

**METRICALLY HOMOGENEOUS GRAPHS:
DYNAMICAL PROPERTIES OF
THEIR AUTOMORPHISM GROUPS
AND THE CLASSIFICATION OF TWISTS**

by

REBECCA COULSON

A dissertation submitted to the
School of Graduate Studies
Rutgers, The State University of New Jersey
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy
Graduate Program in Mathematics

Written under the direction of
Gregory Cherlin
and approved by

New Brunswick, New Jersey

January, 2019

ABSTRACT OF THE DISSERTATION

Metrically Homogeneous Graphs: Dynamical Properties of their Automorphism Groups and the Classification of Twists

By REBECCA COULSON

Dissertation Director:

Gregory Cherlin

We investigate the properties of graphs which are homogeneous in the sense of Fraïssé when considered as metric spaces with the graph metric (metrically homogeneous graphs), and particularly the metrically homogeneous graphs of generic type constructed by Cherlin.

We first consider the properties of the associated automorphism groups, viewed as topological groups. For a large class of metrically homogeneous graphs of generic type, we show that the automorphism groups have ample generics, and therefore have a variety of topological properties such as the small index property and automatic continuity. We also show that the automorphism groups of the generic expansions of these graphs by linear orders are extremely amenable, and describe the universal minimal flow for the full automorphism group. Using standard model theoretic and descriptive set theoretic methods, these results are derived from the study of combinatorial properties of the associated classes of finite partial substructures.

Turning to more algebraic questions, we determine the twisted automorphism groups of metrically homogeneous graphs, and more generally the twisted isomorphisms between such

graphs; these are isomorphisms up to a permutation of the natural language. Returning to the standard automorphism group, we then study the algebra of the associated age in the sense of Peter Cameron, showing that in most cases this algebra is a polynomial algebra. For this, we apply a criterion of Cameron based on a unique decomposition theorem.

Acknowledgements

I would like to thank the Rutgers mathematics department, both staff and faculty, including the graduate program directors Professors Han, Miller, and Borisov. They always provided me with the support and guidance I needed, and occasionally went above and beyond to ensure that I would have a successful graduate career.

Thank you to my committee members, Professors Patel, Sargsyan, Thomas, and Cherlin. Thanks in particular to Professor Thomas, for introducing me to the beauty of logic and for some of the most amusing classes I've taken. If you hadn't been such an engaging teacher, I don't know if I would have gotten to know my husband.

Thank you to my parents for their uncompromising dedication to my education and growth. To say that I would not have accomplished this without your support is an understatement. Your example of autodidacticism is a precious gift, and one I will try to impart to my children.

Thank you to my husband Bud, for his positivity, and endless and selfless support. You have always encouraged me to be the best version of myself, and for that I am eternally grateful.

Thank you finally to my advisor, Gregory Cherlin. Your patience and generosity throughout this process has been astounding, and has been an education to me in so many ways. It has been an honor and a pleasure.

Dedication

In loving memory of my father, Mark R. Gordon

Table of Contents

Abstract	ii
Acknowledgements	iv
Dedication	v
1. Introduction	1
1.1. Topology and Dynamics of Automorphism groups	2
1.2. Algebraic Results	7
2. Background: Fraïssé theory, automorphism groups, and combinatorics	14
2.1. Homogeneity and Fraïssé Theory	16
2.2. Metrichally Homogeneous Graphs	24
2.3. Topological Properties of Automorphism Groups	46
2.4. Dynamical Properties of Automorphism Groups	54
2.5. Structural Ramsey Theory	57
2.6. Partial Structures and Finite Constraint	60
2.7. Twisted Isomorphisms	66
3. Topological Results	71
3.1. Finite constraint	72
3.1.1. $[\delta]$ -metric spaces	72
3.1.2. S -metric spaces	75
3.1.3. Metrichally Homogeneous Graphs	79
3.1.3.1. The set of constraints \mathcal{F}	81
3.1.3.2. Necessity	83
3.1.3.3. A completion process	86

3.1.3.4.	The completion of a candidate configuration	95
3.1.3.5.	Proof of the main theorem	108
3.2.	Ample generics and the topological group $\text{Aut}(\Gamma)$	112
3.3.	The Ramsey property and extreme amenability	114
4.	Algebraic Results	117
4.1.	Twists	117
4.1.1.	Non-generic type	118
4.1.2.	Generic type	121
4.2.	Twistable graphs	127
4.2.1.	Necessity of the restrictions on the parameters	128
4.2.2.	Sufficiency of the restrictions of the parameters	139
4.3.	Algebra of an Age	148
4.3.1.	The case of metrically homogeneous graphs	151
4.3.2.	The bipartite antipodal case	155
References	158

Chapter 1

Introduction

In this thesis, we study a number of questions concerning *metrically homogeneous graphs* and their automorphism groups.

A metric space is *homogeneous* in the sense of Urysohn and Fraïssé if and only if every isometry between finite subspaces extends to an isometry of the whole space [Ury27, Fra54]—that is, the metric and group theoretic notions of congruence coincide. A connected graph is said to be *metrically homogeneous* if it is homogeneous when considered as a metric space, with the path metric.

The finite metrically homogeneous graphs were completely classified by Cameron [Cam76]. The problem of classification in the infinite case was raised by Moss [Mos92] and Cameron [Cam98]. Cherlin gave a catalog of the known examples in [Che11], with some evidence for its completeness; additional work in this direction is found in [ACM16, Che17]. We deal here primarily with the properties of the *known* metrically homogeneous graphs, with some exceptions when the existing classification theory with modest extensions suffices to carry through an analysis on an a priori basis.

Our results fall under two main headings: *topological* and *algebraic*.

The topological results concern the automorphism groups of certain metrically homogeneous graphs of known type. As we will see, these groups are *Polish groups*; that is, they carry a natural topology, and are even complete with respect to a metric compatible with that topology. In Section 1.1, we discuss the topological properties, and the topological dynamics, of a substantial family of metrically homogeneous graphs of known type. Theorem 1 below deals with the relationship between topological properties of the automorphism group and its properties as an abstract group. Theorem 2 concerns dynamical properties of the automorphism group, derived from consideration of the dynamical properties of the automorphism group of a

closely related structure (Proposition 1). Here we rely on the model theoretic and descriptive set theoretic methods of [KPT05] and [KR07], which apply more generally to automorphism groups of *homogeneous structures* in the sense of Fraïssé (Section 2.1), reducing topological and dynamical problems to combinatorial problems which are themselves the subject of a rich theory. A key combinatorial ingredient of our analysis is the study of a *canonical completion process* taking a finite partial (i.e. weak) substructure of a given metrically homogeneous graph to an induced substructure of the same graph (viewed as a *metric space*). Another completion process which applies in greater generality was found slightly later by [ABH⁺17]; this allows the topological and dynamical results to be extended to a correspondingly broader class.

The algebraic results are of two kinds. On the one hand, following a line inspired by work of Cameron and Tarzi [CT17], we classify the *twisted automorphisms* between metrically homogeneous graphs. These are by definition isomorphisms up to a permutation of the associated language (Definition 2.7.2). These results, formulated as Theorem 3 below, apply to all metrically homogeneous graphs, and not just the known ones; however to achieve this, one must be careful about the way in which the associated metrically homogeneous graphs are characterized. A noteworthy feature of this analysis is that the permutations of the language which arise have previously been encountered in the study of finite association schemes with multiple P -polynomial structures [BB80, Gar80].

Finally, we leave the subject of automorphisms and twisted automorphisms briefly, to consider an algebraic invariant introduced by Peter Cameron, the *algebra of an age*. This is a graded algebra attached to a metrically homogeneous graph of finite diameter (or to any *oligomorphic group* in the sense of Definition 2.1.6). Our Theorem 4 concerns cases in which we can show that the associated algebra is a polynomial algebra, for the most part (with one interesting exception) in infinitely many variables.

We proceed now to listing our results.

1.1 Topology and Dynamics of Automorphism groups

Our work on the topological and dynamical properties of the automorphism groups of metrically homogeneous graphs relies on a substantial body of work which applies generally to

homogeneous structures for relational languages in the sense of Fraïssé. Classically, Fraïssé's theory, in the case of relational languages, relates the properties of a homogeneous structure Γ to the class of finite structures embedding in Γ (Section 2.1). Powerful methods for the study of the topological properties of automorphism groups of homogeneous structures were introduced by Hodges, Hodkinson, Lascar, and Shelah, extended by Herwig and Lascar, and systematized by Kechris and Rosendal [HHLS93, HL99, KR07], allowing one to work combinatorially in the associated class of finite structures, focusing on extension properties for partial automorphisms. In another direction, Kechris, Pestov, and Todorćević have given a reduction of certain dynamical properties of the automorphism groups to other combinatorial properties of these classes, closely related to the classical Ramsey theorem. As we will see in Chapter 3, while the topological and dynamical properties are not obviously related, and the corresponding combinatorial properties have no obvious connection, in the cases of interest here the same combinatorial analysis (given in Section 3.1) will suffice in both cases.

We begin with the topological results, following the approach of Kechris and Rosendal, which centers on the study of *ample generics*, that is, finite sequences of automorphisms whose conjugacy class under the action of the automorphism group by conjugation are co-meager in the sense of Baire category (Section 2.3). Our first result states that a broad class of metrically homogeneous graphs has ample generics, from which a number of striking topological consequences then follow. The precise statement runs as follows; we will explain both the hypotheses and the conclusions in more detail afterward.

Theorem 1. *Let $(\delta, K_1, K_2, C, C', S)$ be an admissible parameter sequence with δ and K_1 finite, for which $C > 2\delta + \max(K_1, \delta/2)$, $K_2 \geq \delta - 1$ and $C' = C + 1$, and let $\Gamma = \Gamma_{K_1, K_2, C, C', S}^\delta$ be the associated primitive metrically homogeneous graph of generic type.*

Then the Polish group $\text{Aut}(\Gamma)$ has ample generics, and therefore has the following properties.

- *The small index property and automatic continuity;*
- *uncountable cofinality;*
- *the fixed point properties (FA) and (FH).*

The small index property and automatic continuity express in different ways that the structure of the automorphism group as an abstract group determines the topology. The notions of uncountable cofinality and the fixed point properties (FA) and (FH) (concerning actions of G on trees or on Hilbert space) are properties of the abstract group G which can be derived from its topological properties.

Very precise assumptions are made concerning the metrically homogeneous graphs under consideration, which relate to the classification results reviewed in Section 2.2. These results in turn rely on the Fraïssé theory previously mentioned. The class of metrically homogeneous graphs divides naturally into those of *generic type* (Definition 2.2.4) and the rest. We focus here on the metrically homogeneous graphs of generic type.

While these graphs have not been fully classified, the ones which are associated to some set of forbidden triangles are characterized by five numerical parameters $(\delta, K_1, K_2, C, C')$, where δ is the diameter, and all known metrically homogeneous graphs of generic type are obtained from a graph in this family by imposing a set \mathcal{S} of so-called “Henson constraints” in the sense of Definition 2.2.9. In the statement of Theorem 1, we make additional assumptions on the values of these numerical constraints, which we discuss further below.

We take up next the dynamical properties of the same class of metrically homogeneous graphs. Here we apply the methods of Kechris, Pestov, and Todorčević presented in Section 2.4. In topological dynamics, one considers the *flows*, or continuous actions on compact spaces, of a given topological group, and more particularly the (set-theoretically) *minimal* flows (Definition 2.4.1). By a theorem of Ellis, there is a *universal minimal flow*, unique up to homeomorphism, having any other minimal flow as a continuous image. The universal minimal flow can be quite wild, and indeed there are examples of relatively simple groups G , e.g. countable discrete groups, which have non-metrizable universal minimal flows (see, for example, [KPT05, page 1]). At the opposite extreme, a topological group is said to be *extremely amenable* (Definition 2.4.2) if the universal minimal flow reduces to a single point; in other words, every flow contains a fixed point.

The results of Kechris, Pestov, and Todorčević relate the dynamical property of extreme amenability to the purely combinatorial Ramsey-theoretic properties studied by *structural Ramsey theory* (Section 2.5), and more generally related metrizability of the universal flow, and a

more concrete identification of the universal minimal flow, to the existence of suitable *Ramsey expansions* of a given structure.

In applying this theory, we first expand our metrically homogeneous graphs by a *generic linear order* (Definition 2.5.4) using the Fraïssé theory. This gives a different structure, and thus a different automorphism group, to which the results of Kechris, Pestov, and Todorčević apply in their most direct form, giving the following.

Proposition 1. *Let $(\delta, K_1, K_2, C, C', S)$ be an admissible parameter sequence with δ and K_1 finite, for which $C > 2\delta + \max(K_1, \delta/2)$, $K_2 \geq \delta - 1$ and $C' = C + 1$. Let $\Gamma = \Gamma_{K_1, K_2, C, C', S}^\delta$ be the associated primitive metrically homogeneous graph of generic type, and let $(\Gamma, <)$ be the generic expansion of Γ by a linear order.*

Then $\text{Aut}(\Gamma, <)$ is extremely amenable.

A further application of the theory of Kechris, Pestov, and Todorčević together with a general line of argument as formulated by Bodirsky then yields the following.

Theorem 2. *Let $(\delta, K_1, K_2, C, C', S)$ be an admissible parameter sequence with δ and K_1 finite, for which $C > 2\delta + \max(K_1, \delta/2)$, $K_2 \geq \delta - 1$ and $C' = C + 1$, and let $\Gamma = \Gamma_{K_1, K_2, C, C', S}^\delta$ be the associated primitive metrically homogeneous graph of generic type.*

Then the universal minimal flow of $\text{Aut}(\Gamma)$ is metrizable.

Furthermore, the universal minimal flow of $\text{Aut}(\Gamma)$ is the space $L(\Gamma)$ of all linear orderings of Γ .

Here, $L(\Gamma)$ is viewed as a closed subset of $2^{\Gamma \times \Gamma}$ in the product topology, and hence as a compact topological space.

The dynamical properties of $\text{Aut}(\Gamma)$ are connected with the Ramsey theoretic properties of the expansion of Γ by a generic linear order. One of the most powerful tools in structural Ramsey theory, the *partite method*, works systematically with partial substructures. This is particularly visible in the work of Nešetřil on the Ramsey property for metric spaces, where partial metric spaces are viewed as edge-labeled graphs, and there is a canonical completion process given by the path metric. This method is systematized by Hubička and Nešetřil in [HN16]. Under mild assumptions (strong amalgamation, Definition 2.1.9), a sufficient condition for the Ramsey property is that the collection of partial structures which do not embed

in Γ contains finitely many minimal elements, up to isomorphism; we call this property *finite constraint* (Definition 2.6.2). The conditions given in [HN16] are considerably more general.

The following is proved in Section 3.1.

Proposition 2. *Primitive metrically homogeneous graphs $\Gamma_{K_1, K_2, C, C', S}^\delta$ of generic type for which $C > 2\delta + \max(K_1, \delta/2)$, $C' = C + 1$, and $K_2 \geq \delta - 1$, are finitely constrained.*

This is the essential point, and the numerical assumptions imposed here and in our main results derive from the structure of the proof of this proposition.

