we have a well-defined function from $\Gamma_i(u)$ to $\Gamma_i(u')$, and interchanging u, u' we see that this is a bijection.

Now if $i = \delta/2$, apply Lemma 7.3: with $i = \delta/2$, the set $\{i, \delta - i\}$ is a singleton. \square

Taking $n > 1$, we will first eliminate some small values of δ .

LEMMA 7.5. *Let Γ be a metrically homogeneous antipodal graph, of diameter $\delta \geq 3$, and let $\Gamma_\delta \cong K_n^{(\delta)}$ with $1 < n \leq \infty$. Then $\delta \geq 5$.*

PROOF. First suppose that $\delta = 3$. Fix a basepoint in Γ . Define an equivalence relation \sim on Γ_2 by

$$x \sim y \text{ iff there is a vertex } u \in \Gamma_3 \text{ with } x, y \in \Gamma_1(u)$$

By homogeneity, any two vertices of Γ_2 at distance 2 are in the same class. If C_1, C_2 are two distinct equivalence classes in Γ_2 then each vertex of C_1 is at distance 3 from at most one vertex of C_2 , and is adjacent to the remainder. In particular there are some edges between C_1 and C_2 , and therefore there are none within C_1 or C_2 , in other words Γ_1 is an independent set, and Γ omits K_3 .

Let $a_1, a_2 \in \Gamma_3$ be distinct. We have shown that for any vertex u in $\Gamma_1(a_1)$ there is at most one vertex in $\Gamma_1(a_2)$ not adjacent to u . Taking a_3 a vertex at distance 3 from a_1, a_2 , we may take $u_1 \in \Gamma_1(a_1), u_2 \in \Gamma_1(a_2)$ adjacent. Then the number of vertices of $\Gamma_1(a_3)$ not adjacent to both u_1 and u_2 is at most 2, but since Γ_3 is triangle free, there are no vertices adjacent to both u_1 and u_2 , so $|\Gamma_1| \leq 2$. But then Γ is a cycle, and this contradicts $n > 1$.

So now suppose the diameter is 4. Take $a_1, a_2, a_3 \in \Gamma$ pairwise at distance 3, and $u_1, v_1 \in \Gamma_1(a_1)$ with $d(u_1, v_1) = 2$. Then the distances occurring in the triangle (a_3, u_1, v_1) are 3, 3, 2.

Consider $v_2 \in \Gamma_1(a_2)$ with $d(v_1, v_2) = 4$. As $d(u_1, v_1) = 2$ and $2 = \delta/2$ we find that $d(u_1, v_2) = 2$. Hence the triangles (a_3, u_1, v_1) and (a_3, u_1, v_2) are isometric. By metric homogeneity there is an isometry carrying (a_3, u_1, v_1, a_1) to (a_3, u_1, v_2, b_1) for some b_1 . Then $u_1, v_2 \in \Gamma_1(b_1)$ and $d(b_1, a_3) = 4$. But then $d(a_1, b_1), d(a_2, b_1) \leq 2$ while a_1, a_2, a_3, b_1 are all in the same antipodality class, forcing $a_1 = b_1 = a_2$, a contradiction. \square

Now we need to extend Lemma 7.3 to some cases involving distances which may be greater than $\delta/2$.

LEMMA 7.6. *Let Γ be metrically homogeneous and antipodal, of diameter $\delta \geq 3$, and let $\Gamma_\delta \cong K_n^{(\delta)}$ with $1 < n \leq \infty$. Suppose $d(a, a') = \delta$ and $i < \delta/2$. Suppose $u \in \Gamma_i(a)$, $u' \in \Gamma_i(a')$, with $d(u, u') = \delta$. If $v \in \Gamma_i(a)$ and $d(u, v) = 2i$, then $d(u', v) = \delta - 2i$.*

PROOF. We have $d(a, u') = \delta - i$. Take $v_0 \in \Gamma_i(a)$ so that (a, v_0, u') is a geodesic, that is $d(v_0, u') = \delta - 2i$. As $u, v_0 \in \Gamma_i(a)$ we have $d(u, v_0) \leq 2i$. On the other hand $d(u, u') = \delta$ and $d(v_0, u') = \delta - 2i$, so $d(u, v_0) \geq 2i$. Thus $d(u, v_0) = 2i$.

So we have at least one triple (a, u, v_0) with $v_0 \in \Gamma_i(a)$, $d(u, v_0) = 2i$, and with $d(v_0, u') = \delta - 2i$. Let (a, u, v) be any triple isometric to (a, u, v_0) . Then the quadruples (a, u, v, a') and (a, u, v_0, a') are also isometric since $u, v, v_0 \in \Gamma_i(a)$ with $i < \delta/2$. But as a, u together determine u' , we then have (a, u, v, a', u') and (a, u, v_0, a', u') isometric, and in particular $d(v, u') = \delta - 2i$. \square

After these preliminaries we can prove Theorem 11.

PROOF. We show that $n = 1$, after which the rest follows directly since if u determines u' , then $d(u, v)$ must determine $d(u', v)$.

By Lemma 7.5 we may suppose that $\delta \geq 5$. We fix a_1, a_2, a_3 at mutual distance δ , and fix $i < \delta/2$, to be determined more precisely later.

Take $u_1, v_1 \in \Gamma_i(a_1)$ with $d(u_1, v_1) = 2i$, and then correspondingly $u_2, v_2 \in \Gamma_i(a_2)$, $u_3, v_3 \in \Gamma_i(a_3)$, with u_1, u_2, u_3 and v_1, v_2, v_3 triples of vertices at mutual distance δ .

Now $d(u_1, v_3) = \delta - 2i$, and $d(u_1, u_3) = \delta$, so as usual $d(u_3, v_3) = 2i$. We now consider the following property of the triple (u_1, v_1, a_3) : *For $v_3 \in \Gamma_i(a_3)$ with $d(v_1, v_3) = \delta$, we have $d(u_1, v_3) = \delta - 2i$.* The triple (u_1, v_1, a_3) is isometric with (v_3, u_3, a_2) . It follows that $d(v_3, u_2) = \delta - 2i$. So

$$\delta = d(u_1, u_2) \leq 2(\delta - 2i)$$

This shows that $i \leq \delta/4$, so for a contradiction we require

$$\delta/4 < i < \delta/2$$

and for $\delta > 4$ this is possible. \square

We now give the classification of metrically homogeneous antipodal graphs of diameter 3; this is also treated in [AMP10].

THEOREM 15. *Let G be one of the following graphs: the pentagon (5-cycle), the line graph $E(K_{3,3})$ for the complete bipartite graph $K_{3,3}$, an independent set I_n ($n \leq \infty$), or the random graph Γ_∞ . Let G^* be the graph obtained from G by adjoining an additional vertex adjacent to all vertices of G , and let Γ be the graph obtained by taking two copies H_1, H_2 of G^* , with a fixed isomorphism $u \mapsto u'$ between them, and with additional edges (u, v') or (v', u) , for $u, v \in H_1$, just when (u, v) is not an edge of H_1 . Then Γ is a homogeneous antipodal graph of diameter 3 with pairing the given isomorphism $u \mapsto u'$. Conversely, any connected metrically homogeneous antipodal graph of diameter 3 is of this form.*

PROOF. Let Γ be connected, metrically homogeneous, and antipodal, of diameter 3. Fix a basepoint $* \in \Gamma$ and let $G = \Gamma_1(*)$. Then G is a homogeneous graph, which may be found in the Lachlan/Woodrow catalog given in §4.

Let $H_1 = G \cup \{*\}$. Then the pairing $u \leftrightarrow u'$ on Γ gives an isomorphism of H_1 with $H_2 = \Gamma_1(*)' \cup \{*\}'$. Furthermore, for $u, v \in \Gamma_1(*)$, we have $d(u, v') = \delta - d(u, v)$,

so the edge rule in Γ is the one we have described. It remains to identify the set of homogeneous graphs G for which the associated graph Γ is metrically homogeneous.

We claim that for any homogeneous graph G , the associated graph Γ has the following homogeneity property: if A, B are finite subgraphs of Γ both containing the point $*$, then any isometry $A \rightarrow B$ fixing $*$ extends to an automorphism of Γ . Given such A, B , we first extend to \hat{A}, \hat{B} by closing under the pairing $u \leftrightarrow u'$, then reduce to G by taking $\tilde{A} = \hat{A} \cap G$, $\tilde{B} = \hat{B} \cap G$. Then apply the homogeneity of G to get an isometry extending the given one on \tilde{A} to all of H_1 , fixing $*$, which then extends canonically to Γ . It is easy to see that this agrees with the given isometry on A .

This homogeneity condition implies that for such graphs Γ , the graph will be metrically homogeneous if and only if $\text{Aut}(\Gamma)$ is transitive on vertices. For any of these graphs Γ , whether metrically homogeneous or not, we have the pairing $u \leftrightarrow u'$. Furthermore, we can reconstruct Γ from $\Gamma_1(v)$ for any $v \in \Gamma$. So the homogeneity reduces to this: $\Gamma_1(v) \cong G$ for $v \in G = \Gamma_1(*)$; here we use the pairing to reduce to the case $v \in G$.

Now $\Gamma_1(v)$ is the graph obtained from the vertex $*$, the graph $G_1(v)$ induced on the neighbors of v in G , and the graph $G_2(v)$ induced on the non-neighbors of v in G , by taking the neighbors of $*$ to be $G_1(v)$, and switching the edges and nonedges between $G_1(v)$ and $G_2(v)$. Another way to view this would be to replace v by $*$, and then perform the switching between $G_1(v)$ and $G_2(v)$. So it is really only the latter that concerns us.

We go through the catalog. In the degenerate cases, with G complete or independent, there is no switching, so the corresponding graph Γ is homogeneous. But when G is complete this graph is not connected, so we set that case aside.

When G is imprimitive, we switch edges and non-edges between the equivalence classes not containing the fixed vertex v , and the vertices in the equivalence class of v other than v itself. As a result, the new graph becomes connected with respect to the equivalence relation on G , so this certainly does not work.

When G is primitive, nondegenerate, and finite, we have just the two examples mentioned above for which the construction does work, by inspection.

Lastly, we consider the Henson graphs $G = \Gamma_n$, generic omitting K_n , their complements, and the random graph Γ_∞ . The Henson graphs and their complements will not work here. For example, if $G = \Gamma_n$, then $G_2(v)$ contains K_{n-1} , and switching edges and nonedges with $G_1(v)$ will extend this to K_n . The complementary case is the same. So we are left with the case of the Rado graph. This is characterized by extension properties, and it suffices to check that these still hold after performing the indicated switch; and using the vertex v as an additional parameter, this is clear. \square

There is a good deal more to the general analysis of [AH06], Proposition 2.5 through Corollary 2.10, all with some parallels in our case, but the main examples in the finite case do not satisfy our conditions, while the main examples in our case have no finite analogs, so the statements gradually diverge, and it is better for us to turn to the consideration of graphs which are exceptional in another sense, and only then come back to the bipartite case.

In particular the main result of [AH06] is the following.

FACT 7.7 ([AH06, Theorem 3.3]). *An antipodal and bipartite finite distance transitive graph of diameter 6 and degree at least 3 is isomorphic to the 6-cube.*

We have infinite connected metrically homogeneous graphs of any diameter $\delta \geq 3$ which are both bipartite and antipodal. In the case $\delta = 6$, the associated graph Γ_2 is the generic bipartite graph of diameter 4, and the associated graph Γ_3 is isomorphic to Γ itself. So even in the context of Smith's Theorem, the two pictures eventually diverge. But graphs of odd diameter which are both bipartite and antipodal are quite special.

THEOREM (12). *Let Γ be a connected metrically homogeneous graph of odd diameter $\delta = 2\delta' + 1$ which is both antipodal and bipartite. Then $B\Gamma$ is connected, and Γ is the bipartite double cover of $B\Gamma$. The graph $B\Gamma$ is a metrically homogeneous graph with the following properties:*

- (1) $B\Gamma$ has diameter δ' ;
- (2) No triangle in $B\Gamma$ has perimeter greater than $2\delta' + 1$;
- (3) $B\Gamma$ is not antipodal.

Conversely, for any metrically homogeneous graph G of diameter with the three stated properties, there is a unique antipodal bipartite graph of diameter $2\delta' + 1$ such that $B\Gamma \cong G$.

We remark that conditions (1-3) on $B\Gamma$ imply that $B\Gamma_d$ is a clique of order at least 2. One exceptional case included under this theorem is that of the $(4d+2)$ -gon, of diameter $2d + 1$, associated with the $(2d + 1)$ -gon, of diameter d .

PROOF. Let A, B be the two halves of Γ . Then the metric on Γ is determined by the metrics on A and B and the pairing

$$a \leftrightarrow a'$$

between A and B determined by $d(a, a') = 2d + 1$, since

$$d(a_1, a'_2) = 2d + 1 - d(a_1, a_2)$$

So the uniqueness is clear.

Let us next check that the conditions on $B\Gamma$ are satisfied. The first is clear. For the second, suppose we have vertices (a_1, a_2, a_3) in $B\Gamma$ forming a triangle of perimeter at least $2d + 2$; we may construe these as vertices of A forming a triangle of perimeter $P \geq 4d + 4$. Then looking at the triangle (a_1, a_2, a'_3) , we have

$$\begin{aligned} d(a_1, a'_3) + d(a_2, a'_3) &= (4d + 2) - [d(a_1, a_3) + d(a_2, a_3)] \\ &= (4d + 2) - P + d(a_1, a_2) \\ &< d(a_1, a_2), \end{aligned}$$

contradicting the triangle inequality. Finally, consider an edge (a_1, a_2) of $B\Gamma$, which we construe as a pair of vertices of A at distance 2. Then there must be a vertex a such that a' is adjacent to both, and this means that $a_1, a_2 \in \Gamma_{2d}(a)$, that is $a_1, a_2 \in B\Gamma_d(a)$.

Conversely, suppose G is a metrically homogeneous graph of diameter d , and Γ is the metric space on $G \times \{0, 1\}$ formed by doubling the metric of G on $A = G \times \{0\}$ and on $B = G \times \{1\}$, pairing A and B by $(a, \epsilon)' = (a, 1 - \epsilon)$, and defining

$$d(a_1, a'_2) = 2d + 1 - d(a_1, a_2) \text{ for } a_1, a_2 \text{ in } A$$

The triangle inequality follows directly from the bound on the perimeters of triangles. We claim that Γ is a homogeneous metric space. Furthermore, the pairing $a \leftrightarrow a'$ is an isometry of Γ , and is recoverable from the metric: $y = x'$ if and only if $d(x, y) = 2d + 1$.

Let X, Y be finite subspaces of Γ , and f an isometry between them. Then f extends canonically to their closures under the antipodal pairing. So we may suppose X and Y are closed under the antipodal pairing; and composing f with the antipodal pairing if necessary, we may suppose f preserves the partition of Γ into A, B . Then restrict f to $A \cap X$, extend to A by homogeneity, and then extend back to Γ . Thus Γ is a homogeneous metric space.

Finally, we claim that the metric on Γ is the graph metric, and for this it suffices to show that vertices at distance 2 in the metric have a common neighbor in Γ . So let a_1, a_2 be two such vertices, taken for definiteness in A ; write $a_i = (v_i, 0)$. Taking $v \in G$ with $v_1, v_2 \in G_d(v)$, and $a = (v, 1)$, we find that a is adjacent to a_1 and a_2 . \square

Corollary 12.1 made this result explicit for the case in which $B\Gamma$ is already in the catalog.

COROLLARY (12.1). *Let Γ be an antipodal bipartite graph of diameter 5. Then $B\Gamma$ is either a pentagon, or the generic homogeneous graph omitting I_3 (the complement of the Henson graph G_3) and Γ is its bipartite double cover.*

PROOF. Let $G = B\Gamma$. Then G has diameter 2, so it is a homogeneous graph, on the list of Lachlan and Woodrow.

Furthermore, by the theorem, G contains I_2 but not I_3 . As $d = 2$ here, G contains a path of length 2 as well as a vertex at distance 2 from both vertices of an edge. By the Lachlan/Woodrow classification, in the finite case G is a 5-cycle and in the infinite case it must be the generic graph omitting I_3 . \square

The bipartite double cover of C_5 is C_{10} , and the bipartite double cover of G_3^c is $\Gamma_{1,4;12,11;\emptyset}^5$, the generic antipodal bipartite graph of diameter 5.

8. Exceptional Metrically Homogeneous Graphs

We turn now from the imprimitive case to the case in which the graph Γ_1 is exceptional.

If Γ is a metrically homogeneous graph, then Γ_1 is a homogeneous graph, and must occur in the short list of such graphs described in §4. There are three possibilities for Γ_1 which are compatible with the Fraïssé constructions we have described: an infinite independent set, the Henson graphs, and the Rado graph. Thus we make the following definition.

DEFINITION 8.1. Let Γ be a metrically homogeneous graph.

1. Γ is of *generic type* if the graph Γ_1 is of one the following three types:

- (1) an infinite independent set I_∞ ;
- (2) generic omitting K_n for some $n \geq 3$;
- (3) fully generic (the Rado graph).

2. If Γ is not of generic type, then we say it is of *exceptional type*.

All the bipartite graphs (indeed, all the triangle free ones) are included under “generic type” and will have to be treated under that heading. But it is useful to

dispose of the classification of the exceptional metrically homogeneous graphs as a separate case.

Thus in the exceptional case we have the following possibilities for Γ_1 .

- $\Gamma_1 \cong C_5$ or the line graph of $K_{3,3}$.
- $\Gamma_1 \cong m \cdot K_n$ or $K_m[I_n]$ with $1 \leq m, n \leq \infty$; and not of the form I_∞ .
- Γ_1 the complement of a Henson graph

We will classify those falling under the first two headings and show that the third case does not occur. The case in which Γ_1 is finite is covered by [Cam80, Mph82], split under two headings: (a) Γ finite; or (b) Γ infinite and Γ_1 finite. Those proofs use considerably less than metric homogeneity. We include a treatment of those cases here, but making full use of our hypothesis.

8.1. Exceptional graphs with Γ_1 finite and primitive.

LEMMA 8.2. *Let Γ be a connected metrically homogeneous graph of diameter at least 3 and degree at least 3, and suppose that Γ_1 is one of the primitive finite homogeneous graphs containing both edges and nonedges, that is C_5 or the line graph of $K_{3,3}$. Then Γ is the antipodal graph of diameter 3 obtained from Γ_1 in the manner of Theorem 15.*

PROOF. We fix a basepoint $*$ in Γ so that Γ_i is viewed as a specific subgraph of Γ for each i . The proof proceeds in two steps.

- (1) There is a $*$ -definable function from Γ_1 to Γ_2 .

We will show that for $v \in \Gamma_1$, the vertices of Γ_1 not adjacent to v have a unique common neighbor v' in Γ_2 .

If Γ_1 is a 5-cycle then this amounts to the claim that every edge of Γ lies in two triangles, and this is clear by inspection of an edge $(*, v)$ with $*$ the basepoint and $v \in \Gamma_1$.

Now suppose Γ_1 is $E(K_{3,3})$. We claim that every induced 4-cycle $C \cong C_4$ in Γ has exactly two common neighbors.

Consider $u, v \in \Gamma_1$ lying at distance 2, and let $G_{u,v}$ be the metric space induced on their common neighbors in Γ . This is a homogeneous metric space. Since these common neighbors consist of the basepoint $*$, the two common neighbors a, b of u, v in Γ_1 , and whatever common neighbors u, v may have in Γ_2 , we see that pairs at distance 1 occur, and the corresponding graph has degree 2 (looking at $*$) and is connected (looking at $(a, *, b)$). So $G_{u,v}$ is a connected metrically homogeneous graph of degree 2, and furthermore embeds in $\Gamma_1(u) \cong \Gamma_1$. So $G_{u,v}$ is a 4-cycle, and therefore (u, a, v, b) has exactly 2 neighbors, as claimed. This proves (1).

Now let $f : \Gamma_1 \rightarrow \Gamma_2$ be $*$ -definable. By homogeneity f is surjective, and as Γ_1 is primitive, it is bijective. It also follows from homogeneity that for $u, v \in \Gamma_1$, $d(u, v)$ determines $d(f(u), f(v))$, so f is either an isomorphism or an anti-isomorphism. Since Γ_1 is isomorphic to its complement, $\Gamma_1 \cong \Gamma_2$ in any case. Hence the vertices of Γ_2 have a common neighbor v , and $v \in \Gamma_3$. We claim that $|\Gamma_3| = 1$.

By homogeneity all pairs (u, v) in $\Gamma_2 \times \Gamma_3$ are adjacent. In particular for $v_1, v_2 \in \Gamma_3$ we have $\Gamma_1(v_1) = \Gamma_1(v_2)$ and $d(v_1, v_2) \leq 2$. Since we have pairs of vertices u_1, u_2 in Γ_1 at distance 1 or 2 for which $\Gamma_1(u_1) \neq \Gamma_1(u_2)$, we find $|\Gamma_3| = 1$.

It now follows that Γ is antipodal of diameter 3 and the previous analysis applies. \square

8.2. Exceptional graphs with $\Gamma_1 \cong K_m[I_n]$. We will prove the following.

PROPOSITION 8.3. *Let Γ be a connected metrically homogeneous graph of diameter δ , and suppose $\Gamma_1 \cong K_m[I_n]$ with $1 \leq m, n \leq \infty$. Then one of the following occurs.*

- (1) $\delta \leq 2$, Γ is homogeneous and found under the Lachlan/Woodrow classification.
- (2) $\Gamma \cong C_n$, an n -cycle, for some n ;
- (3) $m = 1$, $\delta \geq 3$, and one of the following occurs.
 - (a) n is finite, and Γ is the bipartite complement of a perfect matching between two sets of order $n + 1$.
 - (b) $\delta = \infty$, Γ is a regular tree of degree n , with $2 \leq n \leq \infty$.
 - (c) $n = \infty$, and any two vertices at distance 2 have infinitely many common neighbors.

LEMMA 8.4. *Let Γ be a connected metrically homogeneous graph of diameter at least 3, and suppose that Γ_1 is a complete multipartite graph of the form $K_m[I_n]$ (the complement of $m \cdot K_n$). Then $m = 1$.*

PROOF. Suppose that $m > 1$. Fix a geodesic $(*, u, v)$ of length 2 and let $\Gamma_i = \Gamma_i(*)$. Let A be the set of neighbors of u in Γ_1 . Then $A \cong K_{m-1}[I_n]$, and the neighbors of u in Γ include $*, v$, and A . Now “ $d(x, y) > 1$ ” is an equivalence relation on $\Gamma_1(u)$, and $*$ is adjacent to A , so v is adjacent to A . Now if we replace u by $u' \in A$ and argue similarly with respect to $(*, u', v)$, we see that the rest of Γ_1 is also adjacent to v , that is $\Gamma_1 \subseteq \Gamma_1(v)$. Now switching $*$ and v , by homogeneity $\Gamma_1(v) \subseteq \Gamma_1$. But then the diameter of Γ is 2, a contradiction. Thus $m = 1$. \square

We next consider the case $m = 1$, that is, Γ_1 is an independent set. We suppose in this case that the set of common neighbors of any pair of points at distance 2 is finite, which covers the case in which n is finite but also picks up the case of an infinitely branching tree. When $n = 2$ we have either a cycle or a 2-way infinite path, so we will leave this case aside in our analysis.

LEMMA 8.5. *Let Γ be a metrically homogeneous graph of diameter at least 2, with $\Gamma_1 \cong I_n$, $3 \leq n \leq \infty$. Suppose that for $u, v \in \Gamma$ at distance two, the number k of common neighbors of u, v is finite. Then either $k = 1$, or $n = k + 1$.*

PROOF. We consider Γ_1 and Γ_2 with respect to a fixed basepoint $* \in \Gamma$. For $u \in \Gamma_2$, let I_u be the k -set consisting of its neighbors in Γ_1 . Any k -subset of Γ_1 occurs as I_u for some u . For $u, v \in \Gamma_2$ let $u \cdot v = |I_u \cap I_v|$. If $u \cdot v \geq 1$ then $d(u, v) = 2$. Now $\text{Aut}(\Gamma)_*$ has a single orbit on pairs in Γ_2 at distance 2, while every value i in the range $\max(1, 2k - n) \leq i \leq k - 1$, will occur as $u \cdot v$ for some such u, v . Therefore $k \leq 2$ or $k = n - 1$, with n finite in the latter case.