Namely, the method of the proof is to give a *canonical completion process*, which takes a finite partial Γ -structure and completes it to a finite induced substructure of Γ . More generally, given any finite structure in the appropriate language, the process attempts to complete it and succeeds where possible. Proposition 2 results from an analysis of the obstructions to the completion process.

It turns out that the numerical assumptions needed for our completion process can be relaxed, and almost eliminated, by adopting a modified completion process. This process is given in [ABH⁺17]. In consequence, the topological and dynamical consequences can then be derived in the same manner under considerably weaker hypotheses.

We have explained the combinatorial connection between Proposition 2 and Ramsey theory, and hence with the dynamical properties of automorphism groups. Remarkably, Proposition 2 together with the associated completion process is sufficient to derive the ostensibly unrelated property of ample generics, and thus give the purely topological consequences as well. The reason for this is that results of Herwig and Lascar (Theorem 2.3.4) apply when one works with partial substructures rather than induced substructures; but the finiteness condition, which is usually both true and obvious in the context of *induced* substructures, has to be derived anew in the partial category.

We find it useful to systematize the use of Theorem 2.3.4 in the context of partial structures somewhat beyond what we have seen in the literature. Thus in Lemmas 2.6.1 and 2.6.2 of Section 2.6, we lay out explicitly the relevant connections between the properties of the class of finite induced substructures and the class of finite partial structures. We also formalize the notion of a *canonical* completion process in Definition 2.6.4.

1.2 Algebraic Results

In Chapter 4, we take up two purely algebraic questions which have been considered in the context of homogeneous structures, and investigate them in the context of metrically homogeneous graphs. In Sections 4.1, 4.2, we investigate the *twisted automorphism* groups of metrically homogeneous graphs in the sense of Cameron and Tarzi [CT17], and more generally the classification of twisted isomorphisms between metrically homogeneous graphs. In Section 4.3, following Cameron [Cam97], we investigate the *algebra of an age* of a metrically homogeneous graphs, and, specifically, the question as to whether this is a polynomial algebra (typically, in infinitely many variables).

The groups of *twisted automorphisms* of some analogs of the random graph are investigated by Cameron and Tarzi in [CT17]. These are, in a natural sense, the automorphisms of the structure up to a permutation of the underlying language (Definition 2.7.2). The more general notion of *twisted isomorphism* has arisen in some model theoretic contexts as well, sometimes under the name *permorphism* (isomorphism up to permutation). Some unexpected twisted isomorphisms between non-isomorphic metrically homogeneous graphs played a significant role in simplifying the proof of the classification results in [ACM16], as a classification up to twisted isomorphism immediately gives a classification up to isomorphism. Thus in this context, the classification of twisted isomorphisms has some practical implications.

In Section 4.1 of this paper, we find all possible permutations of the language which transform some metrically homogeneous graph into another metrically homogeneous graph. In Section 4.2 we analyze in each case which pairs of metrically homogeneous graphs are “twistable” to each other by a twisted automorphism affording the specified permutation of the language.

The level of generality of this work is considerably greater than that of Chapter 3, and in fact we give a satisfactorily complete classification of the twisted automorphisms of metrically homogeneous graphs in spite of the absence of a complete classification of the class of graphs in question. However, we do rely heavily on the existing classification theory, both to break down the question into meaningful parts, and to effect the solution even when the relevant class of graphs is not fully determined.

This classification problem for twisted isomorphisms breaks up naturally into two questions, which are treated separately in Sections 4.1 and 4.2, namely the following.

- (A) What non-identity permutations of the language of integer-valued metric spaces (viewed as edge-labeled graphs) can be induced by a twisted isomorphism between two metrically homogeneous graphs?
- (B) For each such permutation, what are the metrically homogeneous graphs which allow a twisted isomorphism of the corresponding type?

Not surprisingly, some elements of classification theory are needed even for Problem (A), and certainly for Problem (B). In fact, one would not expect a meaningful answer to Problem (B) in the absence of a full classification of the metrically homogeneous graphs, but it turns out that the existing classification theory allows a full solution of both problems, if the solution is phrased with care in the case of Problem (B).

Before presenting our results in detail, we review some essential points from the classification theory for metrically homogeneous graphs, given in detail in Section 2.2.

In the classification of metrically homogeneous graphs, a fundamental distinction is made between *generic* and *non-generic* type (Definition 2.2.4). Roughly speaking, the non-generic graphs occur “in nature” and result from explicit constructions, whereas for the most part the generic type graphs are only known via the Fraïssé theory (Section 2.1). Then within the class of metrically homogeneous graphs of generic type, an important subclass consists of the so-called *3-constrained* graphs (Definition 2.2.5), which are defined in terms of the Fraïssé theory of Section 2.1 as the metrically homogeneous graphs whose induced subspaces (as a metric space) are determined by constraints on metric triangles.

The existing classification theory provides the following.

- An explicit, complete list of the isomorphism types of metrically homogeneous graphs of non-generic type.
- An explicit, complete list of the isomorphism types of 3-constrained metrically homogeneous graphs of generic type, in terms of five numerical parameters δ, K_1, K_2, C, C' .

- A *conjectural* classification of the metrically homogeneous graphs of generic type in terms of the 3-constrained ones together with the “Henson constraints” in the sense of Definition 2.2.9.
- A precise definition of the numerical parameters δ, K_1, K_2, C, C' associated with an arbitrary metrically homogeneous graphs, and some partial information about their properties (not sufficient as yet to show that these parameters actually correspond to some 3-constrained homogeneous graph).

It is not difficult to classify the twisted isomorphisms of metrically homogeneous graphs of non-generic type “by inspection,” and there are few non-trivial cases. The precise statement of this result is given in Proposition 4.1.1 of Section 4.1.

What happens in the case of generic type is more interesting. Using the classification theory, one can determine the non-trivial permutations of the language associated with some twisted isomorphism explicitly. It turns out that the diameter δ must be finite and that only four such permutations can occur, for fixed δ . Remarkably, these are the same permutations that were found by Bannai and Bannai [BB80], and also by Tony Gardiner [Gar80], in the study of “ P -polynomial structures” for finite association schemes, or for finite distance regular graphs, respectively.

Problem. *Can one obtain a similar classification of twisted isomorphisms for all distance transitive graphs?*

What we have in mind here is only a classification of the relevant permutations of the language, not a classification of the structures involved. Such a result would represent a considerable strengthening of a major portion of our work on the topic, and clarify its relationship with the results of [BB80, Gar80].

Our main result is as follows.

Theorem 3. *Let $\delta \geq 3$ be fixed (potentially infinite), and let σ be a non-trivial permutation of the language of metrically homogeneous graphs of diameter δ .*

If there is a metrically homogeneous graph Γ of generic type such that Γ^σ is again a metrically homogeneous graph, then δ is finite and σ is one of the permutations $\rho, \rho^{-1}, \tau_0,$ or τ_1 from Proposition 4.1.2.

Conversely, if δ is finite and σ is one of the permutations ρ, ρ^{-1}, τ_0 , or τ_1 , with $\delta \geq 3$, then there is a metrically homogeneous graph Γ for which Γ^σ is again a metrically homogeneous graph. Furthermore, the metrically homogeneous graphs Γ whose images Γ^σ are also metrically homogeneous are precisely those with the numerical parameters K_1, K_2, C, C' as in Table 4.1.

Here the notation Γ^σ represents the canonical homogeneous structure (in the sense of Section 4.1) obtained by permuting the symbols of the language according to σ . If Γ is a metric space, then these symbols are binary relation symbols indexed by the possible non-zero values of the metric, and σ may be viewed more concretely as a permutation of the set of such values.

The explicit classification of the non-trivial permutations $\rho, \rho^{-1}, \tau_0, \tau_1$ associated with twisted isomorphisms between metrically homogeneous graphs relies on some elements of the classification theory to show that certain properties of the known metrically homogeneous graphs of generic hold for all metrically homogeneous graphs of generic type (see, in particular, Proposition 2.2.1 of Section 2.2), and as each of these permutations does occur for a known pair of metrically homogeneous graphs, we arrive at a complete classification of these permutations.

We express the classification of the metrically homogeneous graphs allowing twisted isomorphisms associated with these permutations in terms of the associated numerical parameters δ, K_1, K_2, C, C' . We must now discuss the extent to which this result is in fact a classification.

The underlying point is that the question, whether a given metrically homogeneous graph Γ allows a twisted isomorphism to another such graph which is associated with a specified permutation of the language, is in fact equivalent to a set of conditions on the parameters δ, K_1, K_2, C, C' . This would clearly be the case if these parameters were known to determine precisely the types of metric triangles embeddable in Γ , a point that remains conjectural. So this point requires some additional attention in the absence of a full classification.

Once we have determined that the existence or non-existence of twisted isomorphisms of a specified type is completely determined by the values of the associated numerical parameters, the final statement in Theorem 3 gives a satisfactory solution to the problem posed. It does, however, leave open certain questions that only a fuller classification result would resolve. Namely, the numerical parameters associated with the known metrically homogeneous graphs

of generic type satisfy additional constraints (“admissibility,” Definition 2.2.10). Conjecturally, the numerical parameters associated to any metrically homogeneous graph of generic type are admissible. Regardless, the results of Theorem 3 apply to any choice of numerical parameters which is in fact associated to some metrically homogeneous graph of generic type.

In addition, Theorem 3 as stated does not entirely settle the corresponding problem for *twisted automorphisms*, but it could easily be reworded to fit that case. The Fraïssé theory provides a canonical description of any metrically homogeneous graph of generic type, and in order for a twisted isomorphism to be a twisted automorphism, one must simply check whether the associated description (in terms of minimal forbidden subspaces) is invariant under the corresponding permutation of the language. The twisted isomorphisms associated with twisted automorphisms are the involutions τ_0, τ_1 , witnessed by known metrically homogeneous graphs of generic type.

To transform Theorem 3 into a completely explicit list of metrically homogeneous graphs, their twisted isomorphisms, and their twisted automorphism groups, would require solving some instances of the classification problem, namely those associated with the values of the numerical parameters shown in Table 4.1. This is a large problem in itself, though possibly more tractable than the full classification.

In the context considered by Cameron and Tarzi [CT17], the structure of the permutation group of the language induced by the twisted automorphism group was trivial, but a more subtle question concerning the structure of the full group of twisted automorphisms was taken up: when does this group split over the normal subgroup of ordinary automorphisms? In our context, this amounts to a problem of lifting permutations of order two to twisted automorphisms of order two. We hope to address this problem elsewhere.

Now we come to our second algebraic topic, namely the algebraic invariant introduced by Peter Cameron, called the *algebra of an age* (Definition 4.3.1), and more particularly the question as to when this invariant is a polynomial algebra.

There is very rich theory concerning the so-called *profile* of an oligomorphic permutation group. A permutation group G acting on a set S said it be *oligomorphic* if the number of orbits of G acting on n -element sets of S is finite for every n . The profile is the function giving the number of orbits of the group on (unordered) sets of order n ; thus, if a permutation group is

oligomorphic, then the profile is finite everywhere.

Considerable attention has been paid to the asymptotic behavior of this function, and a more precise description has been sought for the possibilities in the polynomially bounded case.

A powerful algebraic tool was introduced by Peter Cameron in [Cam81], a graded \mathbb{Q} -algebra \mathcal{A}^G (where G is the permutation group in question) whose Hilbert function is the generating function for the profile. The definition of this algebra is given in §4.3. It may be described succinctly as the ring of G -invariants in the incidence algebra of the partially ordered set of finite subsets of the domain Γ on which G acts.

Cameron showed that multiplication by a suitable element of degree 1 gives an injection from degree n to degree $n + 1$, and deduced that the profile function is non-decreasing, [Cam76, Theorem 2.2]. This argument was given prior to the formal introduction of the algebra, in terms of a linear map which is in fact the relevant multiplication map.

Once this algebra is introduced, the following questions are natural.

- When is it an integral domain?
- When is it finitely generated?
- When is it a polynomial ring over \mathbb{Q} (typically with infinitely many generators)?

Already in [Cam81], Cameron conjectured that the algebra \mathcal{A}^G is an integral domain if and only if G has no finite orbits. This was proved by Pouzet [Pou08].

In the case in which the growth rate of the profile is polynomially bounded, MacPherson asked in [Mac85] whether the associated algebra must be finitely generated, and in [Cam90], Cameron asked whether the profile must then be asymptotically polynomial. A positive solution to these questions, and considerably more, has been announced by Falque and Thiery [FT18].

We will prove the following in §4.3.

Theorem 4. *Let $(\delta, K_1, K_2, C, C', S)$ be an admissible parameter sequence with K_1 and δ finite, and let Γ be the corresponding metrically homogeneous graph, with automorphism group G . If $C = 2\delta + 1$, suppose that δ is even. Then the associated algebra \mathcal{A}^G is a polynomial algebra in infinitely many variables.*

This depends on two ingredients: a criterion introduced by Cameron in [Cam97], and a “disjoint sum” operation for metrically homogeneous graphs suggested by the “magic parameter” used in [ABH⁺17]. This magic parameter is used by Aranda et al. in order to complete compatible edge-labeled graphs to metrically homogeneous graphs in $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$. Their approach is a more general version of the completion process we develop in §3.1.3, which was developed for an altogether different purpose than that of Cameron in [Cam97].

The general imprimitive case (Definition 2.1.7) remains open, though seems tractable.


Open Problem. Is the algebra \mathcal{A}^G associated with the automorphism group of an imprimitive metrically homogeneous graph of generic type a polynomial algebra?

The discussion in [Cam97, §3] contains additional open problems of this type.

Chapter 2

Background: Fraïssé theory, automorphism groups, and combinatorics

The present chapter provides the technical tools required for our work in the remainder of the thesis. Most of this material concerns the general theory of *homogeneous structures*. There is now an elaborate “glossary” which provides a systematic method for relating three different points of view on the subject.



An amalgamation class of finite structures provides a setting for combinatorial analysis: examples would be the classes of all finite graphs, all finite tournaments, all finite partial orders, all finite metric spaces, or all finite permutations. The Fraïssé theory reviewed in Section 2.1 associates to each such class satisfying a few key properties, most notably the eponymous *amalgamation property*, a canonical *Fraïssé limit*, usually infinite. For example, the Fraïssé limit of the class of finite linear orders is isomorphic to the rational order, while the Fraïssé limit of the class of finite graphs gives the so-called *random graph*; combining these examples, the Fraïssé limit of the class of finite ordered graphs gives the *random graph with a generic linear order*, or more loosely speaking, the *random ordered graph*.

Fraïssé’s theory goes back to the 1950’s. More recently, the glossary has been substantially expanded to relate topological and dynamical properties of the automorphism groups of homogeneous structures to the combinatorial properties of the associated classes of finite combinatorial structures.

The seminal work of [KPT05], inspired by prior work in dynamics, gives an exact translation of some dynamical properties, notably *extreme amenability* and related notions, into combinatorial equivalents which have been extensively studied under the heading of *structural Ramsey theory* in a combinatorial setting.

In a quite separate development with a similar flavor, a member of loosely related properties of the automorphism groups of homogeneous structures have been shown to follow from a combinatorial condition which, while not equivalent to any or all of the desired properties, occurs quite often in practice. The main combinatorial property required is the so-called *extension property for partial automorphisms* or EPPA. Unlike the case of structural Ramsey theory, the associated combinatorial theory has only been taken up systematically after the connection with the topology of automorphism groups became clear. The technique was introduced in [HHL93], developed more broadly by [HL99], and further systematized by [KR07], and remains the subject of active development.

We present the Fraïssé theory in Section 2.1. In Section 2.3, we present topological theory of Kechris and Rosendal, and then the dynamical theory of Kechris, Pestov, and Todorčević in Section 2.4. In Section 2.3 we also take note of the formulation of a special case in [Sin17] which is particularly convenient for our purposes. In addition, Sections 2.5 and 2.6 present the combinatorial tools which apply once the reductions of Fraïssé, Kechris-Pestov-Todorčević, and Kechris-Rosendal have been made. Specifically, Section 2.5 deals with the structural Ramsey theoretic results used to obtain dynamical applications, which have been further systematized by [HN16], while Section 2.6 deals with a finiteness condition which plays a double role in our work: in the first place, it yields a finiteness condition required for the application of [HN16], and in the second place, it also provides a finiteness condition required to apply results of Herwig and Lascar in the manner of Siniora [Sin17]. The bulk of Section 2.6 fills an expository gap in the literature, connecting the general Fraïssé theory to the more specialized theory of Herwig and Lascar in a systematic way.

A common theme in the combinatorial work relating to the methods of Kechris-Pestov-Todorčević and Kechris-Rosendal is the study of *partial* (or “weak”) substructures, where the Fraïssé theory and the associated topological or dynamical reductions involve *induced* substructures. The use of partial structures is prominent in the application of the so-called *partite*

method in structural Ramsey theory, more so in some cases than others. A typical example in which this feature is particularly visible is the proof of the structural Ramsey theorem for metric spaces in [Neš07], which makes explicit use of *partial metric spaces* and relies on a clear understanding of the consequences of the triangle inequality in the setting of a partial metric space. This is a critical example for the Ramsey theory of metrically homogeneous graphs as well.

The point of working with, say, partial metric spaces rather than metric spaces is that the free amalgam of metric spaces is a partial metric space, and the free amalgam construction is involved in applications of the partite method. Similarly, Herwig-Lascar work with free amalgamation classes, and hence one must also make a transition to partial structures to apply their results in the manner of Kechris-Rosendal.

This thesis also relies heavily on the existing theory of metrically homogeneous graphs, which is review in Section 2.2. In Chapter 3, we investigate the automorphism groups of the known metrically homogeneous graphs, so a close acquaintance with the catalog of known graphs and the associated notation is required. In the first section of Chapter 4, we work more broadly, and so we also require some detailed results emerging from the existing classification theory. At the end of Chapter 4, we will return to the context of known metrically homogeneous graphs.

Finally, in Section 2.7, we will introduce the notion of *twisted isomorphism*, and consider the theory and associated problems relating to twisted automorphism groups from the point of view of the general theory of homogeneous structures, in preparation for the detailed consideration of the classification of twisted isomorphisms between metrically homogeneous graphs undertaken in Sections 4.1 and 4.2.

2.1 Homogeneity and Fraïssé Theory

The study of metrically homogeneous graphs lies within the broader subject of *homogeneous structures* in the sense of Fraïssé, and we will make extensive use of Fraïssé's general theory, and its more recent extensions, throughout the present thesis. Therefore, we begin by presenting the main notions and results of this theory, before turning in the next section to the specific

context of metrically homogeneous graphs. Further information, and proofs, may be found in [Hod97].

Definition 2.1.1. *A structure is homogeneous (in the sense of Fraïssé) if every isomorphism between finitely generated substructures of Γ extends to an automorphism of Γ .*

There are several other notions of homogeneity in use in model theory, permutation group theory, and metric geometry. Fraïssé’s notion is often called *ultrahomogeneity*.

In the present work, the only homogeneous structures under consideration will be countable, and in later sections we omit explicit mention of that condition. However, that restriction plays a fundamental role in Fraïssé’s theory, on which we rely throughout. Furthermore, we will specialize shortly to the case of relational languages (that is, there are no function symbols) in which case “finitely generated” simply means “finite.”

The motivating example for the theory is the case of the rational order $(\mathbb{Q}, <)$, which is a homogeneous linear order, and Cantor’s uniqueness theorem: any two countable dense linear orders without endpoints are isomorphic. This may be generalized as follows.

Theorem 2.1.1. *Any two countable homogeneous structures with the same finitely generated substructures (up to isomorphism) are isomorphic.*

This theorem can be proved in the same manner as Cantor’s isomorphism theorem, using the back-and-forth method introduced by Hausdorff in his exposition of that result [Hau14].

Theorem 2.1.1 suggests a general program of expressing the theory of countably homogeneous structures in terms of the associated classes of finitely generated structures. The fundamental question which Fraïssé’s theory addresses is the characterization of the classes of finitely generated structures which are associated for countably homogeneous structures.

Definition 2.1.2. *The age of a structure Γ is the class of finitely generated structures isomorphic to some finitely generated substructure of Γ .*

If \mathcal{A} is the age of a homogeneous structure Γ , then Γ is called the Fraïssé limit of \mathcal{A} .

For example, the age of any infinite linear order is the class of all finite linear orders, and in particular there is only one countably infinite homogeneous linear order, up to isomorphism.

That is, the Fraïssé limit of the class of finite linear orders is a countable dense linear order without endpoints.

A synonym for the *Fraïssé limit* is the *generic \mathcal{A} -structure*. For example, if \mathcal{A} is the class of all finite linear orders, the Fraïssé limit would be called the *generic linear order*.

It is easy to characterize the classes \mathcal{A} which arise as the age of some countable structure Γ , in terms of the following properties.

Definition 2.1.3. *Let \mathcal{A} be the class of finitely generated structures.*

1. \mathcal{A} is invariant if \mathcal{A} is closed under isomorphism.
2. \mathcal{A} is hereditary if \mathcal{A} is closed under taking induced finitely generated substructures.
3. \mathcal{A} has joint embedding if for any pair of structures A, B in \mathcal{A} there is a structure in \mathcal{A} into which A and B embed.

These conditions are clearly satisfied by the age of a structure. The ages of countable structures are characterized as follows.

Fact 2.1.1. *Let \mathcal{A} be a class of finitely generated structures in a fixed language. Then the following are equivalent.*

- \mathcal{A} is the age of some countable structure.
- \mathcal{A} is invariant, hereditary, has the joint embedding property, and only countably many isomorphism types of structures occur in \mathcal{A} .

This is easily verified. The characterization of the ages of homogeneous structures is more subtle and involves a further property.

Definition 2.1.4. *A class \mathcal{A} of structures has the amalgamation property if and only if for any triple A_0, A_1, A_2 and any embeddings $f_1: A_0 \rightarrow A_1$, $f_2: A_0 \rightarrow A_2$, there is a structure $A \in \mathcal{A}$ and there are embeddings $g_1: A_1 \rightarrow A$, $g_2: A_2 \rightarrow A$ so that the compositions $g_1 \circ f_1$ and $g_2 \circ f_2$ agree on A_0 .*

The structure A is said to be an amalgam of A_1, A_2 over A_0 .

Theorem 2.1.2. *Let \mathcal{A} be a class of finitely generated structures in a fixed language. Then the following are equivalent.*

- \mathcal{A} has a Fraïssé limit.
- \mathcal{A} is invariant, hereditary, has the joint embedding and amalgamation properties, and only countably many isomorphism types of structures occur in \mathcal{A} .

The existence of the Fraïssé limit is proved by building the structure Γ as the direct limit of a sequence of finitely generated structures so as to satisfy the following extension property: for any embedding $f: A \rightarrow B$ with A a finitely generated substructure of Γ and $B \in \mathcal{A}$, there is an embedding $g: B \rightarrow \Gamma$ with $g \circ f$ the identity on A . From this, homogeneity follows by a back-and-forth argument.

Classes \mathcal{A} satisfying all the specified conditions in Theorem 2.1.2 are called *amalgamation classes*; in other words, amalgamation classes are the ages of countable homogeneous structures.

As we have mentioned above, Fraïssé's theory has since been extended to express highly nontrivial properties of the *automorphism groups* of countable homogeneous structures in terms of the associated amalgamation classes.

The level of generality of Fraïssé's theory is very broad, and in a certain sense encompasses all countable structures, as we now indicate.

Definition 2.1.5. *Let Γ be a countable structure in a language \mathcal{L} . The canonical language \mathcal{L}_Γ for Γ consists of one n -ary relation symbol for each orbit of $\text{Aut}(\Gamma)$ on the collection of n -tuples of distinct elements of Γ .*

The associated \mathcal{L}^ -structure Γ^* has the same underlying set as Γ , with each relation of \mathcal{L}^* interpreted as the corresponding orbit, viewed as an n -ary relation.*

Remark 2.1.1. *If Γ is a countable \mathcal{L} -structure, then the associated \mathcal{L}_Γ -structure Γ^* is homogeneous and has the same automorphism group.*

This places some limits on what one may expect from Fraïssé theory, since in a certain sense the theory is applicable to an arbitrary structure. Nonetheless, this theory provides the conceptual framework for everything we do here.

Canonical structures are homogeneous for a relational language. So the condition of homogeneity in a relational language is not very restrictive. However, the condition of homogeneity with respect to a *finite relational language* is very restrictive, as we will now see.

Definition 2.1.6. *A permutation group G acting on a set X is oligomorphic if G has only finitely many orbits in the induced action on X^n , for each n .*

Remark 2.1.2. *A homogeneous structure in a finite relational language \mathcal{L} has an oligomorphic automorphism group.*

This follows easily from the fact that with n fixed, there are only finite many \mathcal{L} -structures of order n , up to isomorphism, and hence also up to ordered isomorphism (i.e. as n -tuples rather than sets).

By the theorem of Ryll-Nardzewski, Engeler, and Svenonius [RN59, Eng61, Sve59], having an oligomorphic automorphism group is one of the characterizations of \aleph_0 -categoricity of the corresponding theory.

It will be instructive to examine the case of *homogeneous graphs* from the point of view of Fraïssé theory. Thus the language is the language of graph theory, with a single symmetric binary relation which we take to be irreflexive as well.

The following substantial result will serve to illustrate several features of the Fraïssé theory, and some concepts of considerable importance to our work here.

Fact 2.1.2. [Gar80, She74, LW80] *Up to isomorphism and graph complementation, the countable homogeneous graphs are as follows.*

- Two “sporadic” examples: the 5-cycle C_5 and the graph on $[3]^2$ in which the edge relation is defined by $E((i, j), (i', j'))$ if and only if $i = i'$ or $j = j'$.
- The graphs $m \cdot K_n$ (the disjoint union of m n -cliques) with $1 \leq m, n \leq \infty$.
- Henson’s generic K_n -free graphs with $3 \leq n < \infty$ [Hen71].
- The generic or “random” graph.

To begin with, we notice that the list is given only up to graph complementation. In the canonical language associated with a homogeneous graph, there is a symbol for the edge relation E and a symbol for the non-edge relation N (also taken to be irreflexive). Forming

the graph theoretic complement involves a permutation of the symbols of the canonical language in which E, N are switched. Permutations of the canonical language carry homogeneous structures to homogeneous structures, so classification results need only be given up to such permutations. This observation leads to the consideration of twisted isomorphisms, to which we return at the end of this chapter.

In dealing with homogeneous structures, it is useful to distinguish the *primitive* and *imprimitive* cases.

Definition 2.1.7. *A homogeneous structure is imprimitive if it carries a non-trivial equivalence relation definable without parameters. Equivalently, there is a non-trivial equivalence relation invariant under the action of the automorphism group.*

A structure which is not imprimitive is said to be primitive.

We remark that in permutation group theoretic usage, the terms “primitive” and “imprimitive” are only applied to structures with transitive automorphism groups, and in fact all the homogeneous structures which concern us will satisfy this transitivity condition.

A first step toward the classification of the homogeneous graphs is the following straightforward result.

Remark 2.1.3. *Up to complementation, an imprimitive homogeneous graph is a disjoint union of at least two non-trivial cliques, of constant size.*

Indeed, by homogeneity, a non-trivial equivalence relation must be the reflexive extension of the edge relation or the non-edge relation, so up to complementation the equivalence classes are cliques, and the rest follows by vertex transitivity.

Next we consider the use of the term “generic” in various senses in connection with Fraïssé theory.

In the context of relational languages, amalgamation classes are invariant hereditary classes of finite structures, and therefore may be characterized in terms of the *minimal forbidden substructures*: these are the minimal structures which are not in the given class \mathcal{A} . Thus one may construct homogeneous structures by specifying a set \mathcal{F} of forbidden substructures, verifying that the class so defined has the amalgamation property, and taking the Fraïssé limit, which is then called the *generic \mathcal{F} -free structure*.

We have used this notion of genericity twice in the statement of Fact 2.1.2. Namely, taking \mathcal{A} to be the class of all finite graphs, one obtains the *generic graph*. Taking \mathcal{A} to be the class of all K_n -free graphs, that is, graphs containing no clique of order n , we obtain the *generic K_n -free graph* [Hen71].

Of course, one must check the amalgamation property in all such cases, and this is not always trivial. However for the class of all finite graphs, or all K_n -free finite graphs, it suffices to use *free amalgamation* in the following sense.

Definition 2.1.8. *The free amalgam of two relational structures A_1, A_2 over their intersection $A_0 = A_1 \cap A_2$ is the set $A_1 \cup A_2$ equipped with the union of the corresponding relations. More generally, given a structure A_0 and two embeddings $f_1: A_0 \rightarrow A_1, f_2: A_0 \rightarrow A_2$, one forms the free amalgam by first replacing A_1, A_2 by isomorphic copies A'_1, A'_2 for which the functions f_1, f_2 correspond to inclusion maps, with $A'_1 \cap A'_2 = A_0$, and then one takes the usual free amalgam of A'_1 with A'_2 over A_0 .*

This notion makes sense also when A_1, A_2 are equipped with some unary functions, but functions of more than one variable would no longer be functions in the free amalgam as we define it.

It is easy to see that the class of all finite graphs, or all finite K_n -free graphs, is closed under free amalgamation. Thus the Fraïssé theory immediately provides the corresponding homogeneous graph. The generic graph is also called the *random graph* because there is a natural probabilistic construction which with probability 1 produces the same graph, up to isomorphism. On the other hand, there is no comparably natural probabilistic construction producing the generic K_n -free graphs.

Note that passage to graph complements leads to a different notion of free amalgamation, in which the free join of two graphs treats the “default” condition on pairs as the *edge* relation rather than the *non-edge* relation. In particular, the complement of the generic K_n -free graph is the generic I_n -free graph, with I_n an independent set of order n , and with the age a free amalgamation class in this dual sense.