The case $k = 2 < n - 1$ is eliminated by a characteristic application of homogeneity. A set of three pairs in Γ_1 which intersect pairwise may or may not have a common element (once $n \geq 4$), so if we choose u_1, u_2, u_3 and v_1, v_2, v_3 in Γ_2 corresponding to these two possibilities for the associated I_{u_j} and I_{v_j} , we get isometric configurations $(*, u_1, u_2, u_3)$ and $(*, v_1, v_2, v_3)$ which lie in distinct orbits of $\text{Aut} \Gamma$. So we have $k = 1$ or $n = k + 1$. \square

LEMMA 8.6. *Let Γ be a connected metrically homogeneous graph of diameter at least 3, with $\Gamma_1 \cong I_n$, and $3 \leq n \leq \infty$. Suppose any pair of vertices at distance 2 have k common neighbors, with $k < \infty$. Then one of the following occurs.*

- (1) $n = k + 1$, and Γ is the complement of a perfect matching, in other words the antipodal graph of diameter 3 obtained by doubling Γ_1 .
- (2) $k = 1$, and Γ is a k -regular tree.

PROOF. We fix a basepoint $*$ and write Γ_i for $\Gamma_i(*)$.

Suppose first that $n = k + 1$. Then any two vertices of Γ_2 lie at distance 2, and there is a $*$ -definable function $f : \Gamma_2 \rightarrow \Gamma_1$ given by the nonadjacency relation. By homogeneity f is surjective, and as Γ_2 is primitive, f is bijective. In particular $\Gamma_2 \cong \Gamma_1$ and there is a vertex $v \in \Gamma_3$ adjacent to all vertices of Γ_2 , hence $\Gamma_1(v) = \Gamma_2$. It follows readily that $|\Gamma_3| = 1$ and Γ is antipodal of diameter 3. The rest follows by our previous analysis.

Now suppose that $k = 1$. It suffices to show that Γ is a tree.

Suppose on the contrary that there is a cycle C in Γ , which we take to be of minimal diameter d . Then the order of C is $2d$ or $2d + 1$.

Suppose the order of C is $2d$. Then for $v \in \Gamma_d$, v has at least two neighbors u_1, u_2 in Γ_{d-1} , whose distance is therefore 2. Furthermore, in Γ_{d-2} there are no edges, and each vertex of Γ_{d-2} has a unique neighbor in Γ_{d-3} , so each vertex of Γ_{d-2} has at least two neighbors in Γ_{d-1} , whose distance is therefore 2.

So for $u_1, u_2 \in \Gamma_{d-1}$, there is a common neighbor in Γ_d , and also in Γ_{d-2} . This gives a 4-cycle in Γ , contradicting $k = 1$.

So the order of C is $2d + 1$. In particular, each $v \in \Gamma_d$ has a unique neighbor in Γ_{d-1} , and Γ_d contains edges.

Let G be a connected component of Γ_d . Suppose u, v in G are at distance 2 in G . As any vertex in Γ_{d-1} has at least two neighbors in Γ_d , the vertices u, v must have a common neighbor in Γ_{d-1} as well as in G , and this contradicts the hypothesis $k = 1$. So the connected components of Γ_d are simply edges.

Take $a \in \Gamma_{d-1}$, $u_1, v_1 \in \Gamma_d$ adjacent to a , and take $u_2, v_2 \in \Gamma_d$ adjacent to u_1, v_1 respectively. By homogeneity there is an automorphism fixing the basepoint $*$ and interchanging u_1 with v_2 ; this also interchanges u_2 and v_1 . Hence $d(u_2, v_2) = 2$. It follows that u_2, v_2 have a common neighbor b in Γ_{d-1} . Now $(a, u_1, u_2, b, v_2, v_1, a)$ is a 6-cycle. Since the minimal cycle length is odd, we have $|C| = 5$ and $d = 2$.

Furthermore, the previous paragraph shows that the element b is determined by $a \in \Gamma_1$ and the basepoint $*$: we take $u \in \Gamma_2$ adjacent to a , $v \in \Gamma_2$ adjacent to u , and $b \in \Gamma_1$ adjacent to v . So the function $a \mapsto b$ is $*$ -definable. However Γ_1 is an independent set of order at least 3, so this violates homogeneity. \square

This completes the proof of Proposition 8.3.

8.3. Exceptional graphs with $\Gamma_1 \cong m \cdot K_n$, $n \geq 2$. We deal here with the tree-like graphs $T_{r,s}$ derived from the trees $T(r, s)$ as described in §5.3. With $r, s < \infty$ these graphs are locally finite (that is, the vertex degrees are finite). Conversely:

FACT 8.7 (Macpherson, [Mph82]). *Let G be an infinite locally finite distance transitive graph. Then G is $T_{r,s}$ for some finite $r, s \geq 2$.*

The proof uses a result of Dunwoody on graphs with nontrivial cuts given in [Dun82].

The graph $T_{r,s}$ has $\Gamma_1 \cong s \cdot K_{r-1}$. We are interested now in obtaining a characterization of these graphs, also with r or s infinite, in terms of the structure of Γ_1 .

Our goal is the following.

PROPOSITION 8.8. *Let Γ be a connected metrically homogeneous graph with $\Gamma_1 \cong m \cdot K_n$, with $n \geq 2$ and $\delta \geq 3$. Then $m \geq 2$, and $\Gamma \cong T_{m,n+1}$.*

Our main goal will be to show that any two vertices at distance two have a unique common neighbor. We divide up the analysis into three cases. Observe that in a metrically homogeneous graph with $\Gamma_1 \cong m \cdot K_n$, the common neighbors of a pair $a, b \in \Gamma$ at distance 2 will be an independent set, since for u_1, u_2 adjacent to a, b , and each other, we would have the path (a, u_1, b) inside $\Gamma_1(u_2)$.

LEMMA 8.9. *Let Γ be a metrically homogeneous graph of diameter at least 3 and suppose that $\Gamma_1 \cong m \cdot K_n$ with $n \geq 3$. Then for $u, v \in \Gamma$ at distance 2, there are at most two vertices adjacent to both.*

PROOF. Supposing the contrary, every induced path of length 2 is contained in two distinct 4-cycles. Fix a basepoint $* \in \Gamma$ and let $\Gamma_i = \Gamma_i(*)$.

Fix $v_1, v_2 \in \Gamma_1$ adjacent. For $i = 1, 2$, let H_i be the set of neighbors of v_i in Γ_2 . The connected components of H_1 and H_2 are cliques. We claim

$$H_1 \cap H_2 = \emptyset$$

Otherwise, consider $v \in H_1 \cap H_2$ and the path $(*, v_1, v)$ contained in $\Gamma_1(v_2)$. We will find $u_1 \in H_1$, and $u_2, u'_2 \in H_2$ distinct, so that

$$d(u_1, u_2) = d(u_1, u'_2) = 1$$

Extend the edge (v_1, v_2) to a 4-cycle (v_1, v_2, u_2, u_1) . Then $u_1, u_2 \notin \Gamma_1 \cup \{*\}$, so $u_1 \in H_1$ and $u_2 \in H_2$. By our hypothesis there is a second choice of u_2 with the same properties.

With the vertices u_1, u_2, u'_2 fixed, let A, B, B' denote the components of H_1 and H_2 , respectively, containing the specified vertices. Observe that B and B' are distinct: otherwise, the path (u_1, u_2, v_2) would lie in $\Gamma_1(u'_2)$.

The relation “ $d(x, y) = 1$ ” defines a bijection between A and B

With $u \in A$ fixed, it suffices to show the existence and uniqueness of the corresponding element of B . The uniqueness amounts to the point just made for u_1 , namely that $B \neq B'$.

For the existence, we may suppose $u \neq u_1$. Then $d(u, v_2) = d(u, u_2) = 2$. We have an isometry

$$(*, v_1, u, v_2, u_2) \cong (*, v_2, u_2, v_1, u)$$

and hence the triple (u, u_1, u_2) with $u_1 \in H_1$ corresponds to an isometric triple (u_2, u', u) with $u' \in H_2$.

Thus we have a bijection between A and B definable from $(*, v_1, v_2, u_1, u_2)$ and hence we derive a bijection between B and B' definable from $(*, v_1, v_2, u_1, u_2, u'_2)$. Using this, we show

$$n = 2$$

The graph induced on $B \cup B'$ is $2 \cdot K_n$, and any isometry between finite subsets of $B \cup B'$ containing u_2, u'_2 which fixes u_2 and u'_2 will be induced by $\text{Aut}(\Gamma)$. So if there is a bijection between B and B' invariant under the corresponding automorphism group, we find $n = 2$. \square

LEMMA 8.10. *Let Γ be a metrically homogeneous graph of diameter at least 3 and suppose that $\Gamma_1 \cong m \cdot K_n$ with $n \geq 2$. Let $u, v \in \Gamma$ lie at distance 2, and suppose u, v have finitely many common neighbors. Then they have a unique common neighbor.*

PROOF. We fix a basepoint $*$, and for $u \in \Gamma_2$ we let I_u be the set of neighbors of u in Γ_1 . Our assumption is that $k = |I_u|$ is finite. Then any independent subset of Γ_1 of cardinality k occurs as I_u for some $u \in \Gamma_2$.

We consider the $k - 2$ possibilities:

$$|I_u \cap I_v| = i \text{ with } 1 \leq i \leq k - 1$$

As $n \geq 2$, all possibilities are realized, whatever the value of m . However in all such cases, $d(u, v) \leq 2$, so we find $k - 1 \leq 2$, and $k \leq 3$.

We claim that for $u, v \in \Gamma_2$ adjacent, we have $|I_u \cap I_v| = 1$.

There is a clique v, u_1, u_2 with $v \in \Gamma_1$ and $u_1, u_2 \in \Gamma_2$. As u_1, u_2 are adjacent their common neighbors form a complete graph. On the other hand I_{u_1} and I_{u_2} are independent sets, so their intersection reduces to a single vertex. By homogeneity the same applies whenever $u_1, u_2 \in \Gamma_2$ are adjacent, proving our claim.

Now suppose $k = 3$. Then $|I_u \cap I_v|$ can have cardinality 1 or 2, and the case $|I_u \cap I_v| = 2$ must then correspond to $d(u, v) = 2$.

Now $k \leq m$, so we may take pairs (a_i, b_i) for $i = 1, 2, 3$ lying in distinct components of Γ_1 . Of the eight triples t formed by choosing one of the vertices of each of these pairs, there are four in which the vertex a_i is selected an even number of times. Let a vertex $v_t \in \Gamma_2$ be taken for each such triple, adjacent to its vertices. Then the four vertices v_t form a complete graph K_4 . It follows that K_3 embeds in Γ_1 , that is $n \geq 3$. So we may find independent triples I_1, I_2, I_3 such that $|I_1 \cap I_2| = |I_2 \cap I_3| = 1$ while $I_1 \cap I_3 = \emptyset$. Take $u_1, u_2, u_3 \in \Gamma_2$ with $I_i = I_{u_i}$ and with $d(u_1, u_2) = d(u_2, u_3) = 1$. Then $I_{u_1} \cap I_{u_3} = \emptyset$, while $d(u_1, u_2) \leq 2$, a contradiction. Thus $k = 2$.

With $k = 2$, suppose $m > 2$. For $u \in \Gamma_2$, let \hat{I}_u be the set of components of Γ_1 meeting I_u . We consider the following two properties of a pair $u, v \in \Gamma_2$:

$$|I_u \cap I_v| = 1; |\hat{I}_u \cap \hat{I}_v| = i \text{ (} i = 1 \text{ or } 2\text{)}$$

These both occur, and must correspond in some order with the conditions $d(u, v) = 1$ or 2. But just as above we find u_1, u_2, u_3 with $d(u_1, u_2) = d(u_2, u_3) = 1$ and $I_{u_1} \cap I_{u_3} = \emptyset$, and as $d(u_1, u_3) \leq 2$ this is a contradiction.

So we come down to the case $m = k = 2$. But then for $v \in \Gamma_2$, all components of $\Gamma_1(v)$ are represented in Γ_1 , and hence v has no neighbors in Γ_3 , a contradiction. \square

LEMMA 8.11. *Let Γ be a metrically homogeneous graph of diameter at least 3 and suppose that $\Gamma_1 \cong m \cdot K_n$ with $n \geq 2$. Let $u, v \in \Gamma$ lie at distance 2, and suppose u, v have infinitely many common neighbors. Then $n = \infty$.*

PROOF. We fix a basepoint $*$, and for $u \in \Gamma_2$ we let I_u be the set of neighbors of u in Γ_1 . Our assumption is that I_u is infinite. Then any finite independent subset of Γ_1 is contained in I_u for some $u \in \Gamma_2$.

For $u \in \Gamma_2$, let \hat{I}_u be the set of components of Γ_1 which meet I_u , and let \hat{J}_u be the set of components of Γ_1 which do not meet I_u . We show first that

$$\hat{J}_u \text{ is infinite.}$$

Supposing the contrary, let $k = |\hat{J}_u| < \infty$ for $u \in \Gamma_2$. Any set of k components of Γ_1 will be \hat{J}_u for some $u \in \Gamma_2$, and the $k + 1$ relations on Γ_2 defined by

$$|\hat{J}_u \cap \hat{J}_v| = i$$

for $i = 0, 1, \dots, k$ will be nontrivial and distinct. Furthermore, for any preassigned k components \hat{J} , and any vertex $a \in \Gamma_1$ not in the union of \hat{J} , there is a vertex u with $\hat{J}_u = \hat{J}$ and $a \in I_u$, so our $(k + 1)$ relations are realized by pairs $u, v \in \Gamma_2$ with $I_u \cap I_v \neq \emptyset$, and hence $d(u, v) \leq 2$. Hence $k + 1 \leq 2$, $k \leq 1$.

Suppose $k = 1$ and fix a vertex $v_0 \in \Gamma_1$. Then for $u, v \in \Gamma_2$ adjacent to v_0 , the two relations $\hat{J}_u = \hat{J}_v$, $\hat{J}_u \neq \hat{J}_v$ correspond in some order to the relations $d(u, v) = 1$, $d(u, v) = 2$, and since the first relation is an equivalence relation, they correspond in order.

With $u \in \Gamma_2$, $v_0, v_1 \in I_u$ distinct, there are u_0, u_1 in Γ_2 with u_0 adjacent to u and v_0 , and with u_1 adjacent to u and v_1 . The neighbors of u form a graph of type $\infty \cdot K_n$, so $d(u_0, u_1) = 2$. However $d(u, u_0) = d(u, u_1) = 1$ and hence $\hat{J}_{u_0} = \hat{J}_{u_1}$, a contradiction.

So $k = 0$ and for $u \in \Gamma_2$, the set I_u meets every component of Γ_1 . That is, $\Gamma_1(u)$ meets every component of Γ_1 , and after switching the roles of u and the basepoint $*$, we conclude Γ_1 meets every component of $\Gamma_1(u)$, which is incompatible with the condition $\delta \geq 3$. So \hat{J}_u is infinite for $u \in \Gamma_2$.

Now we claim

For $u, v \in \Gamma_2$ adjacent, $\hat{I}_u \setminus \hat{I}_v$ is infinite

Supposing the contrary, for all adjacent pairs $u, v \in \Gamma_2$, the sets \hat{I}_u and \hat{I}_v coincide up to a finite difference.

Take $u \in \Gamma_2$, $v_0, v_1 \in \Gamma_1$ adjacent to u , and u_0, u_1 adjacent to u, v_0 or u, v_1 respectively. Then, as above, $d(u_0, u_1) = 2$, while $\hat{J}_{u_0} = \hat{J}_u = \hat{J}_{u_1}$. Thus for $u, v \in \Gamma_2$ with $d(u, v) \leq 2$ we have $\hat{J}_u = \hat{J}_v$. Furthermore, the size of the difference $|\hat{J}_u \setminus \hat{J}_v|$ is bounded, say by ℓ . But we can fix $u \in \Gamma_2$ and then find $v \in \Gamma_2$ so that I_u meets I_v but \hat{I}_v picks up more than ℓ components of \hat{J}_u . Since I_u meets I_v we have $d(u, v) \leq 2$, and thus a contradiction. This proves our claim.

We make a third and last claim of this sort.

For $u, v \in \Gamma_2$ adjacent, $\hat{I}_u \cap \hat{I}_v$ is infinite

Supposing the contrary, let k' be $|\hat{I}_u \cap \hat{I}_v|$ for $u, v \in \Gamma_2$ adjacent, fix $u, v \in \Gamma_2$ adjacent, and let I be a finite independent subset consisting of representatives for more than k' components in $\hat{I}_v \setminus \hat{I}_u$.

Take $a \in I_u \cap I_v$ and take $v' \in \Gamma_2$ adjacent to a and to u , with $J \subseteq \hat{I}_{v'}$ and $I_v \neq I_{v'}$.

Now u, v, v' are adjacent to a , and v, v' are adjacent to u , so v and v' are adjacent. But by construction $|\hat{I}_{v'} \setminus \hat{I}_v| > k'$. This proves the third claim.

Now, finally, suppose n is finite. Take $u \in \Gamma_2$, $v \in I_u$, and let A, B be components of Γ_1 which are disjoint from I_u but meet $I_{u'}$ for some u' adjacent to u, v . Then

We can then find neighbors $w_{a,b}$ of u, v in Γ_2 for which the intersection of I_w with A and B respectively is an arbitrary pair of representatives a, b . But this requires $n - 1 \geq n^2$, and gives a contradiction. \square

COROLLARY 8.11.1. *Let Γ be a metrically homogeneous graph of diameter at least 3 and suppose that $\Gamma_1 \cong m \cdot K_n$ with $n \geq 2$. Then for $u, v \in \Gamma$ with $d(u, v) = 2$, there is a unique vertex adjacent to both.*

PROOF. Apply the last three lemmas. If u, v have infinitely many common neighbors, then n is infinite. In particular, $n \geq 3$. But then they have at most two common neighbors. So in fact u, v have finitely many neighbors, and we apply Lemma 8.10. \square

After this somewhat laborious reduction, we can complete the proof of Proposition 8.8.

Proof of Proposition 8.8. If $m = 1$ then evidently Γ is complete, contradicting the hypothesis on δ . So $m \geq 2$.

By definition, the blocks of Γ are the maximal 2-connected subgraphs. Any edge of Γ is contained in a unique clique of order $n + 1$. It suffices to show that these cliques form the blocks of Γ , or in other words that any cycle in Γ is contained in a clique.

Supposing the contrary, let C be a cycle of minimal order embedding into Γ as a subgraph, and not contained in a clique. The cycle C carries two metrics: its metric d_C as a cycle, and the metric d_Γ induced by Γ . We claim these metrics coincide. In any case, $d_\Gamma \leq d_C$.

If the metrics disagree, let $u, v \in C$ be chosen at minimal distance such that $d_\Gamma(u, v) < d_C(u, v)$, let $P = (u, \dots, v)$ be a geodesic in Γ , and let $Q = (u, \dots, v)$ be a geodesic in C . Let v' be the first point of intersection of P with Q , after u . Then $P \cup Q$ contains a cycle C' smaller than C , and containing u, v' . Hence by hypothesis the vertices of C' form a clique in Γ , and in particular u, v' are adjacent in Γ .

If u, v' are adjacent in C , then $d_C(v, v') = d_C(u, v) - 1$, $d_\Gamma(v, v') = d_\Gamma(u, v) - 1$, so $d_\Gamma(v, v') < d_C(v, v')$, and this contradicts the choice of u, v . So they are not adjacent, and the edge (u, v') belongs to two cycles C_1, C_2 whose union contains C , both smaller than C . So the vertices of C_1 and C_2 are cliques. Since C is not a clique, there are $u_1 \in C_1$ and $u_2 \in C_2$ nonadjacent in Γ . Then $d(u_1, u_2) = 2$, and there is a unique vertex adjacent to u_1 and u_2 ; but u, v' are two such, a contradiction.

Thus the embedding of C into Γ respects the metric. In particular, C is an induced subgraph of Γ .

Let d be the diameter of C , so that the order of C is either $2d$ or $2d + 1$. Fix a basepoint $* \in C$, and let $\Gamma_i = \Gamma_i(*)$.

For $v \in \Gamma_i$ with $i < d$, we claim that there is a unique geodesic $[*, v]$ in Γ . Otherwise, take $i < d$ minimal such that Γ contains vertices u, v at distance i with two distinct geodesics P, Q from u to v . By the minimality, these geodesics are disjoint. Their union forms a cycle smaller than C , hence they form a clique. As they are geodesics, $i = 1$ and then in any case the geodesic is unique.

Let $v \in \Gamma_{d-1}$, and let $H = \Gamma_1(v)$, a copy of $m \cdot K_n$. Then H contains a unique component meeting Γ_{d-2} . We claim that no other component of H meets Γ_{d-1} .

Suppose on the contrary that u_1, u_2 are adjacent to v with $u_1 \in \Gamma_{d-1}$, $u_2 \in \Gamma_{d-2}$, and $d(u_1, u_2) = 2$. Then taking P_1, P_2 to be the unique geodesics from u_1 or u_2 to $*$, $P_1 \cup P_2$ contains a cycle smaller than C , and containing the path (u_1, v, u_2) , and hence is not a clique. This is a contradiction.

Since Γ_d contains at least one component of H , in particular there is an edge in Γ_d whose vertices have a common neighbor in Γ_{d-1} . Using this, we eliminate the case $|C| = 2d + 1$ as follows.

If $|C| = 2d + 1$ then $C \cap \Gamma_d$ consists of two adjacent vertices v_1, v_2 , whose other neighbors u_1, u_2 in C lie in Γ_{d-1} . Furthermore v_1, v_2 have a common neighbor u in Γ_{d-1} , and $u \neq u_1, u_2$. Let P, P_1 be the unique geodesics in Γ connecting $*$ with u, u_1 respectively. Then $P \cup P_1 \cup \{v_1\}$ contains a cycle shorter than C , which contains the path (u_1, v, v_1) , and we have a contradiction.


Thus $|C| = 2d$, in other words the vertices $v \in \Gamma_d$ are connected to the basepoint $*$ by at least two distinct geodesics, and any two such geodesics will be disjoint.

Take $u \in \Gamma_{d-1}$, and $u' \in \Gamma_1$, with $d(u, u') = d - 2$. Take $v' \in \Gamma_1$ with $d(u', v') = 2$. We claim

$$d(u, v') = d$$

Otherwise, with P, Q the geodesics from u to u' and v' respectively, we have $|P \cup Q| \leq 2d - 1 < |C|$, and hence $(u', *, v')$ is contained in a cycle smaller than C , a contradiction since u', v' are nonadjacent.

Thus $d(u, v') = d$, and the extension of the geodesic from u to u' by the path $(u', *, v')$ gives a geodesic Q from u to v' . There is a second geodesic Q' from u to v' , disjoint from Q . Let u_1 be the unique neighbor of u in Γ_{d-2} ; this lies on Q . Let u'_1 be the neighbor of u in Q' . As the cycle $Q \cup Q'$ satisfies the same condition as the cycle C , the metric on this cycle agrees with the metric in Γ , and in particular $d(u_1, u'_1) = 2$. By the minimality of C , we have $u'_1 \in \Gamma_d$. Let u'_2 be the following neighbor of u'_1 in Q' . Then $d(u, u'_2) = 2$, with $u, u'_2 \in \Gamma_{d-1}$.



Suppose $m \geq 3$. Then there is u^* adjacent to u and at distance 2 from u_1 and u'_1 . In view of the minimality of C , the vertex u^* cannot be in Γ_{d-2} or Γ_{d-1} , hence lies in Γ_d . Hence any two vertices of Γ_d at distance 2 have a common neighbor in Γ_{d-1} . This applies to u and u'_2 and contradicts Corollary 8.11.1. We conclude that

$$m = 2$$

Fix an edge a_1, a_2 in Γ_{d-1} . For $i, j = 1, 2$ in some order, set

$$H_{ij} = \{v \in \Gamma : d(v, a_i) = 1, d(v, a_j) = 2\}$$

We claim $H_{ij} \subseteq \Gamma_d$.

As a_1, a_2 have the same unique neighbor in Γ_{d-2} , we have


$$(H_{12} \cup H_{21}) \cap \Gamma_{d-2} = \emptyset$$

Now suppose $v \in H_{12} \cap \Gamma_{d-1}$. Then $d(v, a_2) = 2$ and $v, a_2 \in \Gamma_{d-1}$. Now since $|C| = 2d$, there is $w \in \Gamma_d$ adjacent to v and a_2 . Then a_2, v have the two common neighbors w and a_1 , a contradiction. So $H_{12}, H_{21} \subseteq \Gamma_d$.

Let $v_1 \in H_{12}$, $v_2 \in H_{21}$. We claim that

$$d(v_1, v_2) = 3$$

Otherwise, there is a cycle of length at most 5, not contained in a clique. As $|C|$ is even, it follows that C is a 4-cycle, so vertices at distance 2 have at least two common neighbors, a contradiction. Thus for any other choice $v'_1 \in H_{12}$, $v'_2 \in H_{21}$, we have $(*, a_1, a_2, v_1, v_2)$ isometric with $(*, a_1, a_2, v'_1, v'_2)$.