Thus the catalog of homogeneous graphs reduces to a few easily found “in nature” and

some graphs whose existence is provided most naturally by a generic construction in the sense of Fraïssé. The Fraïssé theory continues to play a role in the proof of Fact 2.1.2. In fact, it is not that easy to extract the statement of this fact from the presentation in [LW80], which gives the statement in terms of the minimal forbidden substructures, after the following reductions.

The finite homogeneous graphs were classified partially by Sheehan [She74], and fully and independently by Gardiner [Gar80]. In the infinite case, one may set aside the imprimitive case, where as we have already remarked, the graphs have the form $m \cdot K_n$ up to complementation, that is, a disjoint union of m n -cliques, where $2 \leq m, n \leq \infty$.

In the case of infinite graphs, an application of Ramsey's theorem guarantees that there is an infinite complete subgraph or an infinite independent set, and up to complementation one may suppose there is an infinite independent set. Therefore what actually remains to be proved to arrive at Fact 2.1.2 is the following: an infinite primitive homogeneous graph which contains an infinite independent set, and which also contains an n -clique K_n , must contain every finite graph which does not contain K_{n+1} . This is the content of [LW80]. The proof is elaborate and goes more deeply into the theory of amalgamation classes, in directions that will not be needed here.

We now sum up some characteristic features of the classification of homogeneous graphs which have some parallels in the known families of metrically homogeneous graphs.

A few exceptional types of homogeneous graphs occur "in nature" as special cases, such as the finite ones and some imprimitive graphs that appear as their natural infinite limits. The infinite primitive homogeneous graphs are most simply constructed in terms of Fraïssé theory, and we have observed that Fraïssé theory also provides an essential tool for the proof of their classification.

We notice further that the homogeneous graphs furnished directly by Fraïssé theory are of a particularly simple and uniform kind. Namely, we may fix the amalgamation procedure in advance: free amalgamation. Having done so, the only possible forbidden graphs are cliques, and we arrive at Henson's constraints.

While there is a highly developed theory of free amalgamation classes which we will make good use of, notably in connection with the automorphism extension property discussed in

Section 2.3, most of the amalgamation classes we deal with are not closed under free amalgamation. Rather they generally satisfy the following much broader condition.

Definition 2.1.9. *An amalgamation class has strong amalgamation if every amalgamation diagram $f_1: A_0 \rightarrow A_1$, $f_2: A_0 \rightarrow A_2$ in \mathcal{A} has an amalgam $g_1: A_1 \rightarrow A$, $g_2: A_2 \rightarrow A$ such that the images of A_1 and A_2 in A meet in the image of A_0 .*

There is a useful connection between the strong amalgamation property in the age of a structure and free amalgamation in the associated class of partial structures, which we will develop in Section 2.6.

Furthermore, strong amalgamation classes allow some further useful constructions, of particular relevance in connection with Ramsey theory as discussed in Section 2.5. The following is immediate, but useful.

Lemma 2.1.1. *Let \mathcal{A}, \mathcal{B} be strong amalgamation classes in disjoint relational languages $\mathcal{L}_1, \mathcal{L}_2$. Let $\mathcal{A} \star \mathcal{B}$ be the class of finite structures in the language $\mathcal{L}_1 \cup \mathcal{L}_2$ whose reducts to $\mathcal{L}_1, \mathcal{L}_2$ are in \mathcal{A} or \mathcal{B} respectively. Then $\mathcal{A} \star \mathcal{B}$ is a strong amalgamation class.*

The corresponding Fraïssé limits will be structures Γ_1, Γ_2 , and Γ , each homogeneous in the appropriate language, with the reducts of Γ to the languages $\mathcal{L}_1, \mathcal{L}_2$ isomorphic to Γ_1 and Γ_2 respectively.

A case of particular interest arises when \mathcal{B} is the class of finite linear orders. Then the resulting structure Γ is called the expansion of Γ_1 by a *generic linear ordering*. Most of the metrically homogeneous graphs we are about to encounter are associated with strong amalgamation classes and therefore have generic expansions by linear orderings, to be investigated in Section 3.3.

2.2 Metrically Homogeneous Graphs

We now discuss the general theory of (countable) metrically homogeneous graphs, the subject matter of the remainder of the thesis. These graphs are considered, notably, in [Mos92] and [Cam98], where the question of a possible classification was more or less explicitly raised. That classification is by no means complete, but a wide range of examples is known, and useful classification results are known in several special cases.

We will be concerned mostly with the known metrically homogeneous graphs in our work in Chapters 3 and 4, with the exception of the work in Sections 4.1, 4.2 where detailed information resulting from the existing classification theory, supplemented by some new information along the same general lines, allows us to treat the general case with some additional effort.

We will present the relevant results from the classification theory below. This includes a few concepts on which we rely throughout as well as some more technical material that will be used more sporadically in the body of the thesis. Of particular importance is the notion of *generic type* and the associated class of examples which arise via the Fraïssé theory of the previous section, depending mainly on five numerical parameters denoted δ, K_1, K_2, C, C' , which play a leading role in the various combinatorial analyses on which all of our results depend. As we will see, these parameters encode information about metric triangles which embed in a given metrically homogeneous graph.

Definition 2.2.1. *A connected graph Γ is metrically homogeneous if the metric space (Γ, d) is homogeneous in Fraïssé's sense (Definition 2.1.1) where d is the associated path metric.*

One could apply the notion more broadly to vertex-transitive graphs whose connected components are metrically homogeneous (which amounts to replacing d by d^* , which takes the value infinity when not already defined). We will not do so.

Note that the definition presupposes that we are treating metric spaces as relational structures. In general, this is done as follows.

Definition 2.2.1, cont. *Let $S \subseteq \mathbb{R}^{>0}$. A metric space with values in S will be viewed as a complete edge-labeled graph with labels in S , satisfying the triangle inequality. Such structures may then be viewed as relational structures carrying symmetric binary relations $(R_s \mid s \in S)$, with*

$$R_s(x, y) \iff d(x, y) = s.$$

It is often the case in examples of interest here that the set S is finite, and in all cases we consider, S will be at worst countable, so that the corresponding language $(R_s \mid s \in S)$ is countable. One subtlety that should be noted in passing is that an S -metric space need not realize every value in S —it may even consist of a single point, and realize no value in S . This case arises frequently in practice, as the Fraïssé theory leads us to consider arbitrary finite

subspaces of a given space.

Metrically homogeneous graphs fall within the broader class of homogeneous integer-valued metric spaces, and are *geodesic metric spaces* in the following sense, which is a discrete analog of the usual definition for \mathbb{R} -valued spaces.

Definition 2.2.2. *An integer-valued metric space is geodesic if every pair of points u, v is connected by a geodesic path ($u = u_0, \dots, u_r = v$) in which successive vertices lie at distance 1. (Thus $r = d(u, v)$.)*

In fact, metrically homogeneous graphs, viewed as metric spaces, and homogeneous integer-valued geodesic metric spaces are the same class of metric spaces. We now give several similar characterizations in this vein, all of them useful. Here we rephrase Proposition 5.1 of [Cam98], where the statement is given in terms of the first two of our four conditions, in the broader context of distance-transitive graphs.

Fact 2.2.1. *Let (Γ, d) be a homogeneous integer-valued metric space and let (Γ, E) be Γ viewed as a graph with edge relation $d(x, y) = 1$. Then the following are equivalent.*

1. (Γ, E) is a metrically homogeneous graph, and d is the graph metric.
2. (Γ, E) is connected.
3. (Γ, d) is a geodesic metric space.
4. (Γ, d) contains geodesics of arbitrary length up to the diameter of Γ .
5. (Γ, d) contains all triangles with edge lengths $(1, k, k + 1)$ with k less than the diameter of (Γ, d) .

The equivalence of items (3, 4) is a typical instance of homogeneity for pairs (i.e. distance transitivity).

Classification: General theory

Now we take up the classification theory for metrically homogeneous graphs. We begin with a general description of the present state of the theory, adopting terminological conventions

introduced in [Che17]. Also, we deal only with the case of metrically homogeneous graphs of diameter at least 3, as the case of smaller diameter falls under the case of homogeneous graphs discussed earlier, and involves some additional exceptional cases.

The theory begins by considering two mutually exclusive but not quite exhaustive possibilities, which we call *local exceptional type* and *generic type*. We begin with the former. In any metrically homogeneous graph Γ , the graph induced on the set of neighbors of a fixed vertex is a homogeneous graph, and by vertex transitivity, the isomorphism type of this graph is an invariant of Γ . By Fact 2.1.2, this graph is of known type, and up to complementation is either finite, imprimitive, generic K_n -free for some n , or fully generic. The first two cases are of a special character while the last two cases may result from general constructions relating to the Fraïssé theory. Accordingly, we make the following definition in the metrically homogeneous case.

Definition 2.2.3. *A metrically homogeneous graph Γ is of local exceptional type if the graph induced on the set of neighbors of a vertex is finite or imprimitive.*

The metrically homogeneous graphs of local exceptional type have been explicitly classified, as we will discuss a little further on in full detail.

Now we turn to the remaining metrically homogeneous graphs. By definition, the associated graph induced on the neighbors of a vertex is infinite and primitive. In the definition of *generic type*, we require somewhat more than this.

Definition 2.2.4. *A metrically homogeneous graph Γ is of generic type if it satisfies the following two conditions.*

- *The graph induced on the set of neighbors of a vertex is primitive.*
- *The graph induced on the set of common neighbors of a pair of vertices at distance 2 contains an infinite independent set.*

This is the most interesting class of metrically homogeneous graphs, and will be the main object of study in the present thesis. Much of the thesis concerns the *known* metrically homogeneous graphs in this class. Occasionally we are able to work more generally.

The classification of metrically homogeneous graphs which are not of generic type is complete. Those of local exceptional type are described below.

Fact 2.2.2. [[Che11](#), Lemmas 8.6, 8.12] *A metrically homogeneous graph Γ which is neither of exceptional local type nor of generic type is a regular tree with infinite branching.*

Moreover, we also deduce the following from [[Che17](#)], reproduced in [[ACM16](#)].

Fact 2.2.3. *Let Γ be an infinite metrically homogeneous graph of finite diameter. Then Γ is of generic type.*

As the statement given in [[Che11](#)] is phrased in a different way, we give the necessary reductions to the statement actually found there. By assumption, the graph Γ_1 induced on the set of neighbors of a fixed vertex v of Γ is infinite and primitive. By the classification of homogeneous graphs, it follows that Γ_1 is either an independent set, a Henson graph, the complement of a Henson graph, or a generic graph. By Lemma 8.12 of [[Che11](#)], Γ_1 contains an independent set, and hence must be an independent set, a Henson graph, or a generic graph. As Γ is not of generic type, the common neighbors of a pair of vertices at distance 2 in Γ_1 contains no infinite independent set, and so Γ_1 cannot be a Henson graph or a generic graph. Therefore Γ_1 is an infinite independent set, and any pair of vertices at distance 2 in Γ have infinitely many neighbors. By Lemma 8.6 of [[Che11](#)], Γ is a regular tree with infinite branching.

Classification: Exceptional local type

We now return to the exceptional local case, giving the complete classification in this case, based on [[Che11](#)]. This consists of two parts—the existence of the graphs in question, and the completeness of the resulting list. The statement is as follows.

Fact 2.2.4. [[Che11](#), Theorem 10] *Let Γ be a metrically homogeneous graph of exceptional local type of diameter $\delta \geq 3$. Then either Γ is finite, or Γ is a tree-like graph $T_{m,n}$ with $m, n \leq \infty$, as described below.*

In fact, Theorem 10 of [[Che11](#)] is given in a slightly more general form, including the case in which the graph induced on the neighbors of a fixed vertex is the complement of a Henson graph.

The graphs $T_{m,n}$ are introduced in [Mac82]. They may be defined in terms of the theory of blocks (maximal 2-connected subgraphs) as graphs in which all blocks are cliques of order $n + 1$ (i.e. infinite if $n = \infty$) and each vertex lies in m distinct blocks. (Recall that a graph is *2-connected* if there does not exist a vertex whose removal disconnects the graph.) A more explicit and more useful definition is the following. Let $T(m, n + 1)$ be a biregular tree with degrees m and $n + 1$; that is, every vertex has degree m or $n + 1$, and vertices of degree m are joined only to vertices of degree $n + 1$ (it is possible however that $m = n + 1$). Let V be the set of vertices of degree m and let $T_{m,n}$ be the graph with vertex set V and edge relation $d(x, y) = 2$, where d is the path metric. To put this in a broader context, we may put the construction in the following framework.

Lemma 2.2.1. *Let Γ be a connected bipartite graph which is homogeneous when considered as a structure in the language consisting of the path metric in Γ together with unary predicates A, B defining the sets in a bipartition of Γ . Let Γ_A be the graph on A with edge relation $d(x, y) = 2$, where d is the path metric in Γ . Then Γ_A is a metrically homogeneous graph.*

Proof. Since Γ is connected and bipartite, Γ_A is connected.

It is clear that Γ_A is homogeneous as a metric space with the metric d_Γ induced from Γ , which takes on only even values. Therefore Γ_A is also homogeneous as an integer-valued metric space with the scaled metric $(1/2)d_\Gamma$. Then Γ_A satisfies the conditions of Fact 2.2.1 (2), and Fact 2.2.1 (1) tells us that Γ_A is metrically homogeneous with path metric $(1/2)d_\Gamma$. \square

We apply this to the biregular tree $T(m, n + 1)$ in a language containing a bipartition by unary predicates. Then the metric homogeneity of $T_{m,n}$ follows from the homogeneity of $T(m, n + 1)$ as a metric space with a bipartition. For this we refer to Lemma 5.3 of [Che11].

Now to complete the classification of the metrically homogeneous graphs of exceptional local type, it suffices to treat the finite case. Here we will once more encounter some particular instances of more general constructions.

Fact 2.2.5. [Cam76] *The finite metrically homogeneous graphs of diameter at least 3 are as follows.*

- n -cycles with $n \geq 6$.

• *Antipodal double covers of finite homogeneous graphs of one of the following forms:*

1. *the 5-cycle C_5 ;*
2. *the graph $[3]^2$ described in Fact 2.1.2; or*
3. *an independent set I_n of order $n \geq 2$.*

The general definition of an antipodal double cover is given in [Che11, Definition 5.4]. When Γ is a finite graph of diameter 2 and order n which is not complete, it produces a graph of diameter 3 and order $2n + 2$ which will be metrically homogeneous if it is vertex transitive. The antipodal double cover of C_5 (order 12) is the set of vertices of an icosahedron with edge relation given by the 1-skeleton. The antipodal double cover of $[3]^2$ (order 20) is the Johnson graph $J(6, 3)$ with vertices the triples from a 6-element set and edge relation given by $|x \cap y| = 2$. The antipodal double cover of I_n is the bipartite complement of a perfect matching between two sets of order $n + 1$.

In these graphs, each vertex lies at distance 3 from a unique “antipodal” vertex. We will discuss the general notion of “antipodal graph” below in the context of metrically homogeneous graphs of generic type.

Classification: Generic type

As we have seen, the classification of the metrically homogeneous graphs of non-generic type is complete. We now take up the classification of metrically homogeneous graphs of generic type.

There is a uniform description of all known metrically homogeneous graphs of generic type, in terms of the class of *3-constrained* metrically homogeneous graphs. This class is defined using Fraïssé theory, as follows.

Definition 2.2.5. *Let \mathcal{L} be a finite relational language. A homogeneous \mathcal{L} -structure is 3-constrained if the minimal \mathcal{L} -structures not in its age have order at most 3.*

In other words, a homogeneous structure is 3-constrained if the associated amalgamation class is 3-constrained in the sense that it is determined by a set of forbidden structures (constraints) of order at most 3. A typical example of such a constraint would be the triangle inequality in metric spaces, which defines a set of forbidden structures of order 3.

The status of the classification theory is best understood in terms of the following conjecture. We consider \mathcal{L} -structures, where \mathcal{L} is the language of integer-valued metric spaces (i.e. the language of edge-labeled graphs with positive integer labels).

Conjecture 2.2.1. *If Γ is a metrically homogeneous graph of generic type and \mathcal{T}_Γ is the set of structures of order at most 3 which do not embed in Γ , then the class of \mathcal{T}_Γ -free structures is an amalgamation class.*

In other words, the conjecture states that every metrically homogeneous graph of generic type has the same forbidden triangles as some 3-constrained metrically homogeneous graph (also of generic type) with the same diameter.

Now we give the classification of 3-constrained metrically homogeneous graphs of generic type. This involves 5 numerical parameters $\delta, K_1, K_2, C_0, C_1$ which may be defined for any metrically homogeneous graphs as follows.

Definition 2.2.6. *Let Γ be an integer-valued metric space.*

The type of a triangle in Γ is the triple (i, j, k) of edge lengths which occur (taken in any order).

We define parameters $\delta, K_1, K_2, C_0, C_1$ as follows.

- δ is the diameter (possibly infinite).
- K_1 is the smallest number k such that there is a triangle of type $(k, k, 1)$ in Γ , and K_2 is the largest such number (possibly infinite). In the event that there are no such triangles realized in Γ , K_1 is set equal to ∞ and K_2 is set equal to 0.
- For $\epsilon = 0$ or 1 , C_ϵ is the smallest number of parity ϵ , greater than 2δ , such that all triangles in Γ with perimeter p having parity ϵ satisfy $p < C_\epsilon$.

Furthermore, we set $C = \min(C_0, C_1)$ and $C' = \max(C_0, C_1)$.

In many cases, $C' = C + 1$, and in such cases there is no real distinction on the basis of parity, as C becomes a uniform bound on the perimeter of all triangles in Γ .

Fact 2.2.6. [[Che17](#), Theorems 12.1 and 13.1] *Let Γ be a 3-constrained metrically homogeneous graph of generic type and diameter at least 3. Then*

Case (a) (bipartite case). $K_1 = \infty$:

- $K_2 = 0$ and $C_1 = 2\delta + 1$.

Case (b) (low C). K_1 finite, $C \leq 2\delta + K_1$.

- $C = 2K_1 + 2K_2 + 1 \geq 2\delta + 1$;
- $K_1 + 2K_2 \leq 2\delta - 1$;
- If $C' > C + 1$ then $K_1 = K_2$ and $3K_2 = 2\delta - 1$.

Case (c) (high C). $C > 2\delta + K_1$.

- $K_1 + 2K_2 \geq 2\delta - 1$ 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 $C \geq 2\delta + K_2$.

Table 2.1: Admissible parameter choices with $\delta \geq 3$ (Definition 2.2.10)

1. The set of triangles embedding in Γ is determined by the numerical parameters δ, K_1, K_2, C_0 , and C_1 .
2. These parameters satisfy one of the three sets of numerical conditions given in Table 2.1.

Conversely, every such parameter sequence is realized by some 3-constrained metrically homogeneous graph.

We will have occasion to refer to the precise conditions on the numerical parameters shown in Table 2.1. We call such sequences of numerical parameters *admissible*.

Fact 2.2.6 involves three issues: the construction of 3-constrained amalgamation classes associated with a given admissible sequence of parameters, the proof of admissibility, and the proof that the numerical parameters completely determine the set of forbidden triangles, in the 3-constrained context. We elaborate on the last of these points.

Definition 2.2.7. Let $\delta, K_1, K_2, C_0, C_1$ be numerical parameters, some of which may be infinite. The associated set $\mathcal{T}(\delta, K_1, K_2, C_0, C_1)$ of forbidden triangle types consists of all triangle types (i, j, k) which violate one or more of the following constraints.

- $i, j, k \leq \delta$;
- If the perimeter $p = i + j + k$ is odd, then $2K_1 < p < 2K_2 + 2 \min(i, j, k)$.

- If the perimeter $p \equiv \epsilon \pmod{2}$, ($\epsilon = 0$ or 1), then $p < C_\epsilon$.

We may reformulate the discussion up to this point in more explicit terms as follows.

- For an admissible sequence of parameters, constraints of the form $\mathcal{T}(\delta, K_1, K_2, C_0, C_1)$ defines an amalgamation class $\mathcal{A}_{K_1, K_2, C_0, C_1}^\delta$ of $[\delta]$ -valued metric spaces, whose Fraïssé limit $\Gamma_{K_1, K_2, C_0, C_1}^\delta$ is then a 3-constrained metrically homogeneous graph of generic type.
- Conversely, any countable 3-constrained metrically homogeneous graph of generic type has the form $\Gamma_{K_1, K_2, C_0, C_1}^\delta$ for some admissible sequence of numerical parameters.
- *Conjecturally*, if Γ is a metrically homogeneous graph of generic type and $\delta, K_1, K_2, C_0, C_1$ are the associated parameters, then both of the following should hold.
 - The sequence of numerical parameters is admissible.
 - The metric triangles not embedding in Γ are those which lie in the set $\mathcal{T}(\delta, K_1, K_2, C_0, C_1)$.

As we saw in connection with the classification of metrically homogeneous graphs, it is useful to separate the primitive and imprimitive cases of the classification problem. The fundamental result concerning this division is the following.

Fact 2.2.7. [[Smi71](#), Theorem 2]; [[Che17](#), Lemma 7.1] *Let Γ be an imprimitive metrically homogeneous graph of diameter at least 3 and degree at least 3. Then Γ is bipartite or antipodal (Definition 2.2.8 below).*

This is stated by D.H. Smith in the context of finite distance transitive graphs but holds in several other forms as well; for finite distance regular graphs [[AH06](#)], with a suitably modified definition of imprimitivity, and for infinite distance transitive graphs, hence also for metrically homogeneous graphs.

Definition 2.2.8. *A graph of finite diameter δ is antipodal if the reflexive closure of the relation $d(x, y) = \delta$ is an equivalence relation.*

In particular, an antipodal graph of diameter 2 is a complete multi-partite graph. As there are no antipodal graphs of infinite diameter, the imprimitive graphs of infinite diameter are bipartite.

Note that there is also a metric definition of bipartite: the definition “ $d(x, y)$ is even” is an equivalence relation. The definition of antipodality may be rephrased in a similar way: the relation “ $d(x, y)$ is divisible by δ ” is an equivalence relation. In fact, the proof of Smith’s theorem starts by showing that the equivalence relation in question has the form “ $d(x, y)$ is divisible by d ” for some parameter d , and then showing that d must be 2 or δ .

There are stronger forms of antipodality which are equivalent to the stated condition in the context of metrically homogeneous graphs, and which are sometimes taken as the definition in that setting. But the definition given above is the appropriate one in the more general setting of distance transitive graphs.

Fact 2.2.8. [[Che17](#), Lemma 6.1 and Theorem 11] *If Γ is a metrically homogeneous graph of diameter at least 3, then the following conditions are equivalent.*

1. Γ is antipodal.
2. For every vertex u in Γ , there is a unique vertex u' at distance δ .
3. There is an automorphism α of Γ satisfying the law

$$d(x, \alpha(y)) = \delta - d(x, y).$$

4. There is no triangle of perimeter greater than 2δ .
5. $C_1 = 2\delta + 1, C_0 = 2\delta + 2$.

Corollary 2.2.1. *Let Γ be an antipodal metrically homogeneous graph of diameter $\delta \geq 3$. Then $K_1 + K_2 = \delta$.*

The main point to be proved in Fact 2.2.8 is the derivation of the strong form of antipodality given in (2) from the usual definition of antipodality, in the metrically homogeneous context. Once one has this, one can define $\alpha(x)$ by the condition $d(x, \alpha(x)) = \delta$ and deduce the formally stronger condition (3). This is the subject of [[Che11](#), Theorem 11]. Condition (4) follows easily from (3) and implies (2); this is the content of [[Che11](#), Lemma 6.1].

The numerical reformulation of (4) as (5) is straightforward but convenient. Assuming the parameters are admissible, (5) reduces to $C = 2\delta + 1$. Similarly, the bipartite case corresponds numerically to $K_1 = \infty$.

There are as yet no general classification results for imprimitive metrically homogeneous graphs. There is a useful reduction of the bipartite case to primitive cases [Che17, Theorem 1.27]. In the case of antipodal graphs, there is a general reduction which applies in the context of distance transitive graphs but which is not helpful in the metrically homogeneous setting as the reduced graph, while still distance transitive, may not be metrically homogeneous.

We now discuss the full catalog of known metrically homogeneous graphs of generic type and their relationship with 3-constrained metrically homogeneous graphs. The main classification conjecture amounts to the statement that this relationship holds in general.

Just as the Henson graphs may be obtained by fixing an amalgamation procedure, namely free amalgamation, and then considering what additional constraints may be imposed which are compatible with that procedure, a broader class of metrically homogeneous graphs can be obtained by starting with one of the 3-constrained cases, considering the associated amalgamation procedure in detail, and then adding in additional constraints compatible with that procedure. This leads to a notion of *Henson constraint* appropriate to the setting of metrically homogeneous graphs of diameter at least 3. In fact, it leads to two such notions: one is appropriate in all cases except the antipodal case, while there is a special notion of *antipodal Henson constraint* arising as an exceptional case.

Under the assumption that the diameter δ is at least 3, it is frequently the case that the amalgamation procedure can be carried out so as to introduce no new pairs at distance 1 or at distance δ . Therefore it is natural to consider metric spaces with values in the set $\{1, \delta\}$ as additional constraints.

Definition 2.2.9. *A $(1, \delta)$ -space is a metric space with values in the set $\{1, \delta\}$.*

When $\delta \geq 3$, we will also refer to a $(1, \delta)$ -space as a Henson constraint, or more precisely, as an ordinary Henson constraint. We will subsequently introduce antipodal Henson constraints (Definition 2.2.11).

Observe that Henson constraints consist of cliques of various sizes lying at distance δ from each other. The extreme cases of a single clique or trivial cliques are permitted.

Definition 2.2.10. *Let $(\delta, K_1, K_2, C, C')$ be an admissible parameter sequence with $\delta \geq 3$ and $C > 2\delta + 1$.*

A set \mathcal{S} with ordinary Henson constraints is said to be admissible with respect to the parameter sequence if the following conditions are satisfied.

- No structure in \mathcal{S} contains any of the forbidden triangles in $\mathcal{T}(\delta, K_1, K_2, C, C')$.
- If $K_1 = \delta$ or $C = 2\delta + 2$, then \mathcal{S} is empty.

A sequence of parameters $(\delta, K_1, K_2, C, C', \mathcal{S})$ with $C > 2\delta + 1$ is admissible if the sequence of numerical parameters is admissible, and if \mathcal{S} is admissible with respect to that sequence.

It would also be reasonable to impose the requirement on \mathcal{S} that no constraint in \mathcal{S} contain any other constraint in \mathcal{S} (that is, \mathcal{S} is an antichain) but we do not require that here.

Now we pass to the consideration of the case of antipodal graphs, that is, $C = 2\delta + 1$, and the appropriate notion of Henson constraint in that context.

Definition 2.2.11. *Let $\delta \geq 3$.*

Let A be a $[\delta]$ -metric space. An antipodal variant of A is obtained by dividing A into two parts A_0, A_1 and replacing the distances $d(x, y)$ between the two parts by the values $\delta - d(x, y)$. In particular, an antipodal variant of an n -clique is a structure consisting of the union of two cliques K_{n_1}, K_{n_2} with $n_1 + n_2 = n$, and all distances between the two cliques equal to $\delta - 1$.

An antipodal Henson constraint is an antipodal variant of an n -clique.

For δ, K_1, K_2, C, C' an admissible sequence of numerical parameters with $C = 2\delta + 1$, a set \mathcal{S} of antipodal Henson constraints is admissible if the following conditions are satisfied.

- *No structure in \mathcal{S} contains any of the forbidden triangles in $\mathcal{T}(\delta, K_1, K_2, C, C')$.*
- *If $\delta = 3$, then \mathcal{S} is empty.*
- *If \mathcal{S} contains a constraint A then it contains all antipodal variants of A .*

This is strictly analogous to the preceding definition in spite of its exceptional character. The first two conditions are as before, and the third condition is necessary if one wishes to specify the constraints for an antipodal graph. Furthermore, in the antipodal case, the only $(1, \delta)$ -spaces which come into consideration are cliques. So we are more or less forced to consider the notion given above if there is to be a notion of Henson constraint in the antipodal context. As it happens, the situation with $\delta = 3$ is already too tight to allow for this.

Since admissible collections of antipodal Henson constraints consist of cliques together with their antipodal variants, there is no loss of generality in requiring \mathcal{S} to consist of a single clique and its antipodal variants. But we do not impose this requirement here.

The first condition on \mathcal{S} simplifies considerably in the antipodal context. If \mathcal{S} is empty, then $K_1 = 1$ and hence $K_2 = \delta - 1$.

Fact 2.2.9. [[Che11](#), Theorems 9 and 14] *Let $(\delta, K_1, K_2, C, C', \mathcal{S})$ be an admissible sequence of parameters; in particular, \mathcal{S} is a set of Henson constraints of the appropriate type. Then $\mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$ is an amalgamation class, and thus one may speak of the associated metrically homogeneous graph $\Gamma_{K_1, K_2, C, C', \mathcal{S}}^\delta$.*

This is given without proof in [[Che11](#)]. The proof is found in [[Che17](#)].

As we have noted, it is conjectured that the classification of metrically homogeneous graphs is complete.

Conjecture 2.2.2. [[Che17](#)] *Let Γ be a countable metrically homogeneous graph of generic type. Then Γ is isomorphic to the Fraïssé limit of an amalgamation class associated to an admissible parameter sequence $(\delta, K_1, K_2, C, C', \mathcal{S})$ where \mathcal{S} is, in particular, a set of Henson constraints of appropriate type—antipodal if $C = 2\delta + 1$, and ordinary if $C > 2\delta + 1$.*

We focus in this thesis on the known metrically homogeneous graphs, but occasionally make use of the classification theory to give results more generally in terms of the associated numerical parameters (and Henson constraints, when relevant).

The classification is also complete in diameter 3.