There is a unique element $u \in \Gamma_1$ at distance $d-2$ from a_1 and a_2 ; namely, the element at distance $d-3$ from their common neighbor in Γ_{d-2} . On the other hand, for $v \in \Gamma_d$, if I_v is the set $\{v' \in \Gamma_1 : d(v, v') = d-1\}$, then by our hypotheses, I_v is a pair of representatives for the two components of Γ_1 . And if $v \in H_{12}$ or H_{21} , one of these representatives will be u . Let B be the component of Γ_1 not containing u . Then the distance from a_1 or a_2 to a vertex of B is d . It follows that all vertices of B will occur as the second vertex of I_v for some $v_1 \in H_{12}$ and for some $v_2 \in H_{21}$. Therefore, we may choose pairs (v_1, v_2) and (v'_1, v'_1) with $v_1, v'_1 \in H_{12}$, $v_2, v'_2 \in H_{21}$, and $I_{v_1} = I_{v_2}$, while $I_{v'_1} \neq I_{v'_2}$. But as $(*, a_1, a_2, v_1, v_2)$ and $(*, a_1, a_2, v'_1, v'_2)$ are isometric, this contradicts homogeneity. \square

8.4. Γ_1 cannot be the complement of a Henson graph. The last exceptional case requiring consideration is the one in which Γ_1 is the complement of a Henson graph.

LEMMA 8.12. *Let Γ be a metrically homogeneous graph of diameter at least 3 with Γ_1 infinite and primitive. Then Γ_1 contains an infinite independent set.*

PROOF. Suppose the contrary. Evidently Γ_1 contains an independent pair. By the Lachlan/Woodrow classification, if n is minimal such that Γ_1 contains no independent set of order n , then for any finite graph G which contains no independent set of order n , G occurs as an induced subgraph of Γ_1 .

We consider a certain amalgamation diagram involving subspaces of Γ . Let A be the metric space with three points a, b, x constituting a geodesic, with $d(a, b) = 2$, $d(b, x) = 1$, $d(a, x) = 3$. Let B be the metric space on the points a, b and a further set Y of order $n-1$ with the metric given by

$$\begin{aligned} d(a, y) &= d(b, y) = 1 \quad (y \in Y) \\ d(y, y') &= 2 \quad (y, y' \in Y \text{ distinct}) \end{aligned}$$

As the diameter of Γ is at least 3, the geodesic A certainly occurs as a subspace of Γ . On the other hand, the metric space B embeds into Γ_1 , and hence into Γ . Therefore there is some amalgam $G = A \cup B$ embedding into Γ as well.

Now for $y \in Y$, the structure of (a, x, y) forces $d(x, y) \geq 2$. On the other hand, the element b forces $d(x, y) \leq 2$. Thus in G , the set $Y \cup \{x\}$ is an independent set of order n . Furthermore this set is contained in $\Gamma_1(b)$, so we arrive at a contradiction. \square

We may sum up this section as follows.

THEOREM (10). *Let Γ be a connected metrically homogeneous graph and suppose that Γ_1 is either imprimitive or contains no infinite independent set. Then Γ_1 is found in our catalog under case I or case II.*

PROOF. By the Lachlan-Woodrow classification of homogeneous graphs and Lemma 8.12, Γ_1 must be finite, imprimitive, or the complement of a Henson graph, and we have just ruled out this last possibility.

Suppose Γ_1 is finite and primitive. Then Lemma 8.2 applies if Γ_1 is neither complete nor an independent set. If Γ_1 is complete, then Γ is complete. Proposition 8.3 covers the case in which Γ_1 is a finite independent set.

Lastly, suppose Γ_1 is imprimitive. If Γ_1 is of the form $K_m[I_n]$ with $m, n \geq 2$ then Proposition 8.3 applies. Alternatively, Γ_1 may have the form $m \cdot K_n$ with $m, n \geq 2$. This is covered by Proposition 8.8. \square

It is now appropriate to return to the bipartite case and look at metrically homogeneous bipartite graphs Γ for which the associated graph $B\Gamma$ falls on the exceptional side.

9. Exceptional Bipartite Metrically Homogeneous Graphs

If Γ is a connected bipartite metrically homogeneous graph, then we consider the graph $B\Gamma$ induced on each half of the bipartition by the edge relation “ $d(x, y) = 2$ ”, as described in §7. In addition to the exceptional graphs considered in the previous section, we wish to consider those for which $B\Gamma$ is itself exceptional in the sense of the previous section. The result in that case will be as follows.

THEOREM (13). *Let Γ be a connected, bipartite, and metrically homogeneous graph, of diameter at least 3, and degree at least 3, and with Γ_1 infinite. Then either $B\Gamma_1$ is isomorphic to the Rado graph, or $B\Gamma$ and Γ are in the catalog under one of the following headings.*

- (1) $B\Gamma \cong T_{\infty, \infty}$, and Γ is an infinitely branching tree $T_{2, \infty}$.
- (2) $B\Gamma \cong K_{\infty}$, and Γ has diameter 3, with Γ either the complement of a perfect matching, or the generic bipartite graph $\Gamma_{\infty, 0; 10, 7; \emptyset}^3$.
- (3) $B\Gamma \cong K_{\infty}[I_2]$, Γ has diameter 4, and $\Gamma \cong \Gamma_{\infty, 0; 10, 9; \emptyset}^4$ is the generic antipodal bipartite graph of diameter 4.
- (4) $B\Gamma \cong G_n^c$, the complement of a Henson graph, for some $n \geq 3$, and $\Gamma \cong \Gamma_{\infty, 0; 14, 9; \{I_n^{(4)}\}}^4$ the generic bipartite graph in which there is no set of n vertices which are pairwise at distance 4.
- (5) $B\Gamma \cong G_3^c$, $\Gamma \cong \Gamma_{\infty, 0; 12, 11; \emptyset}^5$ antipodal bipartite of diameter 5 as in Corollary 12.1.

If Γ_1 is finite then this was already covered in §8. So we may suppose $\Gamma_1 \cong I_{\infty}$. This means that $B\Gamma$ contains an infinite clique and thus $B\Gamma_1$ also contains an infinite clique. Furthermore, it follows from the connectedness and homogeneity of Γ that $B\Gamma$ is connected.

By the Lachlan/Woodrow classification this already reduces the possibilities to the following.

- (1) $B\Gamma_1$ is imprimitive of the form $m \cdot K_{\infty}$ or $K_{\infty}[I_m]$, with $1 \leq m \leq \infty$;
- (2) $B\Gamma_1$ is generic omitting I_n , for some finite $n \geq 2$;

(3) $B\Gamma_1$ is the Rado graph.

The possibilities for the exceptional graph $B\Gamma$ when $B\Gamma_1$ falls under (1) or (2) above are then known, and will be treated in the following order.

- (1) $B\Gamma$ is a clique;
- (2) $B\Gamma \cong K_\infty[I_m]$;
- (3) $B\Gamma \cong T_{r,s}$ with $2 \leq r, \infty \leq \infty$;
- (4) $B\Gamma$ is the complement of a Henson graph.

9.1. Bipartite graphs with $B\Gamma$ a clique.

PROPOSITION 9.1. *Let Γ be a connected, bipartite, and metrically homogeneous graph, of diameter $\delta \geq 3$. Suppose that $B\Gamma$ is an infinite clique. Then $\delta = 3$ and Γ is either the complement of a perfect matching, or a generic bipartite graph.*

These possibilities occur in the catalog as $\Gamma_{a,3}^3$ and $\Gamma_{\infty,0;10,7}^3$, respectively.

PROOF. As $B\Gamma$ is complete, $\delta \leq 3$, and so $\delta = 3$.

Let Γ' be the graph obtained from Γ by switching edges and non-edges between the two halves of Γ . If Γ' is disconnected, then by homogeneity of Γ , the relation “ $d(x, y) = 3$ ” defines a bijection between the halves, and Γ is the complement of a perfect matching. So suppose that Γ' is connected. Then the graph metric on Γ' results from the metric on Γ by interchanging distance 1 and 3.

Any isometry $A \cong B$ in Γ' corresponds to an isometry in Γ and is therefore induced by an automorphism of Γ . These automorphisms preserve the partition of Γ , possibly switching the two sides. So they act as automorphisms of Γ' as well. It follows that Γ' is a metrically homogeneous graph.

If Γ or Γ' has bounded degree, we contradict Theorem 10. Thus if A, B are the two halves of the partition of Γ , for $u \in A$ the set B_u of neighbors of u in B is infinite, with infinite complement. In particular for any two disjoint finite subsets B_1, B_2 of B , there is $u \in A$ such that u is adjacent to all vertices of B_1 and no vertices of B_2 . This is the extension property which characterizes the Fraïssé limit of the class of finite bipartite graphs, so Γ is generic bipartite. \square

9.2. Bipartite graphs with $B\Gamma \cong K_\infty[I_m]$.

PROPOSITION 9.2. *Let Γ be a connected, bipartite, and metrically homogeneous graph. Suppose that $B\Gamma \cong K_\infty[I_m]$ with $2 \leq m \leq \infty$. Then $m = 2$, and $\Gamma \cong \Gamma_{\infty,0;10,9;\emptyset}^4$, the generic antipodal bipartite graph of diameter 4.*

PROOF. The relation $d(x, y) \in \{0, 4\}$ is an equivalence relation on Γ . It follows that Γ is antipodal of diameter 4 and in particular that $m = 2$. We claim that Γ is the generic antipodal graph of diameter 4.

We call a metric subspace of Γ antipodal if it is closed under the pairing $d(x, y) = 4$ in Γ .

It suffices to show that for any antipodal subspace G of Γ and any extension $G \cup \{v\}$ of G to a graph in which all perimeters of triangles are even and bounded by 8, the extension $G \cup \{v\}$ embeds into Γ over G .

Let $X = G \cap A, Y = G \cap B$. We may suppose both are nonempty, and that v is on the B side in the sense that its distances from vertices of X are odd.

Our hypotheses imply that $d(v, y) = 2$ for all $y \in Y$, and that the neighbors of v in X form a set of representatives X_0 for the components of X . Thus to realize

this extension in Γ , it suffices to see that there are infinitely many vertices $v \in \Gamma$ adjacent to all vertices of X_0 .

As X is finite, there is an infinite subset of B consisting of vertices whose neighbors in X are identical. Furthermore, the set of neighbors in question is a set of representatives for the components of X , hence isometric with X_0 . By homogeneity, X_0 has infinitely many common neighbors in B . \square

9.3. Bipartite graphs with $B\Gamma$ treelike.

PROPOSITION 9.3. *Let Γ be a connected, bipartite, and metrically homogeneous graph. Suppose that $B\Gamma \cong T_{r,\infty}$ with $2 \leq r \leq \infty$. Then $r = \infty$ and Γ is an infinitely branching tree.*

PROOF. Let A, B be the two halves of Γ , and identify A with $B\Gamma$. In particular, for $u \in B$, the neighbors of u in A form a clique in the sense of $B\Gamma$.

Suppose that vertices u, u' in B are adjacent to two points v_1, v_2 of A . These points are contained in a unique clique of A , so all the neighbors of u, u' lie in a common clique. Therefore this gives us an equivalence relation on B , with u, u' equivalent just in case they have two common neighbors in A . But the graph structure on B is also that of $B\Gamma$, which is primitive, so this relation is trivial, and distinct vertices of B correspond to distinct cliques in A . In particular, $r = \infty$.

At the same time, every edge in A lies in the clique associated with some vertex in B , so the neighbors of the vertices of B are exactly the maximal cliques of A . Evidently the edge relation in B corresponds to intersection of cliques in A . Thus B is identified with the “dual” of A with vertices corresponding to maximal cliques, and maximal cliques corresponding to vertices. At this point the structure of Γ has been recovered uniquely from the structure of $T_{\infty,\infty}$, and must therefore be the infinitely branching tree. \square

9.4. Bipartite graphs with $B\Gamma$ the complement of a Henson Graph.

In this case, the diameter of Γ is 4 or 5. We take up the case of diameter 4 first.

PROPOSITION 9.4. *If $B\Gamma$ is the generic homogeneous graph omitting I_n and the diameter of Γ is 4, then $\Gamma \cong \Gamma_{\infty,0;14,9;\{I_n^{(4)}\}}$.*

PROOF. Let Δ, Δ' be the two halves of Γ , each isomorphic to $2B\Gamma$ (that is, isomorphic to $B\Gamma$ after rescaling the metric). We know the finite subspaces of $2B\Gamma$. Our claim is that for any finite $A, B \subseteq 2B\Gamma$ and any metric on the disjoint union $A \cup B$ in which all cross-distances between A and B are equal to either 1 or 3, there is an embedding of A into Δ and B into Δ' such that the metric of Γ induces the specified metric on $A \cup B$. Note that $A \cup B$ and Γ are both considered as metric spaces here rather than as graphs.

We will prove this by induction on the order of B . The case $k = 0$ is known so we treat the inductive step. With k fixed, we proceed by induction on the number a_4 of pairs u_1, u_2 in A with

$$d(u_1, u_2) = 4$$

So we consider an appropriate metric space $A \cup B$ with $|B| = k$. We have to treat the cases $a_4 = 0$, and the induction step for a_4 . We begin with the latter.

Suppose $a_4 > 0$ and fix $u_1, u_2 \in A$ with $d(u_1, u_2) = 4$. Take $v \in B$. We may suppose $d(v, u_1) = 3$. We adjoin a vertex a to A as follows:

$$\begin{aligned} d(a, u_1) &= 4, & d(a, u') &= 2 \text{ for } u' \in A, u' \neq u_1 \\ d(a, v) &= 1, & d(a, v') &= 3 \text{ for } v' \in A, v' \neq v \end{aligned}$$

Now the configuration $(A \cup \{a\}, B \setminus \{v\})$ embeds into Γ by induction on $k = |B|$, and the configuration $(A \setminus \{u_1\} \cup \{a\}, B)$ embeds into Γ by induction on a_4 (with k fixed). So there is an amalgam $(A \cup \{a\}, B)$ of these two configurations embedding into Γ , and the metric on this amalgam agrees with the given metric on (A, B) except possibly at the pair (u_1, v) . But we have $d(u_1, a) = 4$, $d(v, a) = 1$, so $d(u_1, v) \geq 3$, $d(u_1, v)$ is odd, and $d(u_1, v) \leq 4$. Thus $d(u_1, v) = 3$ also in the amalgam.

There remains the case in which $a_4 = 0$, so that in Γ , A is an independent set of the form $I_m^{(2)}$.

Suppose first that $k > 1$. Then we extend (A, B) to a finite configuration (A, B_1) with the following properties.

- (1) Some vertex $b \in B_1$ is adjacent to all vertices of A .
- (2) No two vertices of A have the same neighbors in B_1 .

Now consider the configurations $(\{a\}, B_1)$ for $a \in A$. If they all embed into Γ , then some amalgam does as well, and this amalgam must be isomorphic to (A, B_1) since the vertices of A must remain distinct and the metric is then determined. So it suffices to check that these configurations $(\{a\}, B_1)$ embed into Γ .

As there is an automorphism of Γ switching the two halves of its bipartition, it suffices to deal with with the configuration $(B_1, \{a\})$ instead. In this configuration, the value of $k = 1$, so we conclude by induction on k .

Now suppose

$$k = 1$$

Thus $|B| = 1$; and $a_4 = 0$ by our case hypothesis.

Take a basepoint $*$ in Γ and a vertex u in Γ_2 . It suffices to show that the set I_u of neighbors of u in Γ_1 is an infinite and coinfinite subset of Γ_1 . By Theorem 10, the set I_u is infinite. Furthermore there are by assumption vertices u_1, u_2 in Γ with $d(u_1, u_2) = 4$, and we may suppose that the basepoint $*$ lies at distance 2 from both. Then $u_1, u_2 \in \Gamma_2$ and I_{u_1}, I_{u_2} are disjoint. Thus they are coinfinite.

This completes the identification of Γ . \square

Now we turn to diameter 5. The claim in this case is as follows.

PROPOSITION 9.5. *If $B\Gamma$ is the generic homogeneous graph omitting I_n and the diameter of Γ is 5, then $n = 3$ and Γ is the generic antipodal bipartite graph of diameter 5, $\Gamma_{\infty, 0; 12, 11; \emptyset}^5$.*

We treat the cases $n = 3$ and $n > 3$ separately. Once we have Γ antipodal, Corollary 12.1 completes the analysis.

Since we are now taking the diameter to be 5, the following will allow us to simplify the statements of some results.

LEMMA 9.6. *Let Γ be a connected bipartite metrically homogeneous graph of diameter 5. Then one of the following holds.*

- (1) Γ is the cycle C_{10} .
- (2) $B\Gamma$ is the complement of a Henson graph G_n .
- (3) $B\Gamma$ is the Rado graph.

PROOF. Theorem 10 takes care of the case in which Γ_1 is finite. So we may suppose $B\Gamma$ contains an infinite clique. Proposition 9.1 eliminates the possibility that $B\Gamma$ is an infinite clique. As $B\Gamma$ is connected, Proposition 9.2 disposes of the case in which $B\Gamma$ is imprimitive. By the Lachlan/Woodrow classification, this leaves the cases in which $B\Gamma$ is the complement of a Henson graph, or the Rado graph. \square

LEMMA 9.7. *Let Γ be a metrically homogeneous bipartite graph of diameter 5 with $B\Gamma \cong G_3^c$. Then Γ is antipodal.*

PROOF. Suppose $|\Gamma_5| \geq 2$. We will show first that there is a triple (a, b, c) in Γ with $d(a, b) = 4$, $d(a, c) = 5$, and $d(b, c) > 1$.

Take a pair $u_1, u_2 \in \Gamma_5$. Then $d(u_1, u_2)$ is 2 or 4, and if it is 4 then our triple (a, b, c) can be $(u_1, u_2, *)$ with $*$ the chosen basepoint. If $d(u_1, u_2) = 2$ then extend u_1, u_2 to a geodesic (u_1, u_2, u_3) with $d(u_2, u_3) = 1$, $d(u_1, u_3) = 3$. As $d(u_2, u_3) = 1$ we find $u_3 \in \Gamma_4$ and therefore the triple $(*, u_3, u_1)$ will do.

Now fix a triple (a, b, c) with $d(a, b) = 4$, $d(a, c) = 5$, and $d(b, c) = 3$ or 5. Take a triple (b, c, d) with $d(c, d) = 1$ and $d(b, d) = 4$; this will be a geodesic of length 4 or 5, and therefore exists in Γ by homogeneity.

Now $d(a, b) = d(b, d) = 4$, and consideration of the path (a, c, d) shows that $d(a, d) \geq 4$, and $d(a, b)$ is even, so $d(a, d) = 4$ as well, and we have $I_3^{(4)}$ in Γ , a contradiction. \square

It remains to show that in an infinite metrically homogeneous bipartite graph Γ of diameter 5 for which $B\Gamma$ contains an independent set of order 3, $B\Gamma$ contains arbitrarily large independent sets. We will subdivide this case further according to the structure of Γ_5 . Since in this case Γ is not antipodal, it follows that Γ_5 is infinite. Note that if we rescale the metric on Γ_5 by 1/2 we get a homogeneous graph contained in $B\Gamma$. Thus the possible structure on Γ_5 is quite limited. The first case to be considered is that of a clique.

LEMMA 9.8. *Suppose that Γ is bipartite of diameter 5, that $I_3^{(4)}$ embeds into Γ , and that $\Gamma_5 = I_\infty^{(2)}$. Then $I_\infty^{(4)}$ embeds in $B\Gamma$ and $B\Gamma$ is the Rado graph.*

PROOF. Our claim is that $I_n^{(4)}$ embeds into Γ for all n ; then since $\delta = 5$, $B\Gamma$ is a connected homogeneous graph containing an infinite clique and an infinite independent set. Proposition 9.2 implies that $B\Gamma$ is not of the form $K_\infty[I_\infty]$ and then the Lachlan/Woodrow classification leaves only the Rado graph.

We proceed by induction, with the case $n = 3$ assumed.

Suppose $I_n^{(4)}$ embeds into Γ , with $n \geq 3$. Let $I \cong I_{n-1}^{(4)}$ be a metric subspace of Γ . We aim to embed subspaces $A = I \cup \{a, u\}$ and $B = I \cup \{b, u\}$ into Γ , with $I \cup \{a\} \cong I \cup \{b\} \cong I_n^{(4)}$, and with u chosen so that

$$\begin{aligned} d(u, a) &= 1 & d(u, b) &= 5 \\ d(u, x) &= 3 & (x \in I) \end{aligned}$$

Supposing we have this, considering (a, u, b) we see that $d(a, b) \geq 4$ and hence $I \cup \{a, b\} \cong I_{n+1}^{(4)}$.

We treat the second factor $I \cup \{b, u\}$ first. Consider the metric space $I \cup \{b, b'\}$ in which b' lies at distance 2 from each point of $I \cup \{b\}$. The corresponding configuration in $B\Gamma$ is a point b' adjacent to an independent set of order n , and this we have in $B\Gamma$. Thus the space $I \cup \{b, b'\}$ embeds into Γ .

By hypothesis, there is also a triple (u, b, b') with $b, b' \in \Gamma_5(u)$. Amalgamate $I \cup \{b, b'\}$ with (u, b, b') over b, b' . For $x \in I$, considering (u, b', x) , we see that $d(u, x) \geq 3$, and that $d(u, x)$ is odd. Our hypothesis on Γ_5 implies that $d(u, x) \neq 5$, so $d(u, x) = 3$ for all $x \in I$. Thus $I \cup \{b, u\}$ is as desired.

The construction of the factor $I \cup \{a, u\}$ is more elaborate. Consider the metric spaces $A = I \cup \{a\} \cup J$ and $B = J \cup \{a, u\}$, where $J \cong I_n^{(2)}$, and J may be labeled as $\{v^* : v \in I\}$ in such a way that

$$\begin{aligned} d(a, v^*) &= d(v, v^*) = 4 && (v \in I) \\ d(v, w^*) &= 2 && v, w \in I \text{ distinct} \\ d(u, v) &= 5 && (v \in J) \end{aligned}$$

Supposing that A and B embed into Γ , take their amalgam over $J \cup \{a\}$. Then for $v \in I$ the triple (u, a, v) shows that $d(u, v) \geq 3$, and the distance is odd, while the triple (u, v', v) and the hypothesis on Γ_5 shows that this distance is not 5. Thus the space $I \cup \{a, u\}$ will have the desired metric. It remains to construct A and B .

Consider $B = J \cup \{a, u\}$. The graph Γ contains an edge (a, u) as well as a copy of $J \cup \{u\}$, the latter by the hypothesis on Γ_5 . Furthermore, in any amalgam of (a, u) with $J \cup \{u\}$, the only possible value for the distance $d(a, v)$, for $v \in J$, is 4. So this disposes of B .

Now consider $A = I \cup J \cup \{a\}$, in which all distances are even. So we need to look for the rescaled graph $(1/2)A$ in $B\Gamma$. It suffices to check that the maximal independent sets of vertices in $(1/2)A$ have order at most n . This is the case for $I \cup \{a\}$, and any independent set meeting J would have order at most 3. Since $n \geq 3$, we are done. \square

LEMMA 9.9. *Suppose that Γ is metrically homogeneous, bipartite, infinite, not antipodal, and of diameter 5. Then Γ_5 contains a subspace of the form $I_\infty^{(2)}$; in other words, $(1/2)\Gamma_5$ contains an infinite clique.*

PROOF. Supposing the contrary, for each $u \in \Gamma_4$, the set I_u of neighbors of u in Γ_5 is finite and nonempty, of fixed order k . Since any subset of Γ_5 isomorphic to $I_{k+1}^{(2)}$ would have a common neighbor $u \in \Gamma_4$, it follows that the I_u represent maximal cliques of $(1/2)\Gamma_5$.

As $B\Gamma$ is either generic omitting I_n for some $n \geq 3$, or the Rado graph, it follows that $(1/2)\Gamma_4$ is primitive and contains both edges and nonedges. Now the map $u \rightarrow I_u$ induces an equivalence relation on Γ_4 which can only be equality, that is the map is a bijection. Since $(1/2)\Gamma_4$ contains both edges and non-edges, it follows that Γ_5 is primitive. As each vertex $v \in \Gamma_5$ has infinitely many neighbors in Γ_4 , we have $|I_u| > 1$ for $u \in \Gamma_4$. On the other hand if $|I_u| \geq 3$ then for $u, u' \in \Gamma_4$ we have the possibilities $|I_u \cap I_{u'}| = 0, 1, 2$ while there are only two distances occurring in Γ_4 . So $|I_u| = 2$. That is, $(1/2)\Gamma_5$ is generic triangle-free, and the vertices of Γ_4 correspond to edges of $(1/2)\Gamma_5$. It follows that vertices of Γ_4 lie at distance two iff the corresponding edges meet, that is $(1/2)\Gamma_4$ is the line graph of $(1/2)\Gamma_5$. But there are pairs of vertices in the latter graph at distance greater than 2, so $(1/2)\Gamma_4$ is not homogeneous, and we have a contradiction. \square

LEMMA 9.10. *Suppose that Γ is bipartite, metrically homogeneous, and of diameter 5, and Γ_5 contains a pair of vertices at distance 4. Then the relation*

$$"d(x, y) = 0 \text{ or } 4"$$

is not an equivalence relation on Γ_5 .