Fact 2.2.10. [[ACM16](#), Theorem 1] *The metrically homogeneous graphs of diameter 3 are all of known type, that is, either finite or of generic type, and in the latter case of the form $\Gamma_{K_1, K_2, C, C', \mathcal{S}}^\delta$ with admissible parameters.*

Another reduction of interest, though not used in the present thesis, is the reduction of the main conjecture in the case of infinite diameter to the case of finite diameter [[Che17](#), Theorem 1.23].

Other results of a more technical nature which have been developed for the uses of the classification theory will be applied in Chapter 4 in the classification of the twisted isomorphisms

and the associated metrically homogeneous graphs. These results will be presented below.

Another problem connected with the classification problem concerns the strong amalgamation property. This property is particularly important for a number of reasons, notably because it permits us to form generic linear orderings, which are useful in the context of Ramsey theory, and in fact will be one of the key assumptions in the treatment of Ramsey theory given in Section 2.5 below.

Observation 1. *A known metrically homogeneous graph of generic type is the Fraïssé limit of a strong amalgamation class if and only if it is not antipodal.*

Clearly an antipodal graph does not allow strong amalgamation of finite substructures. This is the case whenever we have an imprimitive structure carrying an equivalence relation with finite classes.

We have presented the known metrically homogeneous graphs above, and, in particular, the detailed catalog of known metrically homogeneous graphs of generic type, as well as the known results and conjectures regarding the classification problem. We now pass to the discussion of some more technical points to be used mainly in Chapter 4 when dealing with metrically homogeneous graphs which are not necessarily of known type. These results can all be verified by inspection in the case of the known metrically homogeneous graphs but can also be proved in general by direct, though sometimes lengthy, arguments.

The first result of this kind concerns the significance of the numerical parameter K_1 .

Fact 2.2.11. *[Che17, Lemma 13.15] Let Γ be a metrically homogeneous graph which is not bipartite. Let p be the least odd number which is the perimeter of some triangle in Γ . Then the following conditions hold.*

1. *A p -cycle embeds isometrically in Γ .*
2. *$p \leq 2\delta + 1$.*
3. *$p = 2K_1 + 1$.*

This result has the following consequences concerning the interpretation of the parameter K_1 .

Corollary 2.2.2. *Let Γ be a metrically homogeneous graph with $K_1 > 1$. Then the following hold.*

1. *Any triangle type of odd perimeter less than $2K_1$ is not realized in Γ .*
2. *Every triangle type of perimeter $2K_1 + 1$ is realized in Γ .*

Local analysis

One of the main topics in the classification theory is some knowledge of the *local structure* of a metrically homogeneous graph.

Definition 2.2.12. *If Γ is a graph, u a vertex, and $i \geq 0$, then $\Gamma_i(u)$ denotes the set of vertices at distance i from u with the induced metric structure, given by the path metric on Γ . In particular, the graph structure induced on Γ_i by Γ is given by the relation $d(x, y) = 1$.*

If Γ is vertex-transitive, then we may write Γ_i for $\Gamma_i(u)$, and this is well-defined up to isomorphism.

We note that metrically homogeneous graphs are vertex-transitive.

In general, one hopes to derive information about Γ from information about its local structure, that is, the structure of the parts Γ_i , particularly when working in an inductive setting. While Γ_i is homogeneous as a metric space with the induced structure, one difficulty is that it is possible that this metric structure is not closely related to the graph structure induced on Γ_i . In particular, Γ_i may be an independent set (for example, with $i = 1$ and $K_1 > 1$). The following shows that in all other cases the situation is satisfactory.

Fact 2.2.12. [[Che17](#), Theorem 1.29] *Let Γ be a metrically homogeneous graph of generic type. Suppose $i \leq \delta$, and Γ_i is not an independent set. Then Γ_i is a metrically homogeneous graph.*

Furthermore, Γ_i is primitive and of generic type except in the following two cases, in which $i = \delta$ or $i = \delta/2$.

1. *If $i = \delta$, $K_1 = 1$, $C_0 = 2\delta + 2$, $C_1 = C_0 + 1$, then Γ_δ is an infinite complete graph.*
2. *If δ is even, $i = \delta/2$, and Γ is antipodal, then Γ_i is also antipodal of diameter δ .*

Note that by definition, metrically homogeneous graphs are connected. In fact, the claim that if Γ_i contains at least one edge then it is connected is one of the main points in the proof of Fact 2.2.12.

Neither of the “exceptional” cases is particularly exceptional. If the graph Γ_δ is an infinite complete graph, then it is primitive, and while not of generic type in the strict sense, it is the generic metrically homogeneous graph of diameter 1. On the other hand, if δ is even and Γ is antipodal, then $\Gamma_{\delta/2}$ is again an antipodal graph of diameter δ . In fact, the conjecture in this case would be that $\Gamma_{\delta/2}$ is isomorphic to Γ .

This fact is very helpful whenever it applies, and conjecturally it should apply whenever $K_1 \leq i \leq K_2$, but this point has not been proved. However the following special case is known.

Fact 2.2.13. [Che17, Proposition 1.30] *Let Γ be a metrically homogeneous graph of generic type and diameter δ with*

$$K_1 \leq 2.$$

Then Γ_i contains an edge for $K_1 \leq i \leq \delta - 1$, except in the case in which Γ is antipodal and

$$K_1 = 1 \qquad K_2 = \delta - 2 \qquad i = \delta - 1$$

A technical variant of Fact 2.2.12 not only plays a major role in the proof of that fact, but is occasionally useful in its own right.

Fact 2.2.14. [Che17, Lemma 15.4] *Let Γ be a metrically homogeneous graph of generic type and diameter δ . Suppose that $i \leq \delta$, and if $i = \delta$, suppose also that $K_1 > 1$.*

Then the metric space Γ_i is connected with respect to the edge relation defined by

$$d(x, y) = 2.$$

The following is new.

Lemma 2.2.2. *Let Γ be a metrically homogeneous graph of diameter δ and let δ' be the diameter of Γ_δ . Then for $i \leq (\delta - \delta')/2$, the diameter of $\Gamma_{\delta-i}$ is $\delta' + 2i$.*

Proof. This is proved by induction on i , with the base case $i = 0$ holding by definition.

Suppose $i > 0$, $\delta' + 2i \leq \delta$, and $\Gamma_{\delta-j}$ has diameter $\delta' + 2j$ for $j < i$. We show first that the diameter of $\Gamma_{\delta-i}$ is at least $\delta' + 2i$.

Take $u, v \in \Gamma_{\delta-(i-1)}$ at distance $\delta' + 2i - 2$ and u', v' adjacent to u, v respectively at distance $d(u, v) + 2$. It suffices to show that u', v' are both in $\Gamma_{\delta-i}$ to conclude. Note that

$$u', v' \in \Gamma_{\delta-(i-2)}, \Gamma_{\delta-(i-1)}, \text{ or } \Gamma_{\delta-i}.$$

Assume towards a contradiction that we do not have both u', v' in Γ_{i-1} . Our inductive diameter bounds imply that we do not have $u', v' \in \Gamma_{\delta-(i-1)}$, nor do we have $u', v' \in \Gamma_{\delta-(i-2)}$.

Suppose next that $u' \in \Gamma_{\delta-(i-1)}, v' \in \Gamma_{\delta-(i-2)}$. In that case, we take $v'' \in \Gamma_{\delta-(i-1)}$ adjacent to v' and then $d(u', v'')$ violates the diameter bound of $\Gamma_{\delta-(i-2)}$. Thus we may assume that $u' \in \Gamma_{\delta-i}$.

This leaves us with $v' \in \Gamma_{\delta-j}$ with $j = i - 1$ or $i - 2$. We may take $u'' \in \Gamma_{\delta-j}$ with $d(u', u'') = i - j$, and find

$$d(u'', v') \geq d(u', v') - d(u', u'') = d(u, v) + 2 - (i - j) = \delta' + i + j > \delta' + 2j,$$

which is a violation of the diameter bound on $\Gamma_{\delta-j}$. Thus we have a contradiction. Hence, $u', v' \in \Gamma_{\delta-i}$.

This shows that the diameter of $\Gamma_{\delta-i}$ is at least $\delta' + 2i$, and the reverse inequality follows similarly: if $u, v \in \Gamma_{\delta-i}$ and u', v' are adjacent to u, v respectively and lie in $\Gamma_{\delta-i+1}$, induction and the triangle inequality give $d(u, v) \leq \delta' + 2i$. \square

Fact 2.2.15. [[Che17](#), Lemma 15.6] *Let Γ be a metrically homogeneous graph of generic type and diameter δ which contains a triangle of perimeter $2\delta + d$. Then Γ_δ has diameter at least d .*

In terms of triangle types, the conclusion is that Γ contains a triangle of type (δ, δ, d') with $d' \geq d$.

Fact 2.2.16. [[Che17](#), Lemma 15.5] *Let Γ be a metrically homogeneous graph of generic type and diameter δ . Suppose $1 \leq i \leq \delta$. Then for $u \in \Gamma_{i\pm 1}$, the set $\Gamma_1(u) \cap \Gamma_i$ is infinite.*

The next result shows that in many cases, the parameters C, C' can be directly computed from the diameter of Γ_δ .

Lemma 2.2.3. *Let Γ be a metrically homogeneous graph of generic type and finite diameter δ and let δ' be the diameter of Γ_δ . Then*

$$C' = 2\delta + \delta' + 2.$$

Let D_δ denote the set of distances realized in Γ_δ by pairs of distinct points. Then one of the following occurs.

1. $C = 2\delta + \delta' + 1$.
2. $C = C_1 = 2\delta + 1$, δ' is even, and D_δ consists of the even numbers in the interval $[2, \delta']$.
3. $C = C_1 = 2\delta + d$ with d odd and $\min D_\delta = 2 < d < \delta'$.

In particular, D_δ is not an interval.

Proof. This follows fairly directly from Facts 2.2.12–2.2.15 and one further estimate: $\min D_\delta \leq$

2. So we begin with the last point.

Take $w \in \Gamma_{\delta-1}$. By Fact 2.2.16, there are distinct neighbors u, v of w in Γ_δ . Thus

$$\min D_\delta \leq d(u, v) \leq 2.$$

By definition, there are triangles of perimeter $C_\epsilon - 2$, for $\epsilon = 0, 1$. So by Fact 2.2.15 we have the following:

$$C_\epsilon - 2 \leq 2\delta + \delta' \quad (\epsilon = 0, 1).$$

Thus

$$C \leq 2\delta + \delta' + 1$$

$$C' \leq 2\delta + \delta' + 2.$$

Thus if $C = 2\delta + \delta' + 1$, then $C' = 2\delta + 2\delta' + 2$ and the claim follows.

We may suppose then that $C = 2\delta + d$ for some $d \leq \delta'$. Now Γ contains triangles of type (δ, δ, i) and thus of perimeter $2\delta + i$ whenever $i \in D_\delta$. So

$$d \notin D_\delta.$$

If $1 \in D_\delta$ then by Fact 2.2.12 we have $D_\delta = [1, \delta']$ and we have a contradiction. So

$$\min D_\delta = 2.$$

Case 1. $d = 1$.

Then $C = C_1 = 2\delta + 1$, and there are no triangles of larger odd perimeter, implying that D_δ consists of even integers in the interval $[2, \delta']$. Thus we claim that all of these distances are

realized in Γ_δ . If Γ_δ is connected with respect to the distance 2 relation, then this follows at once. In particular, if $K_1 > 1$ this holds by Fact 2.2.14.

Suppose on the contrary that Γ_δ has at least two connected components with respect to the relation $d(x, y) = 2$, and in particular $K_1 = 1$. Let $v_1, v_2 \in \Gamma_\delta$ lie in distinct components with respect to the distance 2 relation, and take $d_{1,2} = d(v_1, v_2)$ to be minimal. Then $d_{1,2} > 2$ and $d_{1,2}$ is even. Let v'_1, v'_2 be neighbors of v_1, v_2 with $d(v'_1, v'_2) = d_{1,2} - 2$. Then v'_1, v'_2 lie in $\Gamma_{\delta-1}$.

By Fact 2.2.13, $\Gamma_{\delta-1}$ contains an edge and by Fact 2.2.12, $\Gamma_{\delta-1}$ is connected, with the induced metric given by the path metric in $\Gamma_{\delta-1}$. So there is a geodesic in $\Gamma_{\delta-1}$ between v'_1 and v'_2 , and in particular there is a neighbor v''_1 of v'_1 in $\Gamma_{\delta-1}$ with

$$d(v_1, v''_1) = 2 \qquad d(v_2, v''_1) = d_{1,2} - 2.$$

Take a neighbor u of v''_1 in Γ_δ . By the triangle inequality,

$$d(v_1, u) \leq 3 \qquad d(v_2, u) < d_{1,2}.$$

It follows that $d(v_1, u), d(v_2, u) < d_{1,2}$ and thus u lies in the same connected component with respect to the distance 2 relation as both v_1 and v_2 , which gives a contradiction.

Thus Γ_δ is connected with respect to the relation $d(x, y) = 2$ and the claims follow in this case.

Case 2. $d > 1$.

Since d is not in D_δ , we have

$$\min D_\delta < d < \delta'$$

and D_δ is not an interval. As $d < \delta'$ and there are no triangles in Γ of perimeter $2\delta + \delta'$, it follows that C and δ' have opposite parity. Thus C' and δ' have the same parity and in particular $C \geq 2\delta + 2\delta' + 2$. Thus again $C' = 2\delta + \delta' + 2$ in this case.

This completes the proof in all cases. \square

Our final goal is a new result on realizations of triangles, which will be use in Chapter 4. It is conjectured that any triangle of even perimeter at most 2δ will be realized in any metrically homogeneous graph of generic type. The following is a special case of this.

Proposition 2.2.1. *Let Γ be a metrically homogeneous graph of generic type and diameter δ . Suppose that*

$$k \leq i \qquad i + k \leq \delta$$

Then Γ contains a triangle of type $(i, i, 2k)$.

The bulk of the proof is contained in the following technical result.

Lemma 2.2.4. *Let Γ be a metrically homogeneous graph of generic type and diameter δ . Suppose*

$$k \leq \delta/2.$$

Let u_1, u_2 be vertices at distance $2k$ in Γ . Then in $\Gamma_k(u_1) \cap \Gamma_k(u_2)$ there is a geodesic path (v_0, \dots, v_k) with successive vertices v_i, v_{i+1} at distance 2.

Proof. This is proved by induction on k , for $1 \leq k \leq \delta/2$. By metric homogeneity, if such a path exists for a given pair of vertices at distance $2k$, then the same applies to all pairs at distance $2k$.

For $k = 1$, the condition is a direct consequence of the definition of generic type.

Now suppose $2 \leq k \leq \delta/2$ and the result holds for $j < k$. Take an arbitrary basepoint v_* in Γ . Since Γ is connected, of diameter at least $2k$, and metrically homogeneous, there are points v_0, v_k in Γ_k which lie at distance $2k$, with v_* as the midpoint of a geodesic path between them.

Then there are points $a, b \in \Gamma_{k-1}$ adjacent to v_0, v_k respectively. By the triangle inequality, this forces

$$d(a, b) = 2(k - 1).$$

By the induction hypothesis and homogeneity applied to the triple (v_*, a, b) , there is a geodesic path $(a = w_0, w_1, \dots, w_{k-1} = b)$ in Γ_{k-1} with successive points w_i, w_{i+1} at distance 2.

We claim that for $1 \leq i \leq k - 1$, the vertices w_{i-1} and w_i have a common neighbor v_i in Γ_k . By homogeneity, it suffices to find some vertex $v \in \Gamma_k$ having two neighbors in Γ_{k-1} at distance 2. For this, begin with $v \in \Gamma_k$ and take $v' \in \Gamma_{k-2}$ at distance 2 from v . Applying the hypothesis of generic type to the pair v, v' , we find infinitely many common neighbors of v, v' which are pairwise at distance 2, and in particular we have two such.

At this point, we have constructed a path $P_k = (v_0, v_1, \dots, v_k)$ in Γ_i such that each successive pair v_i, v_{i+1} has the common neighbor w_i , and such that the distance $d(v_0, v_k)$ is $2k$. It follows from the triangle inequality that this is a geodesic path. To complete our analysis, it remains to find a vertex $u \in \Gamma_{2k}$ such that the path P_k lies in $\Gamma_k(u)$.

The first part of the argument involved building up successively longer paths in $\Gamma_1, \dots, \Gamma_k$. Now we will build successively longer paths $P_{k+\ell}$ for $0 \leq \ell \leq k$ with the following properties.

- $P_{k+\ell}$ lies in $\Gamma_{k+\ell}$.
- $P_{k+\ell}$ is a geodesic path on $k + 1 - \ell$ vertices with successive vertices at distance 2

The inductive construction of the paths $P_{k+\ell}$ goes much as in the first part of the proof, using the hypothesis that Γ is of generic type. Given $P_{k+\ell}$, we choose the vertices of $P_{k+\ell+1}$ to be adjacent to successive pairs of vertices of $P_{k+\ell}$. If $P_{k+\ell}$ is a geodesic path on $k + 1 - \ell$ vertices of total length $2(k - \ell)$, then this construction produces a path $P_{k+\ell+1}$ on $k - \ell$ vertices whose endpoints lie at distance at least $2(k - \ell - 1)$ and with successive vertices at distance at most 2. So again this must be a geodesic path and all conditions are fulfilled.

In particular, the path P_{2k} consists of a single vertex u . As every vertex of one of the paths $P_{k+\ell}$ has neighbors in the next path $P_{k+\ell+1}$, the distance from any vertex of P_k to u is at most k . On the other hand, P_k is contained in Γ_k and u is in Γ_{2k} , so this distance is at least k . Therefore $P_k \subseteq \Gamma_k(u)$, as required. \square

Proof of Proposition 2.2.1. We have $k \leq i, i + k \leq \delta$. In particular Γ_{i-k} and Γ_{i+k} are both non-empty.

As $2k \leq i + k \leq \delta$, there is a geodesic triangle of type $(i - k, 2k, i + k)$ in Γ . Therefore we may fix $a \in \Gamma_{i-k}$ and $b \in \Gamma_{i+k}$ with $d(a, b) = 2k$.

Applying the previous lemma with basepoint a and with b playing the role of u there, by homogeneity we may find a geodesic path (v_0, \dots, v_k) in $\Gamma_k(a)$ with

$$d(v_\ell, v_{\ell+1}) = 2 \ (\ell < k) \text{ and } d(b, v_\ell) = k \ (\ell \leq k)$$

As $d(a, v_\ell) = d(b, v_\ell) = k$ for $\ell \leq k$, and $a \in \Gamma_{i-k}, b \in \Gamma_{i+k}$, it follows that $v_\ell \in \Gamma_i$. In particular the vertices v_0, v_k and the basepoint of Γ form a triangle of type $(i, i, 2k)$. \square

2.3 Topological Properties of Automorphism Groups

Chapter 3 of this thesis will deal with the properties of automorphism groups of metrically homogeneous structures viewed as topological groups. As we shall see below, the automorphism group of any countable structure can be viewed in a natural way as a *Polish group*, that is, as a separable topological group carrying a complete metric which is compatible with the topology. There is an extensive theory relating properties of an amalgamation class of structures with the properties of the automorphism group of its Fraïssé limit, viewed as a topological group. In fact, there are two such theories. The first of these, which centers on the notion of *ample generics*, will be discussed in this section, generally following the presentation of [Kec13]. When applicable, this shows that in various senses, the algebraic structure of the group completely determines the associated topology, and has other remarkable consequences for the interplay of the algebra and topology. The second theory concerns more specifically the *topological dynamical* properties of the automorphism group, and will be discussed in the next section.

Definition 2.3.1. $\text{Sym}(\mathbb{N})$ is the full (infinitary) symmetric group on \mathbb{N} , with the topology inherited from the product topology on $\mathbb{N}^{\mathbb{N}}$. Two metrics compatible with this topology may be defined as follows.

For $g, h \in \text{Sym}(\mathbb{N})$, define $|g|$ by

$$|g| := \begin{cases} 2^{-\min(n : g(n) \neq n)} & \text{if } g \neq 1 \\ 0 & \text{if } g = 1 \end{cases}$$

Note that $|g| = |g^{-1}|$ and $|gh| \leq \max(|g|, |h|)$ for $g \in \text{Sym}(\mathbb{N})$.

We define the metrics d_0 and d by

$$d_0(g, h) = |g^{-1}h|$$

$$d(g, h) = \min(d(g, h), d(g^{-1}, h^{-1})).$$

Then both d_0 and d define the topology on $\text{Sym}(\mathbb{N})$.

Any d_0 -Cauchy sequence (g_n) in $\text{Sym}(\mathbb{N})$ has a limit g which defines an injective map from \mathbb{N} to \mathbb{N} , and conversely any injective map is the limit of a sequence in $\text{Sym}(\mathbb{N})$. Thus d_0 is not complete.

If (g_n) is a d -Cauchy sequence in $\text{Sym}(\mathbb{N})$, then both (g_n) and (g_n^{-1}) are d_0 -Cauchy, hence there are limits $g = \lim g_n$ and $h = \lim g_n^{-1}$ which are injective maps, and furthermore $h = g^{-1}$, so $g \in \text{Sym}(\mathbb{N})$. Thus $\text{Sym}(\mathbb{N})$ is d -complete and is, in particular, a Polish group. Therefore any closed subgroup of $\text{Sym}(\mathbb{N})$ is a Polish group under the induced topology.

These definitions may be transferred to any countable set Ω , and to the group $\text{Sym}(\Omega)$, by choosing an enumeration of Ω .

It turns out that the closed subgroups of $\text{Sym}(\Omega)$ are exactly the automorphism groups of countable structures with domain Ω . We give this in detail, following [Kec13].

Fact 2.3.1. [Kec13, Theorems 1.1, 1.3]

1. If Γ is a countable structure, then $\text{Aut}(\Gamma)$ is a closed subgroup of $\text{Sym}(\Gamma)$, and in particular $\text{Aut}(\Gamma)$ is a Polish group in the induced topology.
2. Conversely, let G be a closed subgroup of $\text{Sym}(\Omega)$ and let Γ be the canonical structure on Ω associated to G , in which each G -orbit on Ω^m is a distinguished relation, for all n (cf. Definition 2.1.5). Then $G = \text{Aut}(\Gamma)$.

A more abstract characterization of this class of groups is the following.

Definition 2.3.2. A Polish group is non-archimedean if it has a neighborhood basis at the identity consisting of open subgroups.

Fact 2.3.2. [Kec13, Theorem 1.1] A Polish group is topologically isomorphic to a closed subgroup of $\text{Sym}(\mathbb{N})$ if and only if it is non-archimedean.

Thus one way to view our subject is that we study non-archimedean Polish groups by representing them as automorphism groups of countable structures, and then studying the underlying structure. Accordingly, we may view these groups at three successively more concrete levels of generality:

1. As abstract (discrete) groups;
2. As Polish groups;
3. As automorphism groups (closed permutation groups).

It is very natural to consider the question as to when the structure at a more abstract level determines the structure at a more concrete level. There is particular interest in the case in which the abstract group structure determines the topology as a Polish group uniquely.

One way to express is via the notion of *automatic continuity*. There are several closely related notions going under this name in the literature. We take the following definition.

Definition 2.3.3. *A Polish group G has the property of automatic continuity if every homomorphism of abstract groups from G into a Polish group G^* is necessarily continuous.*

This specializes to the following key property, taking $G^* = G$ and the homomorphism to be the identity: G carries a unique topology for which G is a Polish group.

One way in which the latter property may arise is as follows.

Definition 2.3.4. *A Polish group has the small index property (SIP) if and only if its open subgroups are precisely those of index less than 2^{\aleph_0} .*

This implies the more natural condition that the open subgroups are precisely those of *countable index*, but the methods used to prove this generally prove the stronger version as well. Either of these conditions implies that the topology is determined by the abstract group structure. As we will see in Theorem 1, Chapter 3, these two properties hold for a broad class of metrically homogeneous graphs, and thus in these groups the abstract group structure determines the topology.

The remaining points in Theorem 1 depend heavily on the topology of G but concern some remarkable properties of G as an abstract group. These are the properties of *uncountable cofinality* and the fixed point properties (FA) and (FH).

Definition 2.3.5. *The group G has uncountable cofinality if it is not the union of a countable chain of proper subgroups.*

Equivalently: if a countable union of subsets of G generates G , then one of the subsets generates G .

Definition 2.3.6. *Let G be a group.*

1. *G has property (FA) if every action of G on a tree with no inversions of edges fixes a point [Ser74, Ser77].*

2. G has property (FH) if every action of G by affine isometries on a Hilbert space fixes a point.

Here the symbols A and H abbreviate the terms *arbre* (tree) and Hilbert.

One can also formulate a strong form of uncountable cofinality, due to Bergman, which implies properties (FH) and (FA), and which can sometimes be proved by the same methods used to prove uncountable cofinality.

Definition 2.3.7. *The group G has the Bergman property if and only if for any countable sequence of sets X_n whose union generates G , there is an n and a k such that*

$$X_n^k = G.$$

Equivalently, G has uncountable cofinality, and for any set X which generates G , there is some k so that $X^k = G$ (i.e. X boundedly generates G).

Fact 2.3.3. [KR07] *The Bergman property implies the fixed point properties (FA) and (FH).*

This is explained in [KR07, discussion of Theorem 1.8]. We summarize the key points; full references are found in [KR07]. In the first place, it is known that the Bergman property implies that any action of the group G by isometries on a metric space has bounded orbits. This applies in particular to actions of G on trees, viewed as metric spaces, or on Hilbert space by affine isometries. In the case of an action on a tree, it follows that the action fixes a vertex or an edge, and in the case of an action on Hilbert space, it may be shown that a point is fixed.

Ample generics and their applications

The theory of *generic sequences* (“ample generics”) is a powerful tool for deriving topological and algebraic properties of Polish groups. We will apply this theory to the Polish group $\text{Aut}(\Gamma)$ for Γ a suitable metrically homogeneous graph. Our presentation generally follows [KR07].

Definition 2.3.8. *Let G be a Polish group.*

A sequence (g_1, \dots, g_n) of elements of G is n -generic if its orbit under the action of G by conjugation on the Cartesian power G^n is co-meager in G^n .

G has generic sequences if there are n -generic sequences of elements of G for all n .

In the special case in which $G = \text{Aut}(\Gamma)$ is the automorphism group of a countable structure Γ , if the group $\text{Aut}(\Gamma)$ has generic sequences of elements, we say that the structure Γ has *generic sequences of automorphisms* or simply that it has *ample generics*.

We will first discuss some consequences of the existence of ample generics for the topological and algebraic properties of certain classes of Polish groups, and then turn to the combinatorial conditions used to prove the existence of ample generics in automorphism groups.

Theorem 2.3.1. *[KR07, Theorems 1.6, 1.10] Suppose that G is a Polish group with generic sequences of elements. Then the following hold.*

- G has the small index property.
- G has automatic continuity.

Theorems 1.6 and 1.10 of [KR07] prove the small index property and a stronger form of automatic continuity than the one we have given in Definition 2.3.3.

Theorem 2.3.2. *[KR07, Theorem 1.7, Corollary 1.9] Suppose that G is an oligomorphic closed subgroup of $\text{Sym}(\Omega)$, with Ω countable, and that G has generic sequences of elements. Then the following hold.*

- G has uncountable cofinality.
- G has the fixed point properties (FA) and (FH).

Uncountable cofinality in the case of oligomorphic groups is derived from the more general Theorem 1.7 of [KR07] by a further application of a result of Peter Cameron, as explained in [KR07]. The fixed point properties are derived from the Bergman property, proved in [KR07, Theorem 1.8], by an application of Fact 2.3.3.

2.3.1 Establishing the existence of ample generics

Now we turn to the combinatorial side of the theory in the context of automorphism groups of homogeneous structures. The goal here is to reduce the problem of constructing generic sequences of automorphisms for a homogeneous structure Γ to combinatorial issues involving the age of Γ .

The usual approach to this is via the so-called *extension property for partial automorphisms* (EPPA), defined below. Kechris and Rosendal have given an exact reduction of the topological problem of existence of ample generics to combinatorial properties of the corresponding class of finite structures [KR07]. We will make use of the simpler, though less general, criterion of [Sin17], in which the property EPPA plays a central role.

Following [Sin17] (and, ultimately, [HL99]), we introduce the *extension property for partial automorphisms* (EPPA) and the *amalgamation property with automorphisms* (APA).

Definition 2.3.9. [Sin17, Definition 1.5.1] *Let \mathcal{A} be a class of finite structures in a relational language. Then \mathcal{A} has the extension property for partial automorphisms (EPPA) if for all A in \mathcal{A} there is a structure $\hat{A} \in \mathcal{A}$ so that every partial automorphism of A extends to an automorphism of \hat{A} .*

Definition 2.3.10. [Sin17, Def. 2.1.1] *Let \mathcal{A} be a class of finite structures in a relational language. Then \mathcal{A} has the amalgamation property with automorphisms (APA) if for every embedding $A_0 \rightarrow A_1, A_2$ in \mathcal{A} , there is an amalgam A in \mathcal{A} with the following property: for any automorphisms h_1, h_2 of A_1, A_2 respectively which leave the image of A_0 invariant and induce the same automorphism on A_0 , there is an automorphism h of A inducing h_1, h_2 on A_1, A_2 respectively.*

Theorem 2.3.3. [Sin17, Theorem 2.1.5] *Suppose that Γ is a homogeneous structure such that $\text{Age}(\Gamma)$ has both the EPPA and the APA. Then Γ has ample generics.*

In view of Theorem 2.3.3, Theorem 1 will follow once we establish the properties APA and EPPA of the relevant amalgamation classes, after which one deduces the existence of ample generics and one may then apply the general theory reviewed above.