PROOF. Supposing the contrary, we have

$$\Gamma_5 \cong I_\infty^{(2)}[I_k^{(4)}]$$

for some k with $2 \leq k \leq \infty$.

Suppose first $k \geq 3$. Fix two equivalence classes C, C' in Γ_5 , and choose a triple u_1, u_2, u_3 in C_1 and a vertex u'_1 in C' . Choose $v \in \Gamma$ with $d(v, u_1) = d(v, u'_1) = 1$, and let $d_i = d(v, u_i)$ for $i = 2, 3$. We may then choose u'_2, u'_3 in C' so that $d(v, u'_i) = d_i$ for $i = 2, 3$.

Now the permutation of the u_i, u'_i which switches u'_1 and u'_2 and fixes the other elements is an isometry, so there is an element v' with $d(v', u) = d(v, u)$ for $u = u_1, u_2, u_3, u'_3$, but with $d(v', u'_1) = d_2$, $d(v', u'_2) = 1$.

As u_1 is adjacent to v, v' we have $d(v, v') = 2$. Now u_3, v, v' is isometric with u'_3, v, v' , and the equivalence class of u_3 contains a common neighbor of v, v' ; therefore the equivalence class of u'_3 contains a common neighbor of v, v' . But v can have at most one neighbor in an equivalence class, so this contradicts the choice of v' .

So we are left with the case $k = 2$:

$$\Gamma_5 \cong I_\infty^{(2)}[I_2^{(4)}]$$

In this case we will consider a specific amalgamation.

Let $\gamma = (u, v, w)$ be a geodesic with

$$d(u, v) = 1; d(v, w) = 4; d(u, w) = 5$$

Let $A = \gamma \cup \{a\}$, $B = \gamma \cup \{b\}$, with the metrics given by

	u	v	w
a	4	3	5
b	4	5	5

It is easy to see that A, B embed into Γ . In an amalgama of the two, the relation of a and b to v prevents them from being identified in the amalgam. However $a, b, u \in \Gamma_5(w)$ and $d(a, u) = d(b, u) = 4$. So $d(a, b) = 4$ by our assumption, and this contradicts $k = 2$. \square

LEMMA 9.11. *Suppose that Γ is bipartite, metrically homogeneous, and of diameter 5. Then Γ_i is connected with respect to the edge relation given by $d(x, y) = 2$, for $1 \leq i \leq 5$.*

PROOF. This is true for Γ_1 automatically. It is true for $i = 2$ or 4 in view of the structure of $B\Gamma$. It remains to prove it for $i = 3$ or 5.


If Γ_i is disconnected with respect to this relation, then for $u \in \Gamma_{i-1}$, the set I_u of neighbors of u in Γ_i is contained in one of the equivalence classes of Γ_i , and there is more than one such class. Thus we have a function from Γ_{i-1} to the quotient of Γ_i . As $i - 1$ is even, in view of the structure of $B\Gamma$ we know Γ_{i-1} is primitive, so as Γ_i contains more than one equivalence class, this function is 1 - 1. Then the sets I_u for $u \in \Gamma_{i-1}$ must be exactly the equivalence classes of Γ_i , and Γ_{i-1} is in bijection with the quotient. In particular, only one distance occurs in Γ_{i-1} . But in view of the structure of $B\Gamma$, this is not the case. \square

COROLLARY 9.11.1. *Suppose that Γ is bipartite, metrically homogeneous, and of diameter 5, and Γ_5 contains a pair of vertices at distance 4. Then Γ_5 is primitive, infinite, and contains a copy of $I_\infty^{(2)}$.*

LEMMA 9.12. *Suppose that Γ is bipartite, metrically homogeneous, and of diameter 5, that Γ_5 contains a pair of vertices at distance 4, and that $B\Gamma$ is generic omitting I_{n+1} . Then $I_{n-1}^{(4)}$ embeds into Γ_5 .*

PROOF. Let k be maximal so that $I_k^{(4)}$ embeds into Γ_5 , and suppose $k \leq n-2$.

Let $I \cong I_{n-1}^{(4)}$, and suppose a, b, u are additional vertices with $I \cup \{a\} \cong I \cup \{b\} \cong I_n^{(4)}$, and with $d(u, a) = 1$, $d(u, b) = 5$, $d(u, v) = 3$ for $v \in I$. If $I \cup \{a, u\}$ and $I \cup \{b, u\}$ embed into Γ , then so does an amalgam $I \cup \{a, b, u\}$, and the auxiliary vertex u forces $d(a, b) = 4$, and $I \cup \{a, b\} \cong I_{n+1}^{(4)}$, a contradiction as $B\Gamma$ omits I_{n+1} . So it suffices to embed $I \cup \{a, u\}$ and $I \cup \{b, u\}$ into Γ .



Construction of $I \cup \{a, u\}$.

Introduce a metric space $J = \bigcup_{v \in I} J_v$ with $J_v \cong I_k^{(4)}$ and with $d(x, y) = 2$ for $x \in J_v, y \in J \setminus J_v$. Extend to a metric on $I \cup J$ by taking

$$d(v, x) = \begin{cases} 4 & \text{if } x \in J_v \\ 2 & \text{if } x \in J \setminus J_v \end{cases}$$

for $v \in I$.

Give $J \cup \{a, u\}$ the metric with $d(a, x) = 4$, $d(u, x) = 5$ for $x \in J$. We claim that $I \cup J \cup \{a\}$ and $J \cup \{a, u\}$ embed into Γ . Now $B\Gamma$ is generic omitting I_{n+1} , and $(1/2)\Gamma_5$ is generic omitting I_{k+1} . Since the space $I \cup J \cup \{a\}$ does not contain $I_{n+1}^{(4)}$, and all its distances are even, it embeds into Γ . Since J does not contain $I_{k+1}^{(4)}$, it embeds into Γ_5 , so $J \cup \{u\}$ embeds into Γ . In any amalgam of $J \cup \{u\}$ with $\{a, u\}$ we have $d(a, x) = 4$ for $x \in J$, so $J \cup \{a, u\}$ embeds into Γ as well.

Thus an amalgam of $I \cup J \cup \{a\}$ and $J \cup \{a, u\}$ embeds into Γ . For $v \in I$, consideration of (u, a, v) shows that $d(u, v)$ is 3 or 5, and consideration of $J_v \cup \{u, v\}$ shows that $d(u, v)$ is not 5. Thus we have $d(u, v) = 3$ for all $v \in I$ in our amalgam, and thus $I \cup \{a, u\}$ embeds isometrically into Γ .

Construction of $I \cup \{b, u\}$.

Let $J' = \bigcup_{v \in I} J'_v$ with $J'_v \cong I_{k-1}^{(4)}$. Put a metric on $I \cup J' \cup \{b, u\}$ by taking $d(u, x) = 5$, $d(b, x) = 4$ for $x \in J'$, while for $v \in I$ we take $d(v, x) = 4$ for $x \in J'_v$, and $d(v, x) = 2$ for $x \in J \setminus J'_v$.

Introduce an auxiliary vertex b' with $d(b', u) = 5$, $d(b', x) = 2$ for $x \in I \cup J' \cup \{b\}$.

We claim that $I \cup J' \cup \{b, b'\}$ and $J' \cup \{u, b, b'\}$ embed isometrically in Γ . For $I \cup J' \cup \{b, b'\}$ we use the structure of $B\Gamma$, together with the condition $k+1 < n$, and for $J' \cup \{u, b, b'\}$ we use the structure of Γ_5 to check that $J \cup \{b, b'\}$ embeds into Γ_5 .

Therefore some amalgam $I \cup J' \cup \{u, b, b'\}$ embeds into Γ . Let $v \in I$. In the amalgam, the auxiliary vertex b' ensures that $d(u, v)$ is 3 or 5. Consideration of $J'_v \cup \{u, b, v\}$ shows that $d(u, v)$ is not 5. Thus $d(u, v) = 3$ for $v \in I$, and $I \cup \{b, u\}$ embeds isometrically in Γ . \square

We will need some additional amalgamation arguments to complete our analysis, beginning with the following preparatory lemma.

LEMMA 9.13. *Let Γ be bipartite of diameter 5, and not antipodal. Suppose $B\Gamma$ is generic omitting $I_{n+1}^{(4)}$, with $n \geq 3$, and Γ_5 contains a pair of vertices at distance 4. Then the following hold.*

- (1) $I_n^{(4)}$ embeds in Γ_3 ;
- (2) Γ_3 is primitive.

PROOF.

1. $I_n^{(4)}$ embeds in Γ_3 :

We show inductively that $I_m^{(4)}$ embeds into Γ_3 for $m \leq n$.

Let $I \cong I_{m-1}^{(4)}$. Form extensions $I \cup \{u\}$ and $I \cup \{v\}$ with $d(u, x) = 5$, $d(v, x) = 3$ for $x \in I$. Then $I \cup \{u\}$ embeds into Γ since $m - 1 \leq n - 1$, while $I \cup \{v\}$ embeds into Γ by induction on m . So some amalgam $I \cup \{u, v\}$ embeds into Γ with $d(u, v)$ either 2 or 4.

Consider a geodesic $\{u, v, w\}$ with $d(u, w) = 1$, $d(v, w) = 3$, and $d(u, v)$ as specified. There is an amalgam $I \cup \{u, v, w\}$ of $I \cup \{u, v\}$ with $\{u, v, w\}$ over u, v , and consideration of (w, u, x) for $x \in I$ show that $I \cup \{w\} \cong I_m^{(4)}$. As $I \cup \{w\} \subseteq \Gamma_3(v)$, the induction is complete.

2. Γ_3 is primitive:

By Lemma 9.11 we have $(1/2)\Gamma_3$ connected. Suppose now that Γ_3 is disconnected with respect to the edge relation “ $d(x, y) = 4$.”

Fix two connected components C, C' with respect to this relation. By (1) these have order n , and by assumption $n \geq 3$. Fix $u \in C$ and $u'_1, u'_2 \in C'$, and $v_1 \in \Gamma_2$ adjacent to u, u'_1 . With $*$ the chosen basepoint for Γ , consider the isometry of $C \cup C' \cup \{*\}$ which interchanges u'_1 and u'_2 and fixes the remaining vertices. Then this extends to an isometry $C \cup C' \cup \{*, v\} \cong C \cup C' \cup \{*, v'\}$ for some vertex v' . Take $u'_3 \in C'$, distinct from u_1, u_2 . Then the map $(*, u_3, u'_3, v, v') \mapsto (*, u'_3, u_3, v, v')$ is an isometry and therefore extends to Γ ; its extension interchanges C and C' and fixes v, v' . However $d(v, x) = d(v', x)$ for $x \in C$, so the same applies to C' . But $d(v, u'_1) = 1$, $d(v', u'_2) = 1$, and $d(u'_1, u'_2) = 4$, so this is impossible. \square

Now we can assemble these ingredients.


LEMMA 9.14. *If Γ is bipartite of diameter 5 and not antipodal, then $B\Gamma$ is the universal homogeneous graph (Rado’s graph).*

PROOF. The alternative is that $B\Gamma$ is generic omitting I_{n+1} for some $n \geq 3$. By Lemma 9.8, we may suppose that Γ_5 contains a pair of vertices at distance 4, and hence $I_{n-1}^{(4)}$ embeds in Γ_5 by Lemma 9.12.

To get a contradiction, we will aim at an amalgamation of the following form. Let $I \cong I_{n-1}^{(4)}$, let $I \cup \{a\} \cong I \cup \{b\} \cong I_n^{(4)}$, and adjoin a vertex u such that

$$d(u, a) = 1; d(u, b) = 5; d(u, x) = 3 \text{ for } x \in I$$

We will embed $I \cup \{a, u\}$ and $I \cup \{b, u\}$ in Γ , and then in their amalgam we will have $I \cup \{a, b\} \cong I_{n+1}^{(4)}$, a contradiction. Each of the factors $I \cup \{a, u\}$ and $I \cup \{b, u\}$ will require its own construction.



Construction of the first factor, $I \cup \{a, u\}$.

Let $I = I_0 \cup \{c\}$ and introduce a vertex v with

$$\begin{aligned} d(v, a) = 1 & & d(v, c) = 5 & & d(v, u) = 2 \\ d(v, x) = 4 & & (x \in I_0 \cup \{a\}) & & \end{aligned}$$

If $I_0 \cup \{a, u, v\}$ and $I_0 \cup \{c, u, v\}$ embed into Γ , then in their amalgam $I_0 \cup \{a, c, u, v\}$ we have $d(a, c) = 4$ and thus the desired metric space $I \cup \{a, u\}$ is embedded into Γ .

Construction of $I_0 \cup \{a, u, v\}$.

We first embed $I_0 \cup \{a, u\}$ into Γ . Introduce a vertex a' with

$$d(a', u) = 1; d(a', a) = 2; d(a', x) = 2 \text{ for } x \in I_0$$

On the one hand, the geodesic (a, u, a') embeds into Γ ; on the other hand, the metric space $I_0 \cup \{a, a'\}$ embeds into $B\Gamma$. So some amalgam $I_0 \cup \{a, a', u\}$ embeds into Γ , and for $x \in I_0$, consideration of the paths (u, a, x) and (u, a', x) shows that $d(u, x) = 3$, as required. Thus $I_0 \cup \{a, u\}$ embeds into Γ .

Take a as basepoint. Then $u \in \Gamma_1$, $I_0 \subseteq \Gamma_4$, and $d(u, x) = 3$ for $x \in I_0$. Let $I_1 \subseteq I_0$ be obtained by removing one vertex, so $I_1 \cong I_{n-3}^{(4)}$. Consider the sets

$$\begin{aligned} A &= \{u \in \Gamma_1 : d(u, x) = 3 \text{ for } x \in I_1\} \\ B &= \{u \in \Gamma_4 : d(u, x) = 4 \text{ for } x \in I_1\} \end{aligned}$$

The partitioned metric space (A, B) is homogeneous with respect to the metric plus the partition. We consider the structure of (A, B) .

We show first that A is infinite. Assuming the contrary, consider a configuration $I_1 \cup I_2$ in $I_2 \cong I_\infty^{(2)}$, and $I_1 \cup \{x\} \cong I_{n-2}^{(4)}$ for each $x \in I_2$. This configuration embeds into Γ_4 . There is a pair $x, y \in I_2$ such that $I_1 \cup \{x\}$ and $I_1 \cup \{y\}$ have the same vertices at distance 3 in Γ_1 . Now Γ_4 is generic omitting $I_n^{(4)}$. By homogeneity it follows easily that any two subsets of Γ_4 isomorphic to $I_{n-2}^{(4)}$ have the same vertices at distance 3 in Γ_1 . This yields a nonempty subset of Γ_1 definable without parameters, and a contradiction. So A is infinite.

Now B is generic omitting $I_3^{(4)}$. In particular B is primitive. Furthermore, each vertex of B lies at distance 3 from some vertex of A . By primitivity, this vertex cannot be unique. Take $c \in B$ and $u, v \in A$ so that $d(c, u) = d(c, v) = 3$. Then $I_0 \cup \{c, a, u, v\}$ has the desired structure.

Construction of $I_0 \cup \{c, u, v\}$.

We introduce a vertex d at distance 1 from u and v , and distance 4 from c , with the relation of d to I_0 will be determined below.

In any amalgam of $I_0 \cup \{c, d, u\}$ with $I_0 \cup \{c, d, v\}$ over $I_0 \cup \{c, d\}$ we have $d(u, v) = 2$. It remains to construct $I_0 \cup \{c, d, u\}$ and $I_0 \cup \{c, d, v\}$.

We claim first that $I_0 \cup \{c, u\}$ embeds into Γ , in other words that $I_0 \cup \{c\}$ embeds into $\Gamma_3(u)$. This holds by Lemma 9.13. Now we may form $I_0 \cup \{c, d, u\}$ by amalgamating $I_0 \cup \{c, u\}$ with $\{c, d, u\}$ (a geodesic) to determine the metric on $I_0 \cup \{d\}$; all distances $d(x, d)$ will be even for $x \in I_0$. This amalgamation determines the structure of $I_0 \cup \{d\}$ and thereby completes the determination of the second factor $I \cup \{c, d, v\}$ as well.

We claim that $I \cup \{c, d, v\}$ embeds into Γ . Since the distance $d(c, d) = 4$ is forced in any amalgam of $I_0 \cup \{v, c\}$ with $I_0 \cup \{v, d\}$, we consider these two metric spaces separately.

Now $I_0 \cup \{v, d\} \cong I_0 \cup \{u, d\}$, so this is not at issue, and we are left only with $I_0 \cup \{c, v\}$. This last embeds into the second factor $I \cup \{b, u\}$, so we may turn finally to a consideration of this second factor.

Construction of the second factor, $I \cup \{b, u\}$.

We introduce another vertex v satisfying

$$d(v, b) = 1; d(v, u) = 4; d(v, x) = 5 \text{ for } x \in I$$

This will force $d(b, x) = 4$ for $x \in I$. So it will suffice to embed $I \cup \{u, v\}$ and $\{b, u, v\}$ separately into Γ . Since $\{b, u, v\}$ is a geodesic, we are concerned with $I \cup \{u, v\}$.

Introduce a vertex d with

$$d(u, d) = 1; d(v, d) = 5; d(x, d) = 2 \text{ for } x \in I$$

Then amalgamation of $I \cup \{d, u\}$ with $I \cup \{d, v\}$ forces $d(u, v) = 4$. It remains to embed $I \cup \{d, u\}$ and $I \cup \{d, v\}$ into Γ .

The second of these, $I \cup \{d, v\}$, has a simple structure with $I \cup \{d\} \subseteq \Gamma_5(v)$, and since $I \cup \{d\}$ has order n , with all distances even, it embeds into Γ_5 by Lemma 9.12. So we need only construct $I \cup \{d, u\}$.

Taking d as base point, and I contained in Γ_2 , we are looking for a vertex $u \in \Gamma_1$ at distance 3 from all elements of I . For $v \in I$, let I_v be the set of neighbors of v in Γ_1 . Any vertex $u \in \Gamma_1$ which is not in $\bigcup_{v \in I} I_v$ will do. So it remains to be checked that $\bigcup_{v \in I} I_v \neq \Gamma_1$.

The sets I_v for $v \in I$ are pairwise disjoint. Suppose they partition Γ_1 . We may take a second set $J \cong I_{n-1}^{(4)}$ in Γ_2 overlapping with I so that $|I \cap J| = n - 2$, and then the I_v for $v \in J$ will also partition Γ_1 ; so the vertices $v_1 \in I \setminus J$ and $v_2 \in J \setminus I$ have the same neighbors in Γ_1 . As Γ_2 is primitive, it follows that all vertices of Γ_2 have the same neighbors in Γ_1 , a contradiction. \square

At this point, the proof of Proposition 9.5, and also Theorem 13, is complete. We review the analysis.

Proof of Theorem 13. $B\Gamma_1$ falls under the Lachlan/Woodrow classification.

If Γ_1 is finite, then Theorem 10 applies, and as we assume diameter at least 3 and degree at least 3, we arrive at either the complement of a perfect matching or a tree in this case.

With Γ_1 infinite, $B\Gamma$ contains an infinite clique, and hence so does $B\Gamma_1$. As noted at the outset, $B\Gamma$ is connected. So by the Lachlan/Woodrow classification, $B\Gamma$ is either imprimitive of the form $K_\infty[I_m]$ or $m \cdot K_\infty$, $2 \leq m \leq \infty$, or generic omitting I_n for some finite $n \geq 2$, or universal homogeneous (Rado's graph).

When $B\Gamma_1$ is imprimitive, the classification in Theorem 10 applies to $B\Gamma$, and as $B\Gamma$ contains an infinite clique, the result is that $B\Gamma$ is either an imprimitive homogeneous graph of the form $K_\infty[I_m]$ ($2 \leq m \leq \infty$), or one of the tree-like graphs of Macpherson, $T_{r,\infty}$ with $2 \leq r \leq \infty$.

When $B\Gamma$ is of the form $K_\infty[I_m]$ with $m \geq 2$, Lemma 9.2 applies. When $B\Gamma$ is tree-like, Lemma 9.3 applies, and Γ is a tree.

Thus we may suppose that $B\Gamma_1$ is primitive. We have set aside the case in which $B\Gamma_1$ is universal homogeneous as a distinct (and typical) case. So we are left with the possibility that $B\Gamma_1$ is generic omitting I_n with $2 \leq n < \infty$. In view of Theorem 10, $B\Gamma$ must have diameter at most 2, and be homogeneous as a graph. Then our hypothesis on $B\Gamma_1$ implies that $B\Gamma$ is also generic omitting I_n .

In case $n = 2$, Lemma 9.1 applies. If $n > 2$ then the diameter of Γ is 4 or 5. If the diameter is 4, then Lemma 9.4 applies, and leads to case 4 of the theorem.

This leaves us with the case taken up in Proposition 9.5: Γ has diameter 5, and $B\Gamma$ is generic omitting I_n . As the diameter is 5, we have $n \geq 3$. As $B\Gamma_1$ does not contain I_∞ , Γ is antipodal by Lemma 9.14. So Corollary 12.2 applies. \square

10. Graphs of small diameter

In the Appendix to [Che98], we gave an exhaustive list of certain amalgamation classes for highly restricted languages. The languages considered were given by a certain number of irreflexive binary relations, symmetric or asymmetric, with the proviso that every pair of distinct elements satisfies one and only one of the given relations. The cases of interest here are the languages with either 3 or 4 symmetric irreflexive binary relations. The amalgamation classes \mathcal{A} listed were those satisfying the following three conditions:

- (1) The class \mathcal{A} is determined by a set of forbidden triangles.
- (2) The Fraïssé limit of the class is primitive.
- (3) The class in question is not a free amalgamation class.

This last point means that there is no single relation $R(x, y)$ such that every amalgamation problem $A_0 \subseteq A_1, A_2$ can be completed by taking R to hold between $A_1 \setminus A_0$ and $A_2 \setminus A_0$. This excludes some readily identified metrically homogeneous graphs of diameter 3, but none of greater diameter.

For example, in the case of three symmetric relations $A(x, y), B(x, y), C(x, y)$, the only such class (up to a permutation of the language) is the one given by the following constraints:

$$(AAB), (ACC), (AAA)$$

In this notation, AAB represents a triple x, y, z with $A(x, y), A(x, z)$, and $B(y, z)$.

Now for the Fraïssé limit of this class to correspond to a metric space of diameter 3, one of the forbidden configurations must correspond to the triangle (113). Since the configuration corresponding to (113) may be either (AAB) or (ACC) , there are two distinct homogeneous metric spaces of this type, with the following constraints:

$$(113), (122), (111) \text{ or } (233), (113), (333)$$

The first of these has no triangle of odd perimeter 5 or less, that is $K_1 = K_2 = 3$, $C_0 = 10$, $C_1 = 11$. The second has no triangle of perimeter 8 or more, that is $K_1 = 1$, $K_2 = 3$, $C_0 = 8$, $C_1 = 9$. We recognize these as falling within our previous classification with $C > 2\delta + K_1$ and $C' = C + 1$.

Of course, any of our examples of type $\mathcal{A}_{K_1, K_2; C_0, C_1}^3$ will be a free amalgamation class unless one of the forbidden triangles involves the distance 2, which means:

$$K_1 = 3 \text{ or } C \leq 8$$

And as our list was confined to the primitive case it omits bipartite and antipodal examples.

We view the classification of the amalgamation classes determined by triangles as a natural ingredient of a full catalog of “known” types, and a natural point of departure for an attempt at a full classification. In [AMp10] (a working draft) the problem is taken from the other end, in the case of diameter 3, and it is shown that in the triangle-free case (i.e., K_3 -free), the classification does indeed reduce to identifying the combinations of triangle constraints and $(1, 2)$ -space constraints which (jointly) define amalgamation classes.

10.1. Diameter 4. The explicit classification of the metrically homogeneous graphs of diameter 4 whose minimal forbidden configurations are triangles is more complex. Once the diameter δ exceeds 3, the possibility of “free amalgamation” falls by the wayside, as one can use an amalgamation to force any particular distance strictly between 1 and δ , using the triangle inequality. So in diameter 4, the table in [Che98] covers all primitive homogeneous structures with 4 symmetric 2-types which can be given in terms of forbidden triangles.

There are 27 such (up to a permutation of the language) of which 17 correspond to homogeneous metric spaces, some in more than one way (permuting the distances matters to us, if not to the theory). We will exhibit those classes in a number of formats. Table 4, at the end of the present paper, gives all 27 classes in the order they were originally given, using the symbols A, B, C, D for the binary relations involved. The numbering of cases used in the next two tables conforms to the numbering given in that table. Note that the entry “1” means the given configuration is included in the constraint set, and hence is *omitted* by the structures under consideration.

In Table 1 we have converted A, B, C, D to distances, usually in the order 1, 2, 3, 4, or 4, 2, 3, 1, and in one case in the order 2, 4, 3, 1, for those cases in which the result is a metric space. In checking the possibilities, begin by identifying the forbidden triangle $(1, 2, 4)$, involving three distinct distances; in all 27 cases this can only be the triple (ABD) , by inspection. Thus C corresponds to distance 3. After that, look for the forbidden triangle (113) : by inspection, this is either CDD or AAC in each case. Thus A or D corresponds to distance 1, after which there is at most one assignment of distances that produces the constraint (114) .

All primitive metrically homogeneous graphs of diameter 4 whose constraints are all of order 3 are listed in the resulting table. We omit the columns corresponding to the non-geodesic triangles of types $(1, 2, 4)$, $(1, 1, 3)$, and $(1, 1, 4)$, which are of course present as constraints in all cases.

In Table 2 we list these metric spaces together with their defining parameters K_1, K_2, C, C' .

Let us compare the outcome to the statement of Theorem 9.

The table contains no examples with $C \leq 2\delta + K_1$. The case $C' > C + 1$ and $3K_2 = 2\delta - 1$ is impossible with $\delta = 4$, while the case $K_1 = \infty$ is imprimitive and omitted. On the other hand, when $C' = C + 1$, if $C = 2\delta + 1$ then again the graph is imprimitive, while if $C \geq 2\delta + 2$ then $K_1 \geq 2$, and the condition

#	111	122	133	144	223	244	334	444	344	A,B,C,D
23'	0	0	0	1	0	0	0	1	1	1,2,3,4
7	0	0	0	1	0	0	0	1	0	4,2,3,1
3	0	0	0	1	0	0	0	0	0	4,2,3,1
25	1	0	0	0	0	0	0	1	1	1,2,3,4
25'	1	0	0	0	0	0	0	1	1	4,2,3,1
8	1	0	0	1	0	0	0	1	0	4,2,3,1
5	1	0	0	1	0	0	0	0	0	4,2,3,1
18	1	1	0	1	0	0	0	1	0	4,2,3,1
16	1	1	0	1	0	0	0	0	0	4,2,3,1
26'	0	0	0	0	0	1	1	1	1	1,2,3,4
21'	0	0	0	0	0	0	0	1	1	1,2,3,4
4	0	0	0	0	0	0	0	1	0	4,2,3,1
1	0	0	0	0	0	0	0	0	0	4,2,3,1
24	1	0	0	0	0	0	0	1	1	4,2,3,1
6	1	0	0	0	0	0	0	1	0	4,2,3,1
2	1	0	0	0	0	0	0	0	0	4,2,3,1
17	1	1	0	0	0	0	0	1	0	4,2,3,1
15	1	1	0	0	0	0	0	0	0	4,2,3,1
26	1	1	1	0	1	0	0	0	0	2,4,3,1
22	0	0	0	1	0	0	0	0	1	4,2,3,1
22'	0	0	0	1	0	0	0	0	1	1,2,3,4
23	1	0	0	1	0	0	0	0	1	4,2,3,1

TABLE 1. 22 metric spaces, with duplication

$K_1 + 2K_2 \leq 2\delta - 1 = 7$ gives $K_2 = K_1 = 2$, and hence $C = 2(K_1 + K_2) + 1 = 2\delta + 1$ after all.

So what we see here is the range of possibilities illustrating the third case under Theorem 9: $K_1 < \infty$, $C > 8 + K_1$, $K_1 + 2K_2 \geq 7$, $3K_2 \geq 8$, and if $K_1 + 2K_2 = 7$ then $C \geq 10 + K_1$, while if $C' > C + 1$ then $C \geq 8 + K_2$.

If $C' = C + 1$ these constraints amount to $1 \leq K_1 \leq K_2$, $K_2 = 3$ or 4 , $C \geq 9 + K_1$, and if $K_1 = 1$ and $K_2 = 3$ then $C \geq 11$. This corresponds to the first two sections of the table, arranged according to increasing K_2 .

If $C' > C + 1$ then there is the added constraint $C \geq 8 + K_2$ and this implies $C = 11$, $C' = 14$, $K_2 = 3$. Since $C > 2\delta + K_1$ we find $K_1 \leq 2$. This corresponds to the last two lines of the table.

Part III. Extension Properties of Finite Triangle Free Graphs

11. Extension properties

A structure has the *finite model property* if every first order sentence true in the structure is true in some finite structure. A slightly stronger property is the *finite submodel property*, where the finite approximation should be taken to lie within the original structure. There is little difference between the notions in the cases of most immediate concern here. We focus mainly on the finite model property for the generic triangle free graph, and as “triangle free” is part of the first order

#	K_1	K_2	C	C'
23'	1	3	11	12
7	1	3	12	13
3	1	3	13	14
25, 25'	2	3	11	12
8	2	3	12	13
5	2	3	13	14
18	3	3	12	13
16	3	3	13	14
26'	1	4	10	11
21'	1	4	11	12
4	1	4	12	13
1	1	4	13	14
24	2	4	11	12
6	2	4	12	13
2	2	4	13	14
17	3	4	12	13
15	3	4	13	14
26	4	4	13	14
22, 22'	1	3	11	14
23	2	3	11	14

TABLE 2. 20 metric spaces, sorted

theory of this structure, the finite model property and the finite submodel property are equivalent here. We do not necessarily expect the finite submodel property to be true. The evidence in either direction is meager. However that may be, the problem is a concrete one connected with problems in finite geometries, and we will take pains to put it in an explicit form.

While the finite model problem has attracted the attention of a number of combinatorialists and probabilists, the main conclusion to date is that it is elusive, and there is hardly any literature to be found on the subject. Here I assemble and document the main information that has come my way about the problem.

There is some literature on lower bounds for finite approximations to Fraïssé limits, in cases where probabilistic considerations guarantee their existence, e.g. [Sz²65, Bon09, Bon10]. Such explorations may also shed light on cases where the existence of any such approximation remains in question. Below we will discuss some explorations of that sort in the case of triangle free graphs, where we see clear possibilities for some meaningful analysis.

11.1. Extension Properties E_n, E'_n, Adj_n . In general, the theory of a homogeneous locally finite structure is axiomatized by two kinds of axioms: negative axioms defining the “forbidden substructures,” e.g. triangles in the case that most concerns us, and extension properties stating that for each k , any k -generated subset of k elements may be extended to a $(k + 1)$ -generated subset in any way not explicitly ruled out by the negative constraints. We will drop all further mention of k -generated structures here, and suppose that the language is purely relational, so that the issue is one of extending k given elements by one more element.

In the context of triangle free graphs, the natural extension properties are the following three.

(E_n) : for any set A of at most n vertices, and any subset B of A consisting of independent vertices, there is a vertex v adjacent to all vertices of B , and to none of $A \setminus B$.

(E'_n) : (a) any maximal independent set contains at least n vertices; and (b) for any set A of n independent vertices, and any subset B of A , there is a vertex v adjacent to all vertices of B , and to none of $A \setminus B$.

(Adj_n) : Any set of at most n independent vertices has a common neighbor.

Here the properties E_n , together with the axiom stating that triangles do not occur, gives the full axiomatization of the generic triangle free graph. Therefore the problem of the finite model property for the generic triangle free graph is simply the question, whether for each n some finite triangle free graph has the extension property E_n .

The equivalence of E_n and E'_n in triangle free graphs is straightforward, with the former condition more easily applied, and the latter more readily checked. The mutual adjacency condition Adj_n is manifestly weaker, but only because it allows some relatively degenerate examples: the complete bipartite graph, and some less obvious ones—these can all be satisfactorily classified. As a result, we can replace the property E'_n by a mild strengthening of Adj_n , giving us the simplest version to check.

Our first order of business will be to sort out the force of these extension properties. Along the way we will find it useful to classify explicitly all the triangle free graphs which satisfy the condition Adj_n and E_2 but not E_n . For this, a description of triangle free graphs satisfying Adj_3 in terms of combinatorial geometries (or hypergraphs) is useful. These geometries can be given in two ways, either as a bipartite graph with one set called “points” and the other set called “blocks,” calling the edge relation between points and blocks “incidence,” or more concretely by taking the geometry to consist of a set of points together with a collection of sets of points called blocks (or hyperedges). The two points of view are not identical: in the bipartite setting, distinct blocks may be incident with the same set of points; in the more concrete setting where the blocks are taken to be sets of points, each such set can occur only once.

Given a graph G and a vertex $v \in \Gamma$ we form the geometry G_v whose points are the neighbors of v in G and whose blocks are the nonneighbors distinct from v , with the edge relation between points and blocks given by the edge relation in G . Edges between pairs of points and pairs of blocks are ignored. This is of interest to us only if the graph can be reconstructed from the geometry. In the triangle free case there are no edges between points, and there are no edges between blocks which intersect. The graphs that interest us will not only be triangle free but maximal triangle free (i.e., adding an edge produces a triangle). In such a case the edge relation between blocks is determined by the geometry: two blocks are joined by an edge if and only if they are disjoint. This point of view corresponds to the usual construction of the Higman-Sims graph, and while none of our geometries will have a geometry with the elegance of that graph's, the same point of view will still be useful. In general, the geometry obtained will depend both on the graph and on the choice of a base vertex v in the graph.

We refer to triangle free graphs satisfying condition E_n as n -e.c., which stands for “ n -existentially complete” (for the category of triangle free graphs). This terminology comes from model theory.

We will see that there are a number of infinite families of geometries all corresponding to 3-e.c. graphs, but that the families known to date are neither varied nor robust. In the M_{22} geometry which is associated with the Higman-Sims graph, every block has 6 points; but in the infinite families of 3-e.c. graphs known to us, there is always at least one block of order 2 (and possibly just one). At the other extreme, in a geometry associated with a 4-e.c. graph, the minimal block size is at least 19. Barring some breakthrough taking us to finite 4-e.c. triangle free graphs and beyond, it would be interesting to make the acquaintance of more robust 3-e.c. graphs. We will give explicit descriptions of some infinite families, in the hope that this may stimulate someone to find better constructions.

Of the six known non-trivial triangle free *strongly regular* graphs, two provide interesting examples of 3-e.c. graphs. The question as to whether there are more such to be found, and possibly an infinite family, seems to be tied up with fundamental problems in that subject. However following a suggestion of Peter Cameron, we can eliminate the possibility that a strongly regular triangle free graph could be 4-e.c.

The following general principle is immediate.

REMARK 11.1. *If M is a structure with the finite model property, and M' is a structure which can be interpreted in M , then M' inherits the finite model property.*

In particular if the generic K_n -free graph has the finite model property for one value of n , then so does the generic K_m -free graph for $m \leq n$, interpreting the latter as the graph induced on the set of vertices adjacent to $(n - m)$ vertices of the former. Thus the finite model property for the generic triangle free graph is the weakest instance of the problem still open in the case of homogeneous graphs.

One would of course like to have general methods for settling the finite model property in homogeneous structures. The finite model property for the Rado graph and for similarly unconstrained homogeneous structures holds by a simple probabilistic argument. Just as a random (countable) graph will be isomorphic to the Rado graph with probability one, a large finite random graph will have one of the appropriate extension properties with asymptotic probability 1 [Fag76]. But probabilistic constructions behave very poorly in the presence of constraints.

For example, if we use counting measure on the set of triangle free graphs of a given size, a random one will be bipartite with high probability [EKR76] (cf. [KPR87] for the K_n -free case), so that the theory of the random finite triangle free graph does not approximate the theory of the generic triangle free graph at all well.

Some time ago Vershik raised the question of a Borel measure invariant under the full infinite symmetric group and concentrating on the generic triangle free graph, a question answered positively in [PV08] (with a classification of the measures in question). But this does not seem to help with the finite model property.

11.2. Equivalence of E_n and $E_{n'}$. We examine the relationships among the natural extension properties. The property E'_2 merits separate consideration.

LEMMA 11.2. *For triangle free graphs G the following properties are jointly equivalent to E'_2 :*

- (i) G is maximal triangle free;
- (ii) G is indecomposable;
- (iii) G contains an independent set of order 3.

Maximality means that the adjunction of any additional edge would create a triangle, which is the same as the mutual adjacency condition: any two independent vertices have a common neighbor.

A graph is *decomposable* if it carries a nontrivial congruence, that is, an equivalence relation such that for any two classes C_1, C_2 either all pairs in $C_1 \times C_2$ are edges, or none are.

Proof of 11.2. The implication from E'_2 to (i – iii) is immediate, so we deal with the converse direction.

Given (ii, iii) it is easy to see that every maximal independent set has more than one vertex, which is clause (a) of E'_2 . So we consider clause (b): for any independent pair of vertices $A = \{u, v\}$ and any subset B of A , we have a vertex adjacent to all vertices of B and no vertices of $A \setminus B$. There are four cases here, all of them relevant.

Suppose conditions (i, ii) hold, and that $|G| \geq 3$. It is easy to see that if condition $E'_2(b)$ fails, then G is a 5-cycle. We give the details.

Fix an independent pair $A = \{u, v\}$ in G . By the indecomposability of G , there is a vertex u' in G adjacent to exactly one of u and v ; suppose u' is adjacent to u , but not to v . The pair $\{u', v\}$ also has a common neighbor v' , which is therefore adjacent to v but not u . Thus for $B = \{u\}$ or $B = \{v\}$, clause (b) of the property E'_2 holds.

Thus under the hypotheses (i, ii), with $|G| \geq 3$, any violation of E'_2 consists of a pair $A = \{u, v\}$ of independent vertices which is a maximal independent set. Then $G \setminus A$ divides into three subsets over A : the set G_u of vertices adjacent to u but not v , the set G_v of vertices adjacent to v but not u , and the residual set $G' = G \setminus (A \cup G_u \cup G_v)$. At this point the edge relation on G is completely determined by conditions (i, ii): apart from the edges involving u or v , the remaining edges must connect G_u to G_v ; conversely, by maximality, the induced graph on $G_u \cup G_v$ is bipartite. Thus by indecomposability $|G_u| = |G_v| = 1$ and $|G'| \leq 1$. The case $G' = \emptyset$ gives a 4-cycle, which is decomposable, so that is excluded, and we are left with the case $|G'| = 1$, which gives a 5-cycle.

If we assume (iii) as well, we have a contradiction. □

Next we check that the properties E'_n and E_n are equivalent.

LEMMA 11.3. *For G triangle free, and n arbitrary, properties E_n and E'_n are equivalent.*

PROOF. We must show that $E'_n \implies E_n$.

Let A be a vertex set of order at most n , and B an independent subset of A . We must show that there is a vertex adjacent to all vertices of B and no vertices of $A \setminus B$. We proceed by induction on n , and on $|A \setminus B|$.

If there is an edge (a, b) with $a \in A \setminus B$ and $b \in B$, then let $A_0 = A \setminus \{a\}$, and apply induction to n to get v adjacent to all vertices of B , and no vertices of $A_0 \setminus B$, and hence to no vertex of $A \setminus B$.

So we may suppose there is no edge connecting $A \setminus B$ and B , but that there is an edge in A . In particular $|A \setminus B| \geq 2$. We claim then: *There is a vertex $u \notin A$ which is adjacent to some vertex $a \in A$ and to no vertex of B .*

Let $A_0 \subseteq A$ be a maximal independent subset of A containing B , and u a vertex not adjacent to any vertex of A_0 . Then u is not in A , and if u is adjacent to some vertex of A we have our claim. So suppose u is adjacent to no vertex of A . Take $a \in A \setminus B$. Then $B \cup \{a, u\}$ is an independent set. Take a vertex v adjacent to u and a , and to no vertex in B . Then $v \notin A$, and v meets the conditions of our claim.

Now applying the claim with u adjacent to $a \in A$ and to no vertex of B , let $B_1 = B \cup \{u\}$. Let $A_1 = (A \setminus \{a\}) \cup \{u\}$. By induction on $|A \setminus B|$ we find v adjacent to all vertices of B_1 and no vertex of $A_1 \setminus B_1$. Then the set of neighbors of v in A is B , as required. \square

11.3. The strength of Adj_3 . We will show that with few exceptions, graphs having the properties E_2 and Adj_n satisfy the full n -e.c. property E_n . The delicate case arises when $n = 3$, and we first dispose of the others.

LEMMA 11.4. *If G is a 3-e.c. triangle free graph with the mutual adjacency property Adj_n , then G is n -e.c.*

PROOF. It suffices to verify the condition E'_n . We proceed by induction on n .

Fix A an independent set of order at most n . We may assume $|A| = n$ or conclude by induction. Our objective is to show that every subset of A occurs as the set of neighbors in A of some vertex of G .

For $v \in G$, let us write A_v for the set of neighbors in A of the vertex v . Then for v' chosen adjacent to v and to all vertices of $A \setminus A_v$, we find $A_{v'} = A \setminus A_v$. So the collection $\{A_v : v \in G \setminus A\}$ is closed under complementation in A .

Fix $X \subseteq A$. We will show that X is A_v for some vertex v . Taking complements if necessary, suppose $|X| \leq n/2$. Take $a \in A \setminus X$ and u a vertex whose neighbors in $A \setminus \{a\}$ are the vertices in X . We may suppose that $A_u = X \cup \{a\}$.

Let Y be the complement $A \setminus (X \cup \{a\})$ and let v be a vertex with $A_v = Y$. If $|Y| > 1$ then taking u' (inductively) whose set of neighbors in $(A \setminus Y) \cup \{v\}$ is $X \cup \{v\}$, we finish.

We conclude that $|X| = n - 2 \leq n/2$, so $n \leq 4$. As G is 3-e.c., $n \geq 4$. Thus $n = 4$ and $|X| = 2$. In particular any singleton occurs as A_v for some v .

Write $A = \{a_1, a_2, b_1, b_2\}$ and let $X = \{a_1, a_2\}$. Let b'_1 be a vertex with $A_{b'_1} = \{b_1\}$. Let b_2 be a vertex whose unique neighbor in $\{a_1, a_2, b'_1, b_2\}$ is b_2 . Let u be a vertex adjacent to a_1, a_2, b'_1, b'_2 . Then $A_u = X$. \square

Now we take up the graphs satisfying E_2 and Adj_3 , but not E_3 , and we begin with a construction.

11.4. The Linear Order Geometries. Given any triangle free graph G with the mutual adjacency property Adj_3 we associate a combinatorial geometry to each vertex v of G by taking as the set of points P the neighbors of v , and as the blocks of the geometry its non-neighbors. A point p lies on a block b if the pair (p, b) is an edge. The geometry obtained may depend on the vertex chosen, but given one such geometry, the graph G may be entirely reconstructed as follows. As the vertex set for G we take $P \cup B \cup \{v_0\}$ where v_0 is an additional vertex. We take as edges all pairs (v_0, p) with p in P , all pairs (p, b) with p on b , and all pairs (b_1, b_2) with b_1, b_2 disjoint when viewed as subsets of P ; and we symmetrize. This agrees with the original graph G : in particular, if b_1, b_2 are not adjacent in G , then by the

property Adj_3 there is a point p lying on both, and hence they are not adjacent in the reconstructed version of G .

The following uses extremely weak assumptions, but then we intend to apply it to an extremely weak geometry.

LEMMA 11.5. *Let (P, B) be a combinatorial geometry on at least 3 points satisfying the following conditions.*

- (1) *No three blocks are pairwise disjoint.*
- (2) *No pair of distinct blocks correspond to the same subset of P .*
- (3) *For every block b and every point p not in b , there is a block containing p and disjoint from b .*
- (4) *For every pair of points there is a block containing one but not the other.*
- (5) *No block is incident with every point of P .*

Then the associated graph G is triangle free and 2-e.c.

PROOF. One checks that G is maximal triangle free, indecomposable, with an independent set of size at least three.

One point that requires checking is that no block is empty; this is part of the verification that G is maximal triangle free. For this, use the assumptions to get two nonempty blocks which are disjoint, and observe that the empty block would extend this to a pairwise disjoint triple. \square

Let L be a linear order. Let B be a set of proper initial segments of L and proper terminal segments of L satisfying the conditions:

- (1) For all $a < b$ in L , there is an initial segment in B containing a and not b , and a terminal segment in B containing b and not a .
- (2) If I is an initial segment and $a \in B \setminus I$ is a lub for I , then the terminal segment $[a, \infty)$ is in B ; and dually.

One way to meet these conditions is to let B consist of the proper segments of the form $(-\infty, a]$ and $[a, \infty)$, and this is the only way to achieve even the first of them if L is finite. Let us call such a geometry a *linear geometry*.

LEMMA 11.6. *Let (P, B) be a linear geometry on at least 3 points. Then the associated graph satisfies the conditions E_2 and Adj_n for all n (if P is finite, this is vacuous for n larger than $|P|$).*

PROOF. The conditions on the geometry have been written to ensure that our criterion for the property E_2 applies.

Now suppose A is any independent subset of the associated graph. If A contains no points, then A consists of some blocks, and possibly the base point, and as the blocks all meet pairwise, it suffices to take a point common to the minimal initial segment in A , and the minimal terminal segment in A . That point is then a vertex adjacent to all vertices of A .

If A contains a point p , then it cannot contain both initial segments and terminal segments, as they would be separated by p . So suppose for example A contains initial segments, and let I be the greatest among them. Let $a \in A$ be the least point. Then $a \notin I$. One of our two conditions on B then applies to give a terminal segment disjoint from I , and containing a .

Of course if A consists exclusively of points, the base point will suffice as a common neighbor. \square

If L is finite of size n , then the resulting graph G has order $3n - 1$ and can be construed as follows. The elements of G are the integers $0, 1, \dots, 3n - 2$; the edge relation is defined by $|i - j| \equiv 1 \pmod{3}$. The value $n = 2$, which is not permitted as we require 3 points, corresponds to the pentagon, and the first legitimate example is a graph of order 8. The maximal independent sets have size n and are the points of the geometry, with respect to a basepoint which is their common neighbor. The geometry obtained is independent of the base point.

Our next point is that the converse holds: a graph with property E_2 and property Adj_n which does not have property E_n must be obtained in this way from a linear geometry, and in particular if it is finite then it isomorphic to the graph just described explicitly. It will suffice to treat the case $n = 3$, since once E_3 is satisfied, Adj_n implies E_n .

11.5. Graphs with E_2 , Adj_3 , and not E_3 .

DEFINITION 11.7. An independent set of vertices I in a graph G will be said to be *shattered* if every subset of I occurs as $\{a \in I : a, v \text{ are adjacent}\}$ for some vertex v .

We will be concerned for the present with graphs which are 2-e.c. but not 3-e.c., and therefore contain independent triples which are not shattered. We want to show that within the independent set of all neighbors of any fixed vertex of G , if one triple is shattered, then all are. We arrive at this gradually by considering various special cases.

LEMMA 11.8. *Let G be a triangle free graph with properties E_2 and Adj_3 , and suppose that a, b, c is a shattered independent triple. If a', b, c is another independent triple with a, a' adjacent, then a', b, c is also shattered.*

PROOF. Since we have property Adj_3 , and the collection of subsets of the set $A = \{a', b, c\}$ which occur as the set of neighbors in A of vertices of G is closed under complementation, it suffices to consider a single vertex $u \in \{a', b, c\}$ and to show that u occurs as the unique neighbor of some vertex among a', b, c .

For $u = a'$ we take a as the witnessing vertex. For $u = b$ we take u' adjacent to a, b and not c . For $u = c$ proceed similarly. \square

LEMMA 11.9. *Let G be a triangle free graph with properties E_2 and Adj_3 , and suppose that the triple (a, b, c) is independent and not shattered. Then there is a unique vertex u in $\{a, b, c\}$ which does not occur as the unique neighbor among a, b, c of a vertex in G .*

PROOF. There must be at least one such vertex, say b . Taking u adjacent to b and not to a , the neighbors of u among a, b, c will be b and c ; and then a common neighbor of u and a will have a as its unique neighbor among a, b, c ; similarly c will occur as the unique neighbor of some vertex among a, b, c . \square

LEMMA 11.10. *Let G be a triangle free graph with properties E_2 and Adj_3 , and suppose that $I = \{a, b, c, d\}$ is an independent quadruple with a, b, c shattered, while some vertex u has a and d as its only neighbors in I . Then (b, c, d) is shattered.*

PROOF. By Lemma 11.8 the triple u, b, c is shattered, and by another application of the lemma, b, c, d is shattered. \square

LEMMA 11.11. *Let G be a triangle free graph with properties E_2 and Adj_3 , and suppose that $I = \{a, b, c, d\}$ is an independent quadruple with a, b, c shattered, while no vertex u has precisely two neighbors in I . Then the triple b, c, d is shattered.*

PROOF. Fix vertices a', b', c' having as their unique neighbors in a, b, c the vertices a, b , or c respectively. By our hypothesis, none of the vertices a', b', c' is adjacent to d .