Proving the APA and the EPPA

Now we arrive at the combinatorial problem of establishing the two properties APA and EPPA in the cases of interest to us. We will make some further reductions of the problem of a general character in Section 2.6, which involve a close study of the partial substructures associated with an amalgamation class.

The treatment of the APA and the EPPA are based on the notion of a *canonical completion process* (Section 2.6, Definition 2.6.4). In Section 2.6, Lemma 2.6.4 it will be shown how to derive the APA from a suitable canonical completion process for partial structures. In addition, in Lemma 2.6.5 it will be shown that the EPPA can also be derived from a suitable canonical completion process if the EPPA is known to hold in a slightly weaker sense at the level of partial substructures. The rest of this section is devoted to a purely combinatorial result of Herwig and Lascar [HL99, Theorem 3.2] which gives criteria for the required weak form of the EPPA to hold.

This will require some additional terminology. Since the terminology in this subject is extremely variable at this point, we first settle the purely terminological issues. The most important of these issues is that the term “EPPA” in [HL99] has a different meaning from that used in [Sin17] and elsewhere. We will largely follow Siniora’s terminology but indicate the variations in terminology found in the literature.

We will use the following terminology in the present thesis.

Definition 2.3.9(*cf. Definition 2.3.9*). *Let \mathcal{L} be a relational language. A class \mathcal{C} of \mathcal{L} -structures has the weak EPPA if for all finite $A \in \mathcal{C}$, and any sequence of partial automorphisms of A , whenever there is an extension C of A in \mathcal{C} for which the given sequence of partial automorphisms extends to an automorphism of C , then there is also a finite extension \tilde{A} of A in \mathcal{C} for which the same sequence of partial automorphisms extends.*

In one crucial case, the weak EPPA implies the full EPPA.

Example 2.3.1. *Let \mathcal{A} be an amalgamation class, Γ its Fraïssé limit, and $\mathcal{C} = \mathcal{A} \cup \{\Gamma\}$. Then the following are equivalent.*

- *The EPPA for \mathcal{A} .*
- *The weak EPPA for \mathcal{C} .*

Indeed, applying the weak EPPA to the sequence of all partial automorphisms of A , and taking C to be Γ in the definition of the weak EPPA, the full EPPA follows.

As mentioned above, terminology in this area is remarkably variable. Since we are interested in quoting from [HL99], [KR07], and [Sin17], and since [Her98] is an important precursor

of [HL99], we take a moment to pin down the variations in terminology. We present the main points in a table showing our terminology in parallel with terminology used elsewhere.

<i>Thesis</i>	[Her98]	[HL99]	[Sin17]	[KR07]
EPPA	EP	EP	EPPA	HP
weak EPPA	WEP	EPPA	(see below)	—

Here HP abbreviates “Hrushovski property.”

Our terminology follows [Sin17] fairly closely, but he uses the term “weak EPPA” for another condition slightly weaker than the one given in [HL99].

There are also some variations in related terminology which we will come to shortly: there are three notions of *forbidden structure* in common use, all of which are used at one point or another in this thesis. But we will return to this point below.

The only serious issue posed by all of these variations is that the term EPPA is used on different occasions for the weak or strong form of the condition. In working with amalgamation classes, the difference is negligible: Example 2.3.1 above shows that modulo a slight shift in notation, they are essentially equivalent in that context.

However, we will not always be working with amalgamation classes. We rely on the following combinatorial result of [HL99], which is not limited to amalgamation classes.

Theorem 2.3.4. [HL99, Theorem 3.2] *Let \mathcal{L} be a finite relational language and \mathcal{T} a finite set of \mathcal{L} -structures. Then the class of homomorphically \mathcal{T} -free \mathcal{L} -structures has the weak EPPA.*

Here of course Herwig and Lascar write “EPPA” rather than “weak EPPA.” Furthermore, they use the term “ \mathcal{T} -free” where we write “homomorphically \mathcal{T} -free.” This is defined as follows.

Definition 2.3.11. *Let \mathcal{L} be a relational language, and let T and A be \mathcal{L} -structures.*

Then \mathcal{L} is homomorphically \mathcal{T} -free if there is no homomorphism from T into A . In other words, there is no quotient of T which is isomorphic to a partial substructure of A .

For \mathcal{T} a class of \mathcal{L} -structures, A is said to be homomorphically \mathcal{T} -free if it is \mathcal{T} -free for each T in \mathcal{T} .

Siniora, who also uses all three commonly considered notions of \mathcal{T} -freeness, refers to this one as *free with respect to homomorphisms* [Sin17, Section 1.5]. The other two notions of freeness commonly encountered refer to forbidden induced substructures (in the context of Fraïssé theory) and to forbidden partial substructures, a topic we take up in Section 2.6.

2.4 Dynamical Properties of Automorphism Groups

Definition 2.4.1. *Let G be a topological group.*

1. *A G -flow is a continuous action of G on a compact space X .*
2. *A subflow of X is a compact invariant subset with the restriction of the action.*
3. *A flow is minimal if it has no proper subflows or equivalently every orbit is dense.*
4. *A homomorphism between two G -flows X, Y is a continuous G -map $\pi: X \rightarrow Y$, that is, $\pi(g \cdot x) = g \cdot \pi(x)$, $\forall g \in G, x \in X$. An isomorphism is a bijective homomorphism.*

The existence of minimal subflows follows from Zorn's lemma. It is also true that there is a unique (up to isomorphism) *universal minimal flow*:

Theorem 2.4.1. [Ell60, Theorem 2] *For any topological group G , there is a minimal G -flow, $M(G)$, with the following property.*

For any minimal G -flow X , there is a homomorphism $\pi: M(G) \rightarrow X$. Moreover, $M(G)$ is the unique (up to isomorphism) minimal G -flow with this property.

We define $M(G)$ to be the *universal minimal flow* of G .

Definition 2.4.2. *A topological group G is extremely amenable if its universal minimal flow $M(G)$ is trivial. In other words, every G -flow has a fixed point.*

Extreme amenability strengthens the older notion of *amenability*, which requires that every G -flow admits an invariant Borel probability measure. Extreme amenability was viewed initially as a pathological condition. In the classical setting of locally compact groups, if G is non-trivial, then $M(G)$ is non-trivial [Vee77], and apart from the compact case, where $M(G) = G$, the universal minimal flow is not even metrizable [KPT05, Theorem A2.2]. But it turns out that many automorphism groups of homogeneous structures have metrizable universal flows,

and that extreme amenability is closely connected with the subject of structural Ramsey theory. This is the subject of [KPT05], and the issue has been explored systematically by Ben Yaakov, Melleray, Tsankov, and Van Thé [MTT16, YMT17]; see the review by Lupini [Lup17].

We will show in Section 3.1.3 that for many metrically homogeneous graphs Γ , the automorphism group of the generic expansion of Γ by a linear order (Section 2.1) is extremely amenable (Proposition 1). We will show that the universal minimal flow of the automorphism group of Γ is metrizable, and even identify it explicitly (Theorem 2).

Connections of the universal minimal flow to Ramsey theory

Definition 2.4.3. *Let \mathcal{A} be a class of finite structures in a relational language L . For $\mathbf{A} \leq \mathbf{B}$ in \mathcal{A} let*

$$\binom{\mathbf{B}}{\mathbf{A}} = \{\mathbf{A}' \subseteq \mathbf{B} : \mathbf{A}' \cong \mathbf{A}\}$$

be the set of isomorphic copies of \mathbf{A} contained in \mathbf{B} . For $\mathbf{A} \leq \mathbf{B} \leq \mathbf{C}$ in \mathcal{A} and $r \geq 1$, let

$$\mathbf{C} \rightarrow (\mathbf{B})_r^{\mathbf{A}}$$

mean that for any coloring $c: \binom{\mathbf{C}}{\mathbf{A}} \rightarrow \{1, \dots, r\}$, there is $\mathbf{B}' \in \binom{\mathbf{C}}{\mathbf{B}}$ such that c is constant on $\binom{\mathbf{B}'}{\mathbf{A}}$. We say that \mathcal{A} has the Ramsey Property (RP) if for any $\mathbf{A} \leq \mathbf{B}$ in \mathcal{A} and $r \geq 1$, there is $\mathbf{C} \geq \mathbf{B}$ in \mathcal{A} with $\mathbf{C} \rightarrow (\mathbf{B})_r^{\mathbf{A}}$.

We get the following significant result from [KPT05], which is stated there in terms of a larger class of languages.

Theorem 2.4.2. *Let \mathcal{A}' be an amalgamation class of finite structures in the language $\mathcal{L} \cup \{<\}$, where \mathcal{L} is a finite relational language, and $\{<\}$ denotes a linear order. Let Γ' be the Fraïssé limit of \mathcal{A}' . Then the following are equivalent:*

- (i) *$\text{Aut}(\Gamma')$ is extremely amenable.*
- (ii) *\mathcal{A}' has the Ramsey Property.*

Theorem 2.4.2 will be applied to the study of the automorphism group of the generic expansion of a metrically homogeneous graph of finite diameter by a linear order. Our main focus

is the automorphism group of Γ . A variant of Theorem 2.4.2 included in Theorem 2.4.3 below will give us the metrizable of the universal minimal flow for $\text{Aut}(\Gamma)$. More precise results require a further notion, the *ordering property*.

Definition 2.4.4. *Let \mathcal{L} be a relational language, \mathcal{A} a class of finite \mathcal{L} -structures, and \mathcal{A}' a class of ordered \mathcal{L} -structures, that is, structures for the language $\mathcal{L}' = \mathcal{L} \cup \{<\}$, in which the symbol $<$ is interpreted by a linear order. Suppose that \mathcal{A} is the class of all reducts of structures in \mathcal{A}' to the language \mathcal{L} .*

The class \mathcal{A}' has the ordering property if for every \mathcal{L} -structure $A \in \mathcal{A}$ there is a structure $B \in \mathcal{A}$ such that for any two ordered expansions A' of A and B' of B , both lying in \mathcal{A}' , there is an embedding of A' into B' .

Now we may state the result which connects the topological dynamical properties of the automorphism group of a given structure with the topological dynamical properties of an appropriate expansion. The following result is a special case of [KPT05, Theorem 7.5].

Theorem 2.4.3. [KPT05, Theorem 7.5] *Let \mathcal{L} be a finite relational language, and let $\mathcal{L}' = \mathcal{L} \cup \{<\}$ be its expansion by a symbol for a linear order. Let \mathcal{A} be a strong amalgamation class of \mathcal{L} -structures and \mathcal{A}' the class of all expansions of structures in \mathcal{A} to ordered \mathcal{L} -structures in the language \mathcal{L}' . Let Γ be the Fraïssé limit of the class \mathcal{A} . Then the following hold.*

- *If the class \mathcal{A}' has the Ramsey property, then the universal minimal flow of $\text{Aut}(\Gamma)$ is metrizable.*
- *If the class \mathcal{A}' has both the Ramsey property and the ordering property, then the universal minimal flow of $\text{Aut}(\Gamma)$ may be identified with the compact topological space $L(\Gamma)$ of all linear orders of Γ , under the natural action of $\text{Aut}(\Gamma)$.*

Theorem 7.5 of [KPT05] gives somewhat more detail, which we can neglect here, but also works at a substantially greater level of generality, so we take a moment to explain the setting of [KPT05] and to verify that our special case actually falls under the form of the theorem originally given.

The key difference in the two formulations is found in the notion of *reasonable Fraïssé order class* used in [KPT05]. In addition to generic expansions by linear orders, this covers

such cases as expansions of partial orders by linear orders which are required to extend the given partial order. The requirement is simply that for every embedding $A \rightarrow B$ between structures in \mathcal{A} , every expansion of A which is in \mathcal{A}' embeds (via the same map) into some expansion of B which is in \mathcal{A}' . This certainly applies in the context of generic expansions by linear orders. We have also simplified matters by taking $L(\Gamma)$ to consist of all possible linear orders; this again is the special case corresponding to expansions by generic linear orders.

Even more general variants have been considered (and applied) allowing expansions not necessarily given by linear orderings. But this will suffice for our purposes.

We will use the following result by Bodirsky to establish the ordering property. We specialize to the setting of finite relational languages.

Theorem 2.4.4. *[Adapted from [Bod15, Theorem 6.4]] Let \mathcal{L} be a finite relational language, Γ a homogeneous \mathcal{L} -structure, and Γ' a homogeneous ordered \mathcal{L} -structure, that is, a homogeneous \mathcal{L}' -structure in the language $\mathcal{L} \cup \{<\}$ for which $<$ denotes a linear order. Suppose that the age of Γ' is a Ramsey class, and every acyclic 2-type of Γ does not split in Γ' .*

Then the age of Γ' has the ordering property.

There is some terminology to be explained here. A 2-type of Γ is an orbit of $\text{Aut}(\Gamma)$ on ordered pairs from Γ . In particular, it is a binary relation on Γ which is definable without parameters. A 2-type of Γ is *acyclic* if it defines a directed graph on Γ with no oriented cycles. A 2-type of Γ splits in Γ' if it has more than one expansion to a 2-type in \mathcal{L}' which is realized in Γ' .

In the context of metrically homogeneous graphs, all 2-types are symmetric, and hence none are acyclic (there are cycles of length 2). Accordingly, the meaning of Theorem 2.4.4 in this context is that the Ramsey property implies the ordering property.

2.5 Structural Ramsey Theory

As we have seen in the previous section, the key to understanding the dynamical properties of automorphism groups of countable structures lies in structural Ramsey theory. One powerful method for proving Ramsey theoretic results is the so-called *partite method*.

The following is an early and quite general result obtained by this method, which will serve also as a point of departure for our applications. This result comes from [NR77], but we give a formulation found in [HN16], which is more convenient from the point of view of model theory.

Theorem 2.5.1. [HN16, Theorem 3.6] *Let \mathcal{L} be a relational language and \mathcal{F} a possibly infinite family of ordered irreducible \mathcal{L} -structures. Then the class of all finite ordered \mathcal{L} -structures containing no induced substructure isomorphic to a structure in \mathcal{F} has the Ramsey Property.*

The terminology requires some detailed explanation here. In the first place, as in the previous section, we are considering the language $\mathcal{L}' = \mathcal{L} \cup \{<\}$, and an *ordered \mathcal{L} -structure* is actually an \mathcal{L}' -structure in which the symbol $<$ denotes an ordering.

However, the notions of *irreducible \mathcal{L} -structure* and *ordered irreducible \mathcal{L} -structure* are more subtle (and the latter does not mean *irreducible ordered \mathcal{L} -structure*, as we shall see).

Definition 2.5.1. [HN16, Definition 2.1 and pg. 25] *Let \mathcal{L} be a relational language and $\mathcal{L}' = \mathcal{L} \cup \{<\}$.*

An \mathcal{L} -structure is said to be irreducible if every pair of elements occurs within some n -tuple belonging to one of the \mathcal{L} -relations. In model theory, this is often called Gaifman completeness, i.e. the associated “Gaifman graph” is complete.

An \mathcal{L}' -structure is said to be an ordered irreducible \mathcal{L} -structure if on the one hand it is an ordered \mathcal{L} -structure, and on the other hand the reduct to \mathcal{L} is