Now at least two of the vertices a', b', c' are nonadjacent; let u_1, u_2 be two such. Take v adjacent to u_1, u_2 , and d . Then v is adjacent to at most two vertices of a, b, c, d and hence, by our hypothesis, to at most one; that is, v is adjacent to d and not to a, b, c .

Suppose now that $u_1 = a'$. Then $I' = \{a', b, c, d\}$ is an independent quadruple with a', b, c shattered, and v has only a' and d as its neighbors in I' . By the previous lemma, b, c, d is shattered.

So we may suppose that $u_1 = b'$ and $u_2 = c'$. Then b', c' , and v have as their unique neighbors in b, c, d the vertices b, c , and d respectively, and so b, c, d is shattered. \square

LEMMA 11.12. *Let G be a triangle free graph with properties E_2 and Adj_3 , and suppose that $I = \{a, b, c, d\}$ is an independent quadruple with a, b, c shattered. Suppose there are vertices b', c' in G having as their neighbors in I the pairs a, b and a, c respectively. Then b, c, d is shattered.*

PROOF. The triple b', c', d is independent. Let d' be adjacent to b', c', d . Then the vertices b', c', d' show that the triple b, c, d is shattered. \square

LEMMA 11.13. *Let G be a triangle free graph with properties E_2 and Adj_3 , and suppose that $I = \{a, b, c, d\}$ is an independent quadruple with a, b, c shattered. Suppose there are vertices u, u' in G having as their neighbors in I the pairs a, b and c, d , and no other pairs from a, b, c, d occur in this fashion. Then b, c, d is shattered.*

PROOF. Applying Lemma 11.10 to the quadruple (c, a, b, d) , it follows that the triple a, b, d is shattered.

Now if a, c, d is shattered, we argue similarly that b, c, d is shattered by looking at the quadruple (a, c, d, b) .

Take a vertex a' adjacent to a and not adjacent to b, c . By our hypothesis, a' is not adjacent to d either. Now take a vertex v adjacent to a', b, d . Then v is not adjacent to a , so by our hypothesis v is adjacent to c .

We now consider two cases. First, if there is a vertex d' adjacent to d but not to a, b, c , then we apply Lemma 11.8 to the series of independent triples $a, b, d, a, b, d', a, v, d', a, c, d', a, c, d$ to conclude.

Now suppose that there is no such vertex d' . Then there is no vertex whose unique neighbor in b, c, d is d . By Lemma 11.9, there is a vertex c' whose unique neighbor among b, c, d is c . Hence, by our hypothesis, c is the only neighbor of c' among a, b, c, d . If (a', c') is an edge, we conclude by applying Lemma 11.8 to the sequence of independent triples

$$(a, b, d); (a', b, d); (c', b, d); (c, b, d)$$

So we may suppose (a', c') is not an edge. Take a vertex w adjacent to a', b, c' . Then w is adjacent to b but not a or c , and hence not d either. Now apply Lemma 11.8 to the independent triples $(a, b, d), (a, w, d), (a, c', d), (a, c, d)$ to conclude. \square

LEMMA 11.14. *Let G be a triangle free graph with properties E_2 and Adj_3 , and let I be an independent set containing some shattered triple. Then all triples of vertices from I are shattered.*

PROOF. By the foregoing lemmas, if (a, b, c, d) is any independent quadruple containing a shattered triple, then all of its triples are shattered. The general case follows. \square

PROPOSITION 11.15. *Let G be a triangle free graph with properties E_2 and Adj_3 , but not E_3 . Let v be a vertex of G . Then the geometry (P, B) associated to the vertex v in G is a linear geometry, and G is the associated graph.*

PROOF. G is certainly the associated graph, so everything comes down to recognizing the geometry on (P, B) , with P the set of neighbors of v and B the set of non-neighbors.

We first choose v to be a vertex of G having a triple of neighbors which is not shattered, and let (L, B) be the associated geometry. Once we verify that this is a linear geometry, the structure of G is determined, and it follows that the same applies to the geometry at any vertex of G .

We next look for a *linear betweenness relation* on L . This is a ternary relation $\beta(x, y, z)$, irreflexive in the sense that it requires x, y, z to be distinct, which picks out for each triple x, y, z a unique element which is between the other two, i.e., $\beta(x, y, z)$ implies $\beta(z, y, x)$ and not $\beta(y, z, x)$ or $\beta(z, x, y)$. In addition to these basic properties we have the axiom:

For x, y, z, t distinct, $\beta(x, y, z)$ implies $\beta(x, y, t)$ or $\beta(t, y, z)$

Any linear order gives rise to a linear betweenness relation, and the reverse order gives the same betweenness relation; conversely, a betweenness relation determines a unique pair of linear orders which give rise to it (assuming there are at least two points).

We define $\beta(x, y, z)$ on L as follows: $\beta(x, y, z)$ holds if any vertex adjacent to y is adjacent to x or z ; equivalently (taking complements) any vertex adjacent to x and z is adjacent to y . This is symmetric in x and z , and by Lemma 11.9, the relation picks out of each independent triple (x, y, z) which is not shattered, a unique y satisfying $\beta(x, y, z)$ and $\beta(z, y, x)$. Furthermore by our choice of v and Lemma 11.14, none of the triples in L are shattered. So we only have to check the critical axiom: assuming $\beta(x, y, z)$, with t a fourth vertex in L , we claim that $\beta(x, y, t)$ or $\beta(t, y, z)$ holds.

Suppose $\beta(x, y, z)$ holds and $\beta(x, y, t)$ fails. We will show that $\beta(t, y, z)$ holds. As $\beta(x, y, t)$ fails, we have $\beta(x, t, y)$ or $\beta(t, x, y)$.

Suppose $\beta(x, t, y)$ holds. Take a vertex u adjacent to y and not t . Then by $\beta(x, t, y)$, u is not adjacent to x . By $\beta(x, y, z)$, u is adjacent to z . This proves $\beta(t, y, z)$, as claimed.

Now suppose $\beta(t, x, y)$ holds, and take a vertex u adjacent to y , but not adjacent to z . Then by $\beta(x, y, z)$ we have u adjacent to x , and by $\beta(x, t, y)$ we have u adjacent to t . Thus $\beta(t, y, z)$ holds.

Accordingly, β is a linear betweenness relation on L and we may fix a linear ordering giving rise to this relation. By the definition of β , the blocks of B are convex, and are not bounded both above and below. Therefore they are initial and terminal segments of L (reversing the order will of course interchange these

two notions). It remains to check that the blocks are proper, pairwise distinct, and sufficiently dense in L to satisfy our axioms for a linear geometry. This all follows from the assumption that G satisfies E_2 , now that the general shape of the geometry has been established. \square

With this result in hand we can give a reasonably efficient axiomatization of the geometries associated with 3-e.c. graphs.

11.6. Geometries associated with 3-e.c. graphs.

DEFINITION 11.16. An E_3 -geometry is a combinatorial geometry (P, B) satisfying the following axioms.

- I There are no three disjoint blocks.
- II No block is contained in any other.
- III There are at least two points. For any two distinct points, there is a block containing exactly one of them, and a block containing neither.
- IVa If b is a block and p, q are points not in b , then there is a block disjoint from b containing p and q .
- IVb If b, b' are blocks which intersect, and p is a point outside their union, then some block containing p is disjoint from b and b' .
- V If three blocks intersect pairwise, then they either have a point in common, or some block is disjoint from their union.

The E_3 -geometries are just the geometries associated with 3-e.c. triangle free graphs, as we shall show. We avoid the seemingly natural term “3-e.c. geometry”, as the natural interpretation for that term would be a considerably stronger set of conditions in which, notably, in axiom V we would require both a point in common and a disjoint block, that is, we would apply the 3-e.c. condition directly to the geometry, specifying the type of the element of $P \cup B$ realizing the giving condition. The thrust of this is considerably more like 4-e.c. in the corresponding graph; in fact, it is 4-e.c. restricted to quadruples including the base point. No doubt this is an interesting class of geometries in its own right, and more tractable than those associated with 4-e.c. graphs, but still a good deal beyond anything we can construct, or analyze, at present.

LEMMA 11.17. *The geometry associated to any vertex of a 3-e.c. triangle free graph is an E_3 -geometry, and conversely the graph constructed in the usual way from an E_3 -geometry is a 3-e.c. triangle free graph.*

PROOF. One can read off all these axioms directly from the 3-e.c. property (with the triangle free condition accounting for the first of them). The point is to check that conditions (I–V) are strong enough. For that, we use the analysis of the previous subsection. By Axiom II, we exclude the linear geometries of section §11.4, and therefore it suffices to check that the associated graph is triangle free, 2-e.c., and satisfies the adjacency condition Adj_3 , which is more or less what the axioms assert (with IV(a), IV(b), and V corresponding to different instances of Adj_3). \square

We derive some further consequences of the axioms. Observe that Axioms I and III imply that all blocks are nonempty, and that a further application of Axiom III implies that there are at least 3 points.

LEMMA 11.18. *Let (P, B) be an E_3 -geometry. Then the union of two blocks is never P .*

PROOF. Suppose first that b_1, b_2 are two blocks which meet, and let $p \in b_2 \setminus b_1$. There is a block b' containing p and disjoint from b_1 , and as b' cannot be contained in b_2 , it follows that $b_1 \cup b_2 \neq P$.

Now suppose that b_1, b_2 are disjoint and their union is P . We may suppose that $|b_1| \geq 2$. Take two points of b_1 and a block b containing just one of them. As b is not contained in b_1 and $b_1 \cup b_2$ is P , b meets b_2 . By construction $b \cup b_2$ is not P , so there is a block b'' disjoint from $b \cup b_2$. Then b'' is a proper subset of b_1 and we have a contradiction. \square

LEMMA 11.19. *Let (P, B) be an E_3 -geometry. Then any two points belong to some block.*

PROOF. Call two points of P *collinear* if they lie in a common block. We claim first that this is an equivalence relation on P .

Suppose that $p, q \in b_1$, and $q, r \in b_2$. Take a point $a \notin b_1 \cup b_2$, and a block b disjoint from $b_1 \cup b_2$ containing a . As p, r lie outside b , Axiom IVa applies, and p, r are collinear.

By Axiom I, there are at most two equivalence classes for the collinearity relation, and we claim there is only one.

Suppose there are two collinearity classes P_1, P_2 . As $|P| \geq 3$, we may suppose $|P_1| \geq 2$. Take a block b meeting (and hence contained in) P_1 . As there are no inclusions between blocks, it follows from Axiom III that b is a proper subset of P_1 . Therefore by Axiom IVa we have a block meeting P_1 and P_2 , a contradiction. \square

COROLLARY 11.19.1. *Let (P, B) be an E_3 -geometry. Then every block contains at least two points.*

At this point we are through sorting through the basic axioms and we can begin to look more closely at examples of 3-e.c. triangle free graphs and their associated geometries.

We first take up the strongly regular case, then look into the minimal size of an E_4 -geometry.

12. Strongly regular graphs and E_4 -geometries

We will discuss the extension properties of the known strongly regular triangle free graphs and then show that such a graph cannot have property E_4 , following up on an old suggestion of Peter Cameron. We also look at the minimum degree of a graph with property E_4 , or in other words, the minimal size of an E_4 -geometry.

12.1. The Higman-Sims Graph. This is constructed from the M_{22} geometry, defined as follows. Let P_0 be the projective plane over the field of order 4, with 21 points. Adjoin an additional point ∞ to get $P = P_0 \cup \{\infty\}$. The associated blocks will be of two kinds. The first kind are obtained by extending an arbitrary line ℓ of P_0 by the point ∞ : $\ell^* = \ell \cup \{\infty\}$. The second kind are called *hyperovals*. A hyperoval is a set of 6 points in P_0 which meets any line of P_0 in an even number of points. There are 168 hyperovals, and on this set the relation “ $|O_1 \cap O_2|$ is even” is an equivalence relation, with three classes of 56 hyperovals each, permuted among themselves by the automorphism group of the base field. Any one class of 56 hyperovals may be taken, together with the extended lines, as the set of blocks B for the M_{22} geometry, on 22 points. Thus there are 77 blocks, each with 6 points, and 100 vertices in the associated graph, the Higman-Sims graph. Its automorphism

group is vertex transitive and edge transitive, so we get the same geometry from any base point, and any adjacent point can play the role of the “new” point ∞ .

We check that this graph is a 3-e.c. triangle free graph. Given three independent vertices, one may be taken to be the base point v , and the other two will then represent two intersecting blocks of the geometry. We may suppose that they both contain the point ∞ in the associated geometry, and therefore they represent two extended lines, whose intersection has order 2. So we have the condition Adj_3 with multiplicity two, that is there are two points meeting the adjacency conditions in every case. Since the graph visibly satisfies the E_2 condition, and there are no containments between blocks, this completes the verification that it is 3-e.c. On the other hand, it is not 4-e.c. As we know, this comes down to the 4-adjacency property Adj_4 . Taking a triple of points lying on a projective line ℓ in P_0 , and another line meeting ℓ in a different point, a common neighbor of the four vertices involved would be a block containing the given three points and disjoint from the given line; but there is only one block containing three given points, so this is impossible.

The Higman-Sims graph is an example of a *strongly regular* triangle free graph. In general, a graph on n vertices is strongly regular with parameters (n, k, λ, μ) if it is regular of degree k , and any pair of vertices v, v' has λ common neighbors if v, v' are adjacent, and μ common vertices otherwise. In the case of triangle free graphs ($\lambda = 0$), leaving aside the complete bipartite graphs and the pentagon, there are six known examples, which go by the names of the Petersen, Clebsch, Hoffman-Singleton, Gewirtz, M_{22} , and Higman-Sims graphs [Br11, vLW92, CvL91, BvL84]. Two of these graphs are 3-e.c., the Clebsch graph and the Higman-Sims graph.

While there are no other known strongly regular triangle free graphs, there are many “feasible” sets of parameters, that is combinations of parameters which are compatible with all known constraints on such graphs. Following a suggestion of Peter Cameron, we will use that theory to show that there are no 4-e.c. strongly regular triangle free graphs, leaving entirely open the problem whether there are any more, or infinitely many more, 3-e.c. strongly regular triangle free graphs.

We will also take a closer look at the known strongly regular triangle free graphs, notably the Clebsch graph, which serves as the basis for the simplest construction of an infinite family of 3-e.c. graphs, first proposed by Michael Albert.

12.2. Strongly regular graphs and properties E_2, E_3 . Leaving aside the complete bipartite graphs and the pentagon, the known strongly regular triangle free graphs have the following parameters.

- (1) Petersen: $(10, 3, 0, 1)$;
- (2) Clebsch: $(16, 5, 0, 2)$;
- (3) Hoffman-Singleton: $(50, 7, 0, 1)$;
- (4) Gewirtz: $(56, 10, 0, 2)$;
- (5) M_{22} : $(77, 16, 0, 4)$;
- (6) Higman-Sims: $(100, 22, 0, 6)$

The entry “0” here simply says that the graph is triangle free. The Gewirtz graphs and the M_{22} graph can be seen naturally inside the Higman-Sims graph (the Hoffman-Singleton graph, less naturally): the Gewirtz graph is the graph on the hyperovals of the M_{22} geometry, which could be viewed as the set of vertices

in Higman-Sims nonadjacent to two vertices lying on an edge. The M_{22} graph is the graph on the blocks of the M_{22} geometry, and appears as the constituent of Higman-Sims on the non-neighbors of a fixed vertex.

In the Higman-Sims graph, any independent triple of vertices has exactly two common neighbors, as noted previously. In particular if the vertices represent hyperovals with two common points, then their common neighbors are both represented by points, and hence do not lie in the M_{22} graph. Thus neither the Gewirtz graph nor the M_{22} graph can satisfy the condition Adj_3 .

The Hoffman-Singleton and Petersen graphs, with the fourth parameter $\mu = 1$, do not fit into this framework at all: there is no useful geometry induced on the set of neighbors of a fixed vertex, as its blocks would consist of single points. The Petersen graph does play a respectable role as the set of blocks in the geometry associated to the Clebsch graph, and the latter is indeed a 3-e.c. graph. In the case of the Clebsch graph, the geometry is extremely degenerate: it consists of all pairs from a set of order 5. Nonetheless this geometry is an E_3 -geometry.

Thus the Clebsch graph and Higman-Sims graph both are 3-e.c., and the Petersen, Gewirtz, and M_{22} graphs are naturally represented as descriptions of part or all of the associated geometries.

Any strongly regular graph other than the complete bipartite graphs and the pentagon graph will satisfy the condition E_2 , and have no proper inclusion between blocks, so in all other cases the condition E_n will be equivalent to the adjacency condition Adj_n .

However the condition Adj_4 is already incompatible with strong regularity for triangle free graphs, as we now show, following the notation of [Big09], which relies on the eigenvalue theory for strongly regular graphs, expressing everything in terms of the minimal eigenvalue for the adjacency matrix of the graph. So we begin by reviewing that material.

12.3. Eigenvalues and E_4 in the strongly regular case. Let G be strongly regular with parameters (n, k, λ, μ) . Let A be the $n \times n$ adjacency matrix for G , with 0 entries for non-adjacent pairs of vertices, and 1 for adjacent pairs. With J the $n \times n$ matrix consisting entirely of 1's, the condition of strong regularity, with the specified parameters, translates into the matrix condition

$$A^2 + (\mu - \lambda)A - (k - \mu)I = \mu J$$

and as J has the eigenvalues n with multiplicity 1 and 0 with multiplicity $n - 1$, A has three eigenvalues: k with multiplicity 1, and two eigenvalues α, β which are roots of the quadratic equation

$$x^2 + (\mu - \lambda)x - (k - \mu) = 0$$

with multiplicities m_α, m_β satisfying

$$m_\alpha + m_\beta = n - 1; m_\alpha \cdot \alpha + m_\beta \cdot \beta = 0$$

since the trace of A is zero. Specializing to the case $\lambda = 0$ this gives the following formulas in terms of the parameter $s = \sqrt{\Delta}$, Δ being the discriminant $\mu^2 + 4(k - \mu)$, an integer in the nontrivial cases (leaving aside the pentagon and complete bipartite graph):

$$\alpha = \frac{s - \mu}{2}, \beta = \frac{-s - \mu}{2}$$

Following [Big09], we write q for the eigenvalue of minimal absolute value (α , above) and express everything in terms of q and μ as follows.

$$\begin{aligned} k &= (q+1)\mu + q^2 \\ n &= (q^2 + 3q + 2)\mu + (2q^3 + 3q^2 - q) + (q^4 - q^2)/\mu \\ &= (q+2)k + (q^3 + q^2 - q) + (q^4 - q^2)/\mu \\ \mu &\leq q(q+1) \end{aligned}$$

The inequality on μ is not obvious but is derived rapidly from elementary considerations of linear algebra in [Big09].

We note that the extremal values $\mu = q(q+1)$, $k = q^3 + 3q^2 + q$, $n = q^2(q+3)^2$ satisfy all known feasibility constraints and give $k \approx (q+1)^3$, $n \approx (q+1\frac{1}{2})^4$, so that k^4 and n^3 are fairly close, and as we will see in a moment this makes the refutation of condition E_4 a little delicate. Namely, given E_4 , or what amounts to the same thing, Adj_4 , our condition is that the collection of all independent 4-tuples of vertices (u_1, u_2, u_3, u_4) should be covered by the subset consisting of the independent 4-tuples lying in neighborhoods of the vertices, with the former being slightly less than n^4 (at least $n(n - (k+1))(n - 2(k+1))(n - 3(k+1))$) and the latter approximately $n \cdot k^4$, leading to an estimate of roughly the form $k^4 > n^3$. At least for large values of n it is clear that this will not be satisfied at the extreme values and is less likely to hold lower down. But we will work through this more precisely to get the following.

PROPOSITION 12.1. *There is no strongly regular triangle free graph with the property E_4 .*

PROOF. Begin with the estimates

$$\begin{aligned} n &= (q+2)k + (q^3 + q^2 - q) + (q^4 - q^2)/\mu \\ &\geq (q+2)k + (q^3 + q^2 - q) + (q^2 - q) \\ &= (q+2)k + (q^3 + 2q^2 - 2q) \\ &\geq (q+3)k - (\mu + 2q) \\ &\geq (q+3)k - (q^2 + 3q) \end{aligned}$$

Now the number of independent quadruples of vertices in our graph G is at least $n[n - (k+1)][n - 2(k+1) + \mu][n - 3(k+1) + \mu]$ and assuming E_4 , they all occur in neighborhoods of individual vertices, so the number is at most $nk(k-1)(k-2)(k-3)$. So we have

$$[n - (k+1)][n - 2(k+1) + \mu][n - 3(k+1) + \mu] \leq k(k-1)(k-2)(k-3)$$

and we show this is impossible.

We have

$$\begin{aligned} [n - (k+1)] &\geq (q+2)k - (q^2 + 3q + 1) \\ &\geq (q+2)(k - (q+1)) \end{aligned}$$

and

$$\begin{aligned} (n - 2(k+1) + \mu)(n - 3(k+1) + \mu) &\geq [(q+1)k - (2q+2)][qk - (2q+3)] \\ &= q(q+1)(k-2)(k - (2+3/q)) \end{aligned}$$

Furthermore $q(q+1)(q+2) \geq k+q$, so

$$[n-(k+1)][n-2(k+1)+\mu][n-3(k+1)+\mu] \geq (k+q)(k-(q+1))(k-2)[k-(2+3/q)]$$

Suppose $\mu \geq q+1$. Then $k \geq (q+1)^2 + q^2 \geq 2q(q+1)$, so

$$(k+q)(k-(q+1)) = k^2 - k - q(q+1) \geq k(k-1\frac{1}{2})$$

If we take $q \geq 6$ as well then this yields

$$\begin{aligned} [n-(k+1)][n-2(k+1)+\mu][n-3(k+1)+\mu] &\geq k(k-1\frac{1}{2})(k-2)(k-2\frac{1}{2}) \\ &> k(k-1)(k-2)(k-3) \end{aligned}$$

ruling out this case.

For $q < 6$ we may consult the tables in [Big09]. The two cases in which the necessary inequality holds are one for $q = 2$, namely the Higman-Sims graph, already ruled out, and one for $q = 3$, with parameters $n = 324$, $k = 57$, $\mu = 12$, where oddly enough the two sides are exactly equal. One way to eliminate this is to show that the block size is too small: later we will give a lower bound of 19 for the minimal block size in a geometry associated with a 4-e.c. graph.

There remains the marginal case $\mu \leq q$. In this case as $q \leq 2q^2 + q$ we have $n \geq (q+3)k + q^3 - 3q \geq (q+3)k + 3$ for $q \geq 3$ and thus we can use the crude estimate

$$(n-(k+1))(n-2(k+1))(n-3(k+1)) \geq [(q+2)k][(q+1)k][qk] \geq k^4$$

to reach a contradiction. \square

Since we have quoted a lower bound for the block sizes in geometries associated with 4-e.c. triangle free graphs, we will give that next.

12.4. E_4 -geometries: Block size. We will refer to a geometry associated with a 4-e.c. triangle free graph as an E_4 -geometry. We will not try to write out the axioms explicitly. These would consist of conditions encoding the 2-e.c. property as was done in the case of E_3 -geometries, the condition that no block is contained in another, to eliminate the degenerate case of a linear geometry, and finally the main axioms which correspond to the adjacency condition Adj_4 which takes on various forms in the geometrical context depending on how the various vertices are interpreted in the geometry. We may use any instance of the 4-e.c. condition, and are not confined to the special cases corresponding directly to our reduced set of axioms.

We omit the elementary proofs of the next three lemmas.

LEMMA 12.2. *Let (P, B) be an E_4 -geometry, let b, b_1 be intersecting blocks, and let b_2 be any other block. Then*

$$|(b \cap b_1) \setminus b_2| \geq 2$$

LEMMA 12.3. *Let (P, B) be an E_4 -geometry, let b, b_1 be intersecting blocks, Then $|b \cap b_1| \geq 5$.*

LEMMA 12.4. *Let (P, B) be an E_4 -geometry, and let b, b_1, b_2 be blocks with a point in common. Let b_3 be a block meeting b but disjoint from b_1 and b_2 . Then $|b \setminus (b_1 \cup b_2 \cup b_3)| \geq 5$.*

PROPOSITION 12.5. *Let (P, B) be an E_4 -geometry. Then any block contains at least 19 points.*

PROOF. Let b be a block. We claim first that there are b_1, b_2 with $b \cap b_1 \cap b_2 \neq \emptyset$, and with $|(b \cap b_2) \setminus b_1| \geq 4$.

Begin with $b \cap b_1 \cap b_2$ nonempty and with the three blocks distinct, and suppose $|(b \cap b_2) \setminus b_1| \leq 3$. Take a block b_3 so that:

$$|b_3 \cap [(b \cap b_2) \setminus b_1]| = 1; b_3 \cap b_1 = \emptyset$$

Then $|b \cap b_2 \cap b_3| = 1$ and $(b \cap b_3) \setminus b_2 \geq 4$.

So fix blocks b_1, b_2 with $b \cap b_1 \cap b_2 \neq \emptyset$, and with $|(b \cap b_2) \setminus b_1| \geq 4$. Then $|b \cap (b_1 \cup b_2)| \geq 9$. Now take b_3 disjoint from b_1, b_2 and meeting b . Then we have $|b \cap (b_1 \cup b_2 \cup b_3)| \geq 14$. And then by the previous lemma $|b| \geq 19$. \square

It would be good to have a more sophisticated lower bound here. We can convert this bound into a crude but decent lower bound for the number of points in such a geometry.

LEMMA 12.6. *Let (P, B) be an E_4 -geometry. Then there are intersecting blocks b_1, b_2 with $|b_1 \cup b_2| \geq 33$.*

PROOF. Let $m = \min(|b_1 \cap b_2| : b_1 \cap b_2 \neq \emptyset)$. If $m = 5$ we are done.

Suppose $m \geq 6$. Take any two intersecting blocks b_1, b_2 . Take a block b meeting b_1 and disjoint from b_2 . Take a point p in $b_1 \setminus (b_2 \cup b)$, and a block b' disjoint from b_2 containing p .

Then b_1, b, b' meet pairwise and hence by the 4-e.c. condition have a point in common. Furthermore b_1 meets b_2 and b, b' are disjoint from b_2 . So by Lemma 12.4, $|b_1 \setminus (b \cup b' \cup b_2)| \geq 5$. Furthermore $|(b_1 \cap b') \setminus b| \geq 2$. So $|b_1 \setminus (b \cup b_2)| \geq 7$ and $|b_1 \setminus b_2| \geq m + 7 \geq 13$.

Also, in the proof of the previous lemma, the lower bound obtained on the block size is actually $m + 4 + m + 5$ which with $m \geq 6$ would give a block size of at least 21 and a lower bound for $|b_1 \cup b_2|$ of at least 34 in this case. \square

LEMMA 12.7. *Let (P, B) be an E_4 -geometry, $n = |P|$. Then $n \geq 66$.*

PROOF. Take intersecting blocks b_1, b_2 with $|b_1 \cup b_2| \geq 33$. Let

$$m = \min(|b_3 \cap b_4| : b_3 \cap b_4 \neq \emptyset, (b_3 \cup b_4) \cap (b_1 \cup b_2) = \emptyset)$$

If $m = 5$ the result is immediate, so take $m = 6$ and argue as in the previous lemma. \square

Another way of stating all of this is as follows.

COROLLARY 12.7.1. *Let G be a 4-e.c. graph. Then every vertex has degree at least 66, every pair of independent vertices has at least 19 common neighbors, and every triple of independent vertices has at least 5 common neighbors.*

In the Higman-Sims graph the corresponding numbers are 22, 6, and 2.

It might be of interest to explore the known feasible parameter sets for strongly regular triangle free graphs to see which seem compatible with the E_3 -condition. At the extreme value $\mu = q(q + 1)$, the ratio of $k(k - 1)(k - 2)$ to

$$[n - (k + 1)][n - 2(k + 1) + \mu]$$

is q , which in the two known cases of the Clebsch and Higman-Sims graphs actually corresponds to the condition Adj_3 with multiplicity q . There are many other cases where the necessary inequality is satisfied with smaller values of μ . In fact, the majority of the cases listed in the appendix to [Big09] meet this condition. The only case consistent with this inequality in which the Adj_3 condition is known to fail is the case of the M_{22} graph.

In a similar vein, dropping the E_4 condition, we add some comments on the relationship between the multiplicity with which Adj_3 is satisfied, and the multiplicity with which Adj_2 is satisfied.

Let G be a 3-e.c. triangle free graph. Define $\mu_n(G)$ as the minimum over all independent sets $I \subseteq G$ of order n of the cardinality of the set of common neighbors of I . Thus for example in the Higman-Sims graph, $\mu_2(G) = 6$ and $\mu_3(G) = 2$.

LEMMA 12.8. *Let G be a 3-e.c. triangle free graph. If $\mu_3(G) \geq 2$ then $\mu_2(G) \geq 5$.*

PROOF. Fix two vertices $u_1, u_2 \in G$, and v adjacent to both. With v as base point work in the associated geometry. We look for 4 blocks containing the points u_1, u_2 .

Let b_1 be a block containing u_1, u_2 , and $u_3 \in P \setminus b_1$. Let b_2 be a block containing u_1, u_2, u_3 ($\mu_3(G) = 2$), and $u_4 \in P \setminus b_1 \cup b_2$. Let b_3 be a block containing u_1, u_2, u_4 . It suffices to show that $b_1 \cup b_2 \cup b_3 \neq P$.

Let b be a block containing u_4 and disjoint from b_1, b_2 . Then $b \setminus b_3 \subseteq P \setminus (b_1 \cup b_2 \cup b_3)$. \square

So we have what appears to be a sharply descending series of successive weakenings of the E_4 condition, with no known examples of even the weakest condition other than subgraphs of the Higman-Sims graph.

- (1) G is 4-e.c.
- (2) $\mu_3(G) \geq 5$
- (3) $\mu_3(G) \geq 2$
- (4) $\mu_2(G) \geq 5$
- (5) $\mu_2(G) \geq 3$
- (6) In the geometry associated to some base point, every block contains at least 3 points.

What we can do, as mentioned, is produce an infinite family of E_3 -geometries with a *unique* block of order 2, but even in this case we get no bound on the number of blocks of order 2 for *other* geometries associated with the same graph, at different basepoints.

We return now to the case of E_3 -geometries. The first examples of an infinite family of 3-e.c. triangle free graphs was given by Michael Albert. An examination of the corresponding geometries leads naturally to the consideration of a more general class of geometries with rather special properties. From one point of view these examples are degenerate; on the other hand, examples can be constructed naturally from projective geometries.

13. Some E_3 -geometries

13.1. Albert geometries. We first present Michael Albert's original construction. Observe that the Clebsch graph can be represented as a collection of 4

copies of a 4-cycle, related systematically to one another: a vertex v in one copy will be connected only to one vertex in any other copy, namely the vertex corresponding to the one diagonally opposed to v . Evidently the same recipe can be extended to any number of copies of a 4-cycle, and any triple of vertices in one of these extended graphs embeds into a copy of the Clebsch graph; so the “stretched” Clebsch graph inherits the 3-e.c. property from the Clebsch graph.

In terms of the associated combinatorial geometry, the Clebsch graph corresponds to the geometry on 5 points in which every pair is a block. The stretched Clebsch graphs correspond to a geometry on n points, $n \geq 5$, in which the blocks are of two sorts:

- (i) all the pairs containing either of two fixed points;
- (ii) all the sets of points of order $n - 3$ not containing either of those two points.

DEFINITION 13.1.

1. A point p in a combinatorial geometry (P, B) will be said to be *isolated* if every pair of points containing p is a block.

2. An E_3 -geometry will be called an *Albert geometry* if it has at least one isolated point.

LEMMA 13.2. *For $n \geq 5$, there is a unique Albert geometry on n points having two isolated points.*

PROOF. If p, q are two isolated points, and b is a block not containing p or q , then $|b| \geq n - 3$ as otherwise there will be three pairwise disjoint blocks.

Let $a \in P$, $a \neq p, q$. There is a block b disjoint from the blocks $\{a, p\}$ and $\{a, q\}$, and as $|b| \geq n - 3$, $b = P \setminus \{a, p, q\}$. So the identification of the geometry is complete. \square

We will look at some examples of Albert geometries with a unique isolated point. We do not expect that one can classify these without some further restrictions. In general, the geometry in the associated graph will depend on the base point selected, so it is noteworthy that if one of these geometries is an Albert geometry, then they all are. We note that even the number of points in the geometry may depend on the base point, in other words the corresponding graphs are not regular in general. For the specific case of the geometry with two isolated points just described, the corresponding graph is vertex transitive (as is clear from Albert’s original description of it), so the same geometry is obtained from any base point.

If we remove an isolated point from an Albert geometry and look at the geometry induced on the remaining points, we get a reasonable class of geometries. This point of view is useful for the construction of examples.

DEFINITION 13.3. Let (P, B) be an Albert geometry, a an isolated point. The *derived geometry* (P_0, B_0) with respect to a has point set $P_0 = P \setminus \{a\}$, and blocks $B_0 = \{b \in B : a \notin b\}$.

The reconstruction of (P, B) from (P_0, B_0) is immediate. We can phrase the E_3 -conditions directly in terms of the derived geometry as follows.

I-D There are at least two points.

II-D For any three points p, q, r there is a block containing p, q , and not r .

III-D There is no inclusion between distinct blocks.

IV-D For any block, its complement is a block.

V-D If b_1, b_2, b_3 are blocks intersecting pairwise, but with no common point, then $b_1 \cup b_2 \cup b_3 \neq P_0$.

Any geometry satisfying these axioms will be called a *derived geometry*. This terminology is justified by the following lemma.

LEMMA 13.4. *If (P, B) is an Albert geometry with a an isolated point then the derived geometry with respect to a satisfies Axioms I-D through V-D. Conversely, if (P_0, B_0) is a derived geometry, then the combinatorial geometry on $P = P_0 \cup \{a\}$ whose blocks are the blocks of B_0 together with the pairs containing a is an Albert geometry.*

We omit the verification. Using this result, we can give examples in a very convenient form.

EXAMPLES 13.1.

- (1) Let (P_0, H) be a projective geometry with H the set of hyperplanes. Let (P_0, B_0) have as its blocks the elements of H and their complements. This is a derived geometry.
- (2) Let (P_0, L) be a projective plane and let L' be a set of lines satisfying one of the following conditions:
 - (a) Every point lies on at least 3 lines of L' ;
 - (b) L' contains all the lines of L not passing through some fixed point p , and two of the lines passing through p .

Let L^* consist of the lines in L' and their complements. Then (P_0, L^*) is a derived geometry.

- (3) Let (P_0, L) be a projective plane and $\ell \in L$ a fixed line. Let L_ℓ consist of the line ℓ together with the sets $\ell \cup \ell_1 \setminus \ell \cap \ell_1$ as ℓ_1 varies over the remaining lines. Taking these sets and their complements as blocks, we get an Albert geometry, to which we return below.
- (4) Let (P_0, H) be as in (1) and let ∞ be an additional point. Extend (P_0, H) to (P_1, H_1) with $P_1 = P_0 \cup \{\infty\}$, $H_1 = \{h \cup \{\infty\} : h \in H\}$, and let H_1^* consist of the elements of H_1 and their complements. Then (P_1, H_1^*) is a derived geometry.
- (5) Let (P_0, B_0) be a derived geometry, and $a \in P_0$. Let

$$P_a = \{b \in B_0 : a \in b\} \cup \{\infty\}$$

with ∞ a new point. Let $B_a = P_0 \setminus \{a\}$. For $p \in B_a$ let $b_p = \{b \in B_0 : a \in b\} \cup \{\infty\}$. Let B_a^* consist of $\{b_p : p \in B_a\}$ and the complements $\{P_a \setminus b_p : p \in B_a\}$. Then (P_a, B_a^*) is a derived geometry; indeed, it is just the derived geometry obtained by passing from (P_0, B_0) to the corresponding graph, and then using the point p as a base point in place of the original base point. If we repeat this construction to form the geometry $(P_{a\infty}, B_{a\infty})$, we recover (P_0, B_0) .

If (P_0, B_0) is a derived geometry with $|P_0| = n$ and $|B_0| = 2n'$, we will say that (n, n') are the *parameters* of the derived geometry. In Example 13.1 (5) above, if (P_0, B_0) has parameters (n, n') , then (P_a, B_a) has parameters $(n' + 1, n - 1)$. For example: the derived geometry associated with a projective geometry has type (n, n) and the new geometry (P_a, B_a) thus has parameters $(n + 1, n - 1)$.

As we are ultimately interested in graphs, we will want to consider how the geometry varies as we change the basepoint. We may suppose that any change of basepoint takes place in a series of steps, replacing a given basepoint by one adjacent to it, thus iterating the construction in Example 13.1 (5).

PROPOSITION 13.5. *Let (P, B) be an Albert geometry, G the associated graph with base point v and vertex set $\{v\} \cup P \cup B$, and $u \in G$ any vertex. Let (P_u, B_u) be the geometry associated to the graph G with respect to the base point u . Then (P_u, B_u) is again an Albert geometry.*

PROOF. Since G is connected, it suffices to prove the claim when u is adjacent to the base point v of G .

The case in which u is an isolated point of (P, B) must be handled separately. In this case, there is an involution $i \in \text{Aut}(G)$ defined by

$$v \leftrightarrow u; a \leftrightarrow \{a, u\} \ (a \in P \setminus \{u\}); b \leftrightarrow P \setminus b \text{ on blocks}$$

Thus in this case (P_u, B_u) is isomorphic via this automorphism to the original geometry (P, B) .

Suppose now that $p_0 \in P$ is an isolated point, and $u \neq p_0$. Let $p_1 = \{p_0, u\} \in P_u$. We claim that p_1 is an isolated point of P_u .

The vertex $p_0 \in B_u$ is incident with $\{v, p_1\}$ in P_u . For any other $b \in P_u \setminus \{v, p_1\}$, we have $b \in B_0$, $u \in b$, and $b' = P_0 \setminus b$ is in B_u , with b' incident with p_1 and b . Thus p_1 is an isolated point. \square

LEMMA 13.6. *Let (P, B) be an Albert geometry, $p_0 \in P$ an isolated point, and (P_0, B_0) the corresponding derived geometry with parameters (n, n') where $n = |P_0|$ and $n' = |B_0|/2$. Let G be the graph associated with (P, B) , with base point v . Then for any vertex $u \in B_0 \cup \{p_0, v\}$, the associated geometry with base point u has the same parameters (n, n') , while for $u \in P_0 \cup (B \setminus B_0)$, the associated geometry has parameters $(n' + 1, n - 1)$. Thus at most two vertex degrees occur in the associated graph.*

PROOF. Let g be the order of the graph G . We have $g = 1 + (n+1) + (2n' + n) = 2n + 2n' + 2$. For $p \in P_0$, the geometry (P, B_p) has parameters (n_p, n'_p) where $n_p = |P_p| - 1$ with P_p the set of neighbors of p in G , namely $v, \{p, p_0\}$, and the blocks in B_0 which contain p : so $n_p = n' + 1$, and therefore $n'_p = n - 1$.

Making use of the involution $v \leftrightarrow p_0$ which interchanges $B \setminus B_0$ and P_0 , the same parameters are associated with $u \in B \setminus B_0$.

On the other hand, for a block $b \in B_0$, the set P_b of neighbors of b consists of the points p belonging to b , the pairs $\{p_0, p\}$ with $p \in P_0$ not belonging to b , and the complement b' of b in P_0 , leading to $n_b = |P_b| - 1 = |P_0| = n$ and thus $n'_b = n'$. \square

This raises the question as to what sorts of geometries are associated on the one hand with regular graphs, and on the other hand with graphs having just two vertex degrees. These conditions do not seem to be very restrictive, and it may be of interest to impose similar conditions generalizing strong regularity, perhaps allowing some further use of algebraic methods. As an example, if we begin with the Albert geometry based on a projective geometry with a single hyperplane removed, we get a regular graph.

EXAMPLE 13.2. Let G be the graph associated to the Albert geometry whose derived geometry comes from the projective plane. Let v be the base point, $p_0 \in P$ the isolated point, and $\ell \in B$ a line. The associated geometry has points P_ℓ consisting of:

$$\begin{cases} \text{points } p \in P_0 \text{ lying on } \ell \\ \text{co-points } \{p, p_0\} \text{ with } p \notin \ell \end{cases}$$

Then for $\ell_1 \neq \ell$ a line, the points of P_ℓ incident with ℓ as a block in B_ℓ are

$$\{p \in \ell \setminus \ell_1\} \cup \{\hat{p} : p \in \ell_1 \setminus \ell\}$$

In other words, this corresponds to the union of the fixed line ℓ with ℓ_1 , with their common point removed. The block associated with the base point v is $\{p : p \in \ell\}$.

In the specific case of the projective plane of order 2, the resulting geometry again comes from the projective plane of order 2. Otherwise, the geometry is a different one, and the automorphism group of the graph leaves the pair $\{v, p_0\}$ invariant, and is $\text{Aut}(P_0, B_0) \times \mathbb{Z}_2$. But for $q = 2$ the group is transitive on $B_0 \cup \{p_0, v\}$.

13.2. Another series of E_3 -geometries. Moving away from Albert geometries, what we would like to see next is an infinite family of 3-e.c. graphs G with $\mu_2(G) \rightarrow \infty$, ideally even $\mu_3(G) \rightarrow \infty$. But we are far from this. Leaving aside the M_{22} geometry with its remarkably good properties ($\mu_2(G) = 6$, $\mu_3(G) = 2$), in *infinite families* the best we have done to date is to reduce the number of blocks of order 2 to a single one, while all other blocks can be made arbitrarily large. But in this construction we deal with a single geometry, rather than the set of geometries associated with a given graph. So there is much to be improved on even at this weak level.

EXAMPLE 13.3. With $m_1, m_2, m_3 \geq 2$, let the geometry $A(m_1, m_2, m_3)$ be defined as follows. Our pointset is the union of three disjoint sets P_1, P_2, P_3 with $|P_i| = m_i$, together with two distinguished points p', p^- , and the following blocks.

- (1) $b_0 = \{p', p^-\}$ (size 2);
- (2) For $a \in P_i$: $a' = P_i \setminus \{a\} \cup \{p'\}$ (size m_i);
- (3) For $a \in P_i$: $a^+ = \{a\} \cup P_{i+1}$ (addition modulo 3), of size $m_{i+1} + 1$;
- (4) For $a \in P_i$: $a^- = \{a\} \cup P_{i-1} \cup \{p^-\}$ (subtraction modulo 3), of size $m_{i-1} + 2$.

If $m_1, m_2, m_3 \geq 2$ then $A(m_1, m_2, m_3)$ is an E_3 -geometry, and if $m_1, m_2, m_3 \geq 3$ then it has a unique block of order 2, with the other blocks of order at least $\min(m_1, m_2, m_3)$. Let $G(m_1, m_2, m_3)$ be the associated graph.

LEMMA 13.7. *The dihedral group of order 8 acts on $G = G(m_1, m_2, m_3)$ as a group of automorphisms, extending the natural action on the 4-cycle (v, p', b_0, p^-) , with v the base point, $b_0 = \{p', p^-\}$. In this action the classes $P_0 = P_1 \cup P_2 \cup P_3$, $B' = \{a' : a \in P_0\}$, $B^+ = \{a^+ : a \in P_0\}$, and $B^- = \{a^- : a \in P_0\}$ are permuted.*

PROOF. We define two involutions in $\text{Aut}(G)$ by:

$$\begin{array}{l} \iota' : \quad v \leftrightarrow p' \quad b_0 \leftrightarrow p^- \quad a \leftrightarrow a' \quad a^+ \leftrightarrow a^- \\ \iota^- : v \leftrightarrow p^- \quad b_0 \leftrightarrow p' \quad a \leftrightarrow a^- \quad a^+ \leftrightarrow a' \end{array}$$

□

LEMMA 13.8. *The graph $G(m_1, m_2, m_3)$ is regular.*

PROOF. It suffices to check the degree of a vertex $a \in P_i$. The neighbors of a are v, a^+, a^- and a'_1 for $a_1 \in A_i \setminus \{a\}$, b^- for $b \in A_{i+1}$, c^+ for $c \in A_{i-1}$, for a total of $3 + (m_i - 1) + m_{i+1} + m_{i-1}$ which is the degree of the base point. \square

The geometries associated with $G(m_1, m_2, m_3)$ (i.e., those giving an isomorphic graph) are not well behaved.

LEMMA 13.9. *Let $G = G(m_1, m_2, m_3)$ corresponding to $A(m_1, m_2, m_3)$ with point set $P_1 \cup P_2 \cup P_3 \cup \{p', p^-\}$, and take $a \in P_i$. Then the geometry (P_a, B_a) associated with the base point a has at least $m_i + 1$ blocks of order 2, and exactly $2(m_i - 1)$ blocks of order $m_1 + m_2 + m_3 - 1 = n - 2$ with $n = |P|$ the degree of G . On the set $P_i^* = \{a'_1 : a_1 \neq a, a_i \in P_i\} \cup \{a^+, a^-\}$ of order $m_i + 1$, the induced geometry is the Albert geometry with two isolated points a^+, a^- , in the sense that all pairs $\{a^+, a'_1\}$ and $\{a^-, a'_1\}$ occur as blocks of the associated geometry (P_a, B_a) , while all subsets of order $n - 3$ contained in P_a which are disjoint from $\{a^+, a^-\}$ and which contain $P_a \setminus P_i^*$ also occur as blocks.*

Note that in the “restricted” geometry on P_i^* we are taking as blocks, those which lie within P_i^* , and those which contain its complement.

PROOF.

$$P_a = \{v, a^+, a^-, b'(b \in P_i, b \neq a), c^-(c \in P_{i+1}), d^+(d \in P_{i-1})\}$$

The block associated with a' is $\{a^+, a^-\}$. The block associated with b^+ for $b \in A_i$, $b \neq a$ is $\{a^-, b'\}$. The block associated with b^- for $b \in P_i$ is $\{a^+, b'\}$.

Finally, the block associated with $b \in P_i$ ($b \neq a$) is

$$\{v\} \cup P_{i+1}^- \cup A_{i-1}^+ \cup \{b'_1 : b_1 \in P_i, b_1 \neq a, b\}$$

\square

We observe that the smallest geometry for which we are able to get a single block of order 2 has order 11, namely $A(3, 3, 3)$, and this is sharp. We will give some additional information concerning small geometries.

13.3. E_3 -Geometries of order at most 7.

LEMMA 13.10. *Let (P, B) be an E_3 -geometry, and $n = |P|$.*

- (1) *The maximal block size is at most $n - 3$.*
- (2) *If there is a block of order $n - 3$ then all pairs lying in its complement are blocks.*
- (3) *$n \geq 5$.*
- (4) *If p is a point which is contained in at least $n - 3$ blocks of order 2 then p is isolated.*
- (5) *If $n = 5$ or 6 then (P, B) is the Albert geometry with two isolated points.*

PROOF. The first three points are immediate. For the fourth, let q, q' be the two points not known to occur together with p as a block of order 2. Take a block containing p, q and not q' ; it must be $\{p, q\}$. Similarly $\{p, q'\}$ is block.

For the last point, the case $n = 5$ is immediate, so take $n = 6$. Then there is a pair of points p, q which do not constitute a block, and therefore they are contained in two blocks of order 3, say $\{p, q, r\}$ and $\{p, q, s\}$. Let $t \notin \{p, q, r, s\}$. Then the

pairs containing t and neither of p, q are blocks, by our second point. So t is an isolated point. So the geometry has two isolated points, and we are done. \square

We will also carry through the analysis for the case $n = 7$, finding in this case that the geometry is necessarily an Albert geometry, and that there are only two possibilities: the Albert geometry with two isolated points, and one other.

EXAMPLE 13.4. We construct a derived geometry $B'(n_1, n_2)$ as follows. Let A_1, A_2 be sets of orders n_1, n_2 respectively, and c an additional point. Take as blocks:

$$a_1 A_2 (a_1 \in A_1); a_2 A_1 (a_2 \in A_2); c A'_i (A'_i \subseteq A_i, |A_i \setminus A'_i| = 1)$$

Let $B(n_1, n_2)$ be the corresponding Albert geometry, with $n_1 + n_2 + 2$ points.

If $n_1 = n_2 = 2$ then this is the Albert geometry with two isolated points on 6 vertices. Otherwise, it is an Albert geometry with one isolated vertex. In particular we have the case $n_1 = 2, n_2 = 3$ of order 7.

LEMMA 13.11. *Let (P_0, B_0) be the derived geometry associated with an Albert geometry, with $n_0 = |P_0| > 4$. Call the blocks of order 2 in B_0 edges, and view P_0 as a graph with respect to these edges. Then all edges in P_0 have a common vertex.*

PROOF. As $n_0 > 4$ there can be no disjoint edges. So we need only eliminate the possibility that there is a triangle.

Suppose p, q, r form a triangle: any pair is a block. Take further points s, t , and a block b containing r, s but not t . Then b is disjoint from p, q and hence the complement of b is p, q . But then t is in b , a contradiction. \square

LEMMA 13.12. *An Albert geometry on 7 points with one isolated point must be isomorphic with $B(2, 3)$.*

PROOF. We work in the derived geometry (P_0, B_0) on 6 points, and consider the graph on P_0 whose edges are the pairs occurring as blocks in B_0 . If there are three or more edges then their common vertex is a second isolated point, a contradiction.

If there are exactly two edges we identify the geometry $B(2, 3)$ as follows. Let the vertex common to the edges be called c , and let the other vertices on the edges be $A = \{a_1, a_2\}$, while the remaining vertices are $B = \{b_1, b_2, b_3\}$. As there are just the two edges, the other blocks containing c are triples of the form $\{c, b, b'\}$ with $b, b' \in B$, and as we may exclude any element of B , all such triples occur. Thus we know all the blocks containing c and taking complements, we have all the blocks.

Suppose there is at most one edge, and take four points $A = \{a_1, a_2, a_3, a_4\}$ containing no edge, and let c_1, c_2 be the other points. Then every block containing c_1 and meeting A is a triple, thus the blocks containing c_1 meet A in a certain set of pairs E_1 , and for any two points $a, a' \in A$ there is a pair in E_1 containing a and not a' . If E_1 contains no disjoint pairs, it follows that the edge set E_1 forms a triangle in A ; and if we define E_2 similarly, and E_2 contains no disjoint pairs, then E_2 forms a triangle in A . However $E_1 \cup E_2$ covers all pairs in A . So we may suppose that E_1 contains two disjoint pairs, say $\{a_1, a_2\}$ and $\{a_3, a_4\}$. We may suppose then that E_2 contains the pair $\{a_1, a_3\}$, and hence B_0 contains the blocks $\{c_1, a_1, a_2\}, \{c_1, a_3, a_4\}, \{c_2, a_1, a_3\}$ which meet pairwise but have no point in common. But as the union of these blocks is P_0 , we contradict our axioms. \square

And lastly we claim that up to this point no non-Albert geometry occurs.

LEMMA 13.13. *Let (P, B) be an E_3 -geometry on $n \leq 7$ points. Then (P, B) is an Albert geometry.*

PROOF. We have dealt with the cases $n < 7$ and we suppose $n = 7$.

We call a block of order 2 an *edge* in P .

Suppose first that there is some block A of order 4. We claim that any $a \in A$ lies on an edge.

Let $A = \{a, a_1, a_2, a_3\}$ and take blocks b_1, b_2 with $a_i, a_3 \in b_i$, $a \notin b_i$ for $i = 1, 2$. Take b disjoint from $b_1 \cup b_2$ with $a \in b$. Since any pair disjoint from A is an edge, b must also be an edge.

Now as $|A| = 4$ we can find two points a', a'' in A for which there are edges $\{p, a'\}, \{p, a''\}$ with a common neighbor $p \in P \setminus A$. Then p lies on 4 edges and is therefore an isolated point.

From now on suppose that *there is no block of order 4*, and we will arrive at a contradiction.

We show first that every point lies on an edge. Suppose the point p lies on no edge, so that every block containing p has order 3. Take two blocks b_1, b_2 whose intersection is $\{p\}$. Then the complement of $b_1 \cup b_2$ is an edge e .

Take $q \in b_1, r \in b_2$ with $q, r \neq p$, and with q, r not an edge, using the fact that there are no three disjoint edges. Take a block b_3 containing q, r , and not p . Then b_1, b_2, b_3 meet pairwise but have no common point, so they are disjoint from a block, which must be e . Thus $b_3 \subseteq b_1 \cup b_2$. There is a point $s \notin b_3 \cup \{p\} \cup e$. Form a block b containing q, s and not p , and take a block b' disjoint from $b \cup b_3$ and containing p . Then $p \in b' \subseteq \{p\} \cup e$ and thus b' is an edge containing p .

Now consider the graph on P formed by the edges. Every vertex lies on an edge, there are no three disjoint edges, and furthermore no vertex has degree greater than 3, as it would then be an isolated point of the geometry.

By the first two conditions, some vertex p must have degree at least 3, and hence exactly 3. Then it follows by inspection that there is some vertex q not adjacent to p such that every point of P other than p, q is adjacent to one of the two points p, q . Consider a block b containing p, q . Since b cannot contain any edge at p or q , b is $\{p, q\}$, so these points are adjacent and we have a contradiction. \square

13.4. A small non-Albert geometry. We record some further information about small E_3 -geometries. The smallest non-Albert geometry lives on a set with 8 points, and is unique up to isomorphism. This geometry has 7 blocks of order 2.

The smallest geometry in which one has a unique block of order 2 is the geometry $A(3, 3, 3)$ on 11 points. To get an E_3 geometry with no block of order 2 one may take the geometry of lines and hyperovals in the projective plane over a field of order 4, with 21 points. On the other hand, we have checked that such a geometry must have at least 13 points. We would like to know the minimum size of such a geometry, and what the geometry is. We observe at this point some distinct and possibly very substantial gap between the degenerate cases we have discussed and the next level.

By brute force search, all of the E_3 -geometries of order 8 may be identified. There are 11 such geometries, corresponding to 7 graphs. Four of these graphs correspond to two geometries of order 8, two correspond to one geometry of order 8 apiece, while the last graph corresponds to one geometry of order 8 and one of order 9. These geometries are Albert geometries except for one pair of geometries corresponding to a single graph. We list the geometries as follows, including the

block counts (the number of blocks of each size, from the minimum size 2 up to the maximum size). We will refer to one geometry as a “variant” of another if it defines an isomorphic graph.

E_3 -geometries of order 8:

- (1) The Albert geometry with two isolated points. Block count $(13, 0, 0, 6)$.
- (2) Albert geometries with unique isolated points:
 - (a) The geometry whose derived geometry comes from a projective plane minus a line, with a vertex transitive automorphism group. Block count $(7, 6, 6)$.
 - (b) The geometry whose derived geometry comes from a projective plane. In the associated graph, there is also a geometry on 9 points. Block count $(7, 7, 7)$.
 - (c) The geometry $B(2, 4)$ with block count $(9, 4, 4, 2)$ and a variant with block count $(7, 6, 6)$.
 - (d) The geometry $B(3, 3)$ with block count $(7, 6, 6)$ and a variant with block count $(9, 4, 4, 2)$.
 - (e) An Albert geometry with block count $(8, 5, 5, 1)$ and a variant with block count $(7, 6, 6)$.
- (3) A pair of non-Albert geometries with block counts $(6, 10, 1, 2)$ and $(7, 6, 6)$.

14. Appendix: Amalgamation Classes (Tables)

The following table shows the full list of 27 amalgamation classes determined by constraints on triangles, not allowing free amalgamation, and with the associated Fraïssé limit primitive. The structures are assumed to have four nontrivial 2-types, all of them symmetric (e.g., one may think of these structures as complete graphs, with 4 colors of edges). The data are taken from [Che98], with slightly different notation but the same numbering. Those which can be interpreted as homogeneous metric spaces were put into a clearer form in §10.

#	ABD	CDD	AAC	ADD	AAD	BBD	CCA	CCD	BDD	BAA	AAA	DDD
1	1	1	0	1	0	0	0	0	0	0	0	0
2	1	1	0	1	0	0	0	0	0	0	0	1
3	1	1	0	1	1	0	0	0	0	0	0	0
4	1	1	0	1	0	0	0	0	0	0	1	0
5	1	1	0	1	1	0	0	0	0	0	0	1
6	1	1	0	1	0	0	0	0	0	0	1	1
7	1	1	0	1	1	0	0	0	0	0	1	0
8	1	1	0	1	1	0	0	0	0	0	1	1
9	1	1	0	0	1	0	0	0	0	0	0	1
10	1	0	0	1	0	0	0	0	0	0	1	1
11	0	1	0	1	0	1	0	0	0	0	0	1
12	0	1	0	1	1	1	0	0	0	0	0	1
13	0	1	0	1	0	1	0	0	0	0	1	1
14	0	1	0	1	1	1	0	0	0	0	1	1
15	1	1	0	1	0	1	0	0	0	0	0	1
16	1	1	0	1	1	1	0	0	0	0	0	1
17	1	1	0	1	0	1	0	0	0	0	1	1
18	1	1	0	1	1	1	0	0	0	0	1	1
19	1	1	0	0	1	1	0	0	0	0	0	1
20	1	1	0	0	1	1	0	0	0	0	1	1
21	1	1	1	0	1	0	0	0	0	0	0	1
22	1	1	1	1	1	0	0	0	0	0	0	0
23	1	1	1	1	1	0	0	0	0	0	0	1
24	1	1	1	1	0	0	0	0	0	0	1	1
25	1	1	1	1	1	0	0	0	0	0	1	1
26	1	1	1	0	1	0	0	1	1	0	0	1
27	1	0	0	1	1	0	1	1	1	1	1	1

TABLE 3. 27 amalgamation classes [Che98]. See §10

References

- [AH78] F. Abramson and L. Harrington, “Models without indiscernibles” *J. Symbolic Logic* **43** (1978), 572–600.
- [AL95] R. Akhtar and A. Lachlan, “On countable homogeneous 3-hypergraphs,” *Arch. Math. Logic* **34** (1995), 331–344.
- [AH06] M. Alfuraidan and J. Hall, “Smith’s theorem and a characterization of the 6-cube as distance-transitive graph,” *J. Algebraic Combin.* **24** (2006), 195–207.
- [AMp10] D. Amato and H. D. Macpherson, “Infinite metrically homogeneous graphs of diameter 3,” Preprint, April 2010.
- [Ash71] C. Ash, “Undecidable \aleph_0 -categorical theories,” *Notices Amer. Math. Soc.* **18** (1971), 423.
- [Big09] N. Biggs, “Strongly regular graphs with no triangles,” *Research Report, September 2009*. arXiv:0911.2160v1.
- [Bog02] S. A. Bogatyĭ, “Metrically homogeneous spaces,” *Uspekhi Mat. Nauk* **57** (2002), 3–22 (Russian); English translation in *Russian Mathematical Surveys*, bf 57 (2002), 221–240.
- [Bon09] A. Bonato, “The search for n -e.c. graphs” *Contrib. Discrete Math.* **4** (2009), 40–53.
- [Bon10] A. Bonato, “Bounds and constructions for n -e.c. tournaments,” *Contributions to Discrete Mathematics* **5** (2010), 52–66.
- [BM08] J. A. Bondy and U. R. S. Murty, *Graph Theory, Graduate Texts in Mathematics* **244**, Springer, 2008, xii+651 pp.
- [BDM07] C. Bridges, A. Day, and S. Manber, “ k -Saturated graphs,” REU report, Clemson, 2007, <http://www.math.clemson.edu/~kevja/REU/2007/>.
- [Br11] A. Brouwer, “Graph descriptions,” <http://www.win.tue.nl/~aeb/graphs/>.
- [BCN89] A. E. Brouwer, A. M. Cohen, and A. Neumaier, *Distance-Regular Graphs*, Springer-Verlag, New York, 1989, 485 pp.,
- [BH95] A. E. Brouwer and W. Haemers, “Association Schemes,” in *Handbook of Combinatorics, Vol. 1, 2*, 747–771. Elsevier, Amsterdam 1995.
- [BvL84] A. E. Brouwer and J. H. van Lint, *Strongly regular graphs and partial geometries*, in *Enumeration and design (Waterloo, Ont., 1982)*, 85–122, Academic Press, Toronto, ON, 1984.
- [Cam80] P. Cameron, “6-transitive graphs,” *J. Combinatorial Theory, Series B* **28** (1980), ‘68–179.
- [Cam90] P. Cameron, *Oligomorphic permutation groups, London Mathematical Society Lecture Note Series*, **152**. Cambridge University Press, Cambridge, 1990. viii+160 pp.
- [Cam96] P. Cameron, “The random graph,” in J. Nešetřil, R. L. Graham (Eds.), *The Mathematics of Paul Erdős*, Springer, Berlin, 1996, pp. 331–351.
- [Cam97] P. Cameron, “Oligomorphic groups and homogeneous graphs,” *NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci.*, 497, Kluwer Acad. Publ., Dordrecht, 1997.
- [Cam98] P. Cameron, “A census of infinite distance transitive graphs” *Discrete Math.* **192** (1998), 11–26.
- [Cam03] P. Cameron, “Homogeneous permutations,” in *Permutation patterns (Otago, 2003)*, *Electron. J. Combin.* **9** (2002/03), no. 2, Research paper 2, 9 pp.
- [Cam05] P. Cameron, “The random graph has the strong small index property,” *Discrete Math.* **291** (2005), 41–43.
- [CvL91] P. Cameron and J. H. van Lint, *Designs, graphs, codes and their links, London Mathematical Society Student Texts* **22**. Cambridge University Press, Cambridge, 1991. x+240 pp.
- [Che88] G. Cherlin, “Homogeneous tournaments revisited,” *Geom. Ded.* **26** (1988), 231–240.
- [Che93] G. Cherlin, “Combinatorial Problems connected with finite homogeneity,” *Proceedings of the International Conference on Algebra, Part 3 (Novosibirsk, 1989)*, 3–30, *Contemp. Math.* **131**, Part 3, Amer. Math. Soc., Providence, RI, 1992.
- [Che98] G. Cherlin, “The classification of countable homogeneous directed graphs and countable homogeneous n -tournaments.” *Mem. Amer. Math. Soc.* **131** (1998), no. 621, xiv+161 pp.
- [Che99] G. Cherlin, “Infinite imprimitive homogeneous 3-edge-colored complete graphs,” *J. Symbolic Logic* **64** (1999), 159–179.

- [ChH03] G. Cherlin and E. Hrushovski, *Finite Structures with Few Types*, Annals of Math. Studies **152**, 2003.
- [ChL00] G. Cherlin and B. Latka, “Minimal antichains in well-founded quasi-orders with an application to tournaments,” *J. Combin. Theory Ser. B*, **80** (2000), 258–276.
- [CSW93] G. Cherlin, D. Saracino, and C. Wood, “On homogeneous nilpotent groups and rings” *Proc. Amer. Math. Soc.* **119** (1993), 1289–1306.
- [CT02] G. Cherlin and S. Thomas, “Two cardinal properties of homogeneous graphs,” *J. Symbolic Logic* **67** (2002)
- [DNT86] J. Dixon, P. Neumann, and S. Thomas, “Subgroups of small index in infinite symmetric groups,” *Bull. London Math. Soc.* **18** (1986), 580–586.
- [Dun82] M. J. Dunwoody, “Cutting up graphs,” *Combinatorica* **2** (1982), 15–23.
- [Ehr72] A. Ehrenfeucht, “There are continuum \aleph_0 -categorical theories,” *Bull. Acad. Polon. Sci.* **20** (1972), 425–427.
- [EKR76] P. Erdős, D. Kleitman, B. Rothschild, “Asymptotic enumeration of K_n -free graphs,” in *Colloquio Internazionale sulle Teorie Combinatorie (Rome, 1973)*, Tomo II, 19–27. *Atti dei Convegni Lincei* **17** (1976), Accad. Naz. Lincei, Rome.
- [Fag76] R. Fagin, “Probabilities on finite models,” *J. Symb. Logic* **41** (1976), 50–58.
- [Fra54] R. Fraïssé, “Sur l’extension aux relations de quelques propriétés des ordres,” *Ann. Ecole Normale Sup.* **7** (1954), 361–388.
- [Gar76] A. Gardiner, “Homogeneous graphs,” *J. Comb. Th.* **20** (1976), 94–102.
- [Gla71] W. Glassmire, Jr., “There are 2^{\aleph_0} countably categorical theories,” *Bull. Acad. Polon. Sci.* **19** (1971), 185–190.
- [GK09] L. Graham and J. M. Kantor, *Naming Infinity: A True Story of Religious Mysticism and Mathematical Creativity*, Belknap Press of Harvard University Press, 2009.
- [Hen71] C. W. Henson, “A family of countable homogeneous graphs,” *Pacific J. Math.* **38** (1971), 69–83.
- [Hen72] C. W. Henson, “Countable homogeneous relational systems and categorical theories,” *J. Symb. Logic* **37** (1972), 494–500.
- [HHL593] W. Hodges, I. Hodkinson, D. Lascar, and S. Shelah, “The small index property for ω -stable ω -categorical structures and for the random graph,” *J. London Math. Soc.* **48** (1993) 204–218.
- [HN05] J. Hubička and J. Nešetřil, “Finite presentation of homogeneous graphs, posets and Ramsey classes,” *Probability in mathematics, Israel J. Math.* **149** (2005), 21–44.
- [Huš08] M. Hušek, “Urysohn space, its development and Hausdorff’s approach” *Topology and its Applications* **155** (2008), 1493–1501. Special number devoted to the proceedings of a Workshop on the Urysohn space, Beer-Sheva, 2006.
- [KLM89] W. Kantor, M. Liebeck, and H. D. Macpherson, “ \aleph_0 -categorical structures smoothly approximated by finite substructures,” *Proc. London Math. Soc.* **59** (1989), 439–463.
- [KPT05] A. Kechris, V. Pestov, and S. Todorcevic, “Fraïssé limits, Ramsey theory, and topological dynamics of automorphism groups,” *Geom. Funct. Anal.* **15** (2005), 106–189.
- [GGK96] M. Goldstern, R. Grossberg, and M. Kojman, “Infinite homogeneous bipartite graphs with unequal sides,” *Disc. Math.* **149** (1996), 69–82
- [GK78] Ya. Gol’fand and Yu. Klin, “On k -homogeneous graphs,” (Russian), *Algorithmic studies in combinatorics (Russian)* **186**, 76–85, (errata insert), “Nauka,” Moscow, 1978.
- [KPR87] Ph. Kolaitis, H. Prömel, and B. Rothschild, “ $K_{\ell+1}$ -free graphs: asymptotic structure and a 0–1 law,” *Trans. Amer. Math. Soc.* **303** (1987), 637–671.
- [Kom99] P. Komjáth, “Some remarks on universal graphs,” *Discrete Math.* **199** (1999), 259–265.
- [KMP88] P. Komjáth, A. Mekler and J. Pach, “Some universal graphs,” *Israel J. Math.* **64** (1988), 158–168.
- [KT01] D. Kuske and J. Truss, “Generic automorphisms of the universal partial order,” *Proc. Amer. Math. Soc.* **129** (2001), 1939–1948.
- [Lac82] A. Lachlan, “Finite homogeneous simple digraphs,” *Proceedings of the Herbrand symposium (Marseilles, 1981)*, 189–208, *Stud. Logic Found. Math.*, **107**, North-Holland, Amsterdam, 1982.
- [Lac84] A. Lachlan, “Countable homogeneous tournaments,” *Trans. Amer. Math. Soc.* **284** (1984), 431–461.
- [Lac86a] A. Lachlan, “Binary Homogeneous Structures II,” *Proc. London Math. Soc.* **52** (1986), 412–426.

- [Lac86b] A. Lachlan, “Homogeneous structures,” in **Proceedings of the ICM 1986**, American Mathematical Society, 1987, 314–321.
- [LT95] A. Lachlan and A. Tripp, “Finite homogeneous 3-graphs,” *Math. Logic Quart.* **41** (1995), 287–306.
- [Lac96] A. Lachlan, “Stable finitely homogeneous structures: a survey,” *Algebraic model theory (Toronto, ON, 1996)* 145–159, NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci., **496**, Kluwer Acad. Publ., Dordrecht, 1997.
- [LW80] A. Lachlan and R. Woodrow, “Countable ultrahomogeneous undirected graphs,” *Trans. Amer. Math. Soc.* **262** (1980), 51–94.
- [Lat94] B. J. Latka, “Finitely constrained classes of homogeneous directed graphs,” *J. Symb. Logic*, 59:124–139, 1994.
- [Lat02] B. J. Latka, “A structure theorem for tournaments omitting $IC_3(I, I, L_3)$,” Preprint, 1994, revised 2002.
- [Lat03] B. J. Latka, “A structure theorem for tournaments omitting N_5 ,” *J. Graph Theory* **42** (2003), 165–192.
- [Mph82] H. D. Macpherson, “Infinite distance transitive graphs of finite valency,” *Combinatorica* **2** (1982), 63–69.
- [Mph11] H. D. Macpherson, “A survey of homogeneous structures,” Preprint 2011. To appear, *Discrete Mathematics*.
- [Neš07] J. Nešetřil, “Metric spaces are Ramsey,” *European J. Combin.* **28** (2007), 457–468.
- [NR75a] J. Nešetřil and V. Rödl, “A Ramsey graph without triangles exists for any graph without triangles,” *Infinite and Finite Sets (Colloq., Keszthely, 1973; dedicated to P. Erdős on his 60th birthday, Vol. III)*, 1127–1132. *Colloq. Math. Soc. Janos Bolyai* **10**, North-Holland, Amsterdam, 1975.
- [NR75b] J. Nešetřil and V. Rödl, “Type theory of partition properties of graphs,” *Recent Advances in Graph Theory (Proc. Second Czechoslovak Sympos., Prague, 1974)* 405–412, Academia, Prague, 1975.
- [NR76] J. Nešetřil and V. Rödl, “The Ramsey property for graphs with forbidden complete subgraphs,” *J. Combin. Theory Ser. A* **20** (1976), 243–249.
- [NR77a] J. Nešetřil and V. Rödl, “A structural generalization of the Ramsey Theorem,” *Bull. Amer. Math. Soc.* **83** (1977), 127–128.
- [NR77b] J. Nešetřil and V. Rödl, “Partition of relational and set systems,” *J. Combin. Theory Ser. A* **22** (1977), 289–312.
- [NR83] J. Nešetřil and V. Rödl, “Ramsey classes of set systems,” *J. Combin. Theory Ser. A* **34** (1983), 183–201.
- [NR90] J. Nešetřil and V. Rödl, eds., *Mathematics of Ramsey theory*, Algorithms and Combinatorics, 5. Springer-Verlag, Berlin, 1990. xiv+269
- [Pes02] V. Pestov, “Ramsey-Milman phenomenon, Urysohn metric spaces, and extremely amenable groups” *Israel J. Math.* **127** (2002), 317–357.
- [PV08] F. Petrov and A. Vershik, “Invariant measures on the set of graphs and homogeneous uncountable universal graphs,” Preprint, 2008.
- [Rad64] R. Rado, “Universal graphs and universal functions,” *Acta Arith.* **9** (1964), 331–340.
- [Rub94] M. Rubin, “On the reconstruction of \aleph_0 categorical structures from their automorphism groups,” *Proc. London Math. Soc.* **69** (1994), 225–249.
- [SW86] D. Saracino and C. Wood, “Finite homogeneous rings of odd characteristic,” *Logic colloquium '84 (Manchester, 1984)*, 207–224, *Stud. Logic Found. Math.*, **120**, North-Holland, Amsterdam, 1986.
- [SWS93] N. Sauer, R. E. Woodrow, B. Sands, *Finite and Infinite combinatorics in Sets and Logic (Proc. NATO Advanced Study Inst., Banff, Alberta, Canada, May 4, 1991)*, Kluwer-Academic Publishers, 1993.
- [Sch79] J. Schmerl, “Countable homogeneous partially ordered sets,” *Alg. Univ.* **9** (1979), 317–321.
- [She74] J. Sheehan, “Smoothly embeddable subgraphs,” *J. London Math. Soc.* **9** (1974), 212–218.
- [Smi71] D. Smith, “Primitive and imprimitive graphs,” *Quart. J. Math. Oxford* **22**, (1971), 551–557.
- [Sz²65] E. Szekeres and G. Szekeres, “On a problem of Schütte and Erdős,” *Math. Gazette* **49** (1965), 290–293.

- [Tho91] S. Thomas, “Reducts of the random graph,” *J. Symbolic Logic* **56** (1991), 176–181.
- [Tho96] S. Thomas, “Reducts of random hypergraphs,” *Ann. Pure Appl. Logic* **80** (1996), 165–193.
- [Tru85] J. Truss, “The group of the countable universal graph,” *Math. Proc. Cambridge Phil. Soc.* **98** (1985), 213–245.
- [Tru92] J. Truss, “Generic automorphisms of homogeneous structures,” *Proc. London Math. Soc.* **65** (1992), 121–141.
- [Tru09] J. Truss, “On the automorphism group of the countable dense circular order,” *Fund. Math.* **204** (2009), 97–111.
- [Tru03] J. Truss, “The automorphism group of the random graph: four conjugates good, three conjugates better,” *Discrete Math.* **268** (2003), 257–271.
- [TrT08] J. Truss and S. Torrezão de Sousa, “Countable homogeneous coloured partial orders,” *Dissertationes Math. (Rozprawy Mat.)* **455** (2008), 48 pp.
- [Ury25] P. Urysohn, “Sur un espace métrique universel,” *CRAS* **180** (1925), 803–806.
- [Ury27] P. Urysohn, “Sur un espace métrique universel,” *Bull. Sci. Math.* **51** (1927), 43–64 and 74–90.
- [vLW92] J. H. van Lint and R. Wilson, *A course in combinatorics*, Cambridge University Press, Cambridge, 1992.
- [Wat07] S. Waton, “On Permutation Classes Defined by Token Passing Networks, Gridding Matrices and Pictures: Three Flavours of Involvement,” Ph.D. thesis, St. Andrews, 2007.

DEPARTMENT OF MATHEMATICS, RUTGERS UNIVERSITY, 110 FRELINGHUYSEN RD., PISCATAWAY, NJ 08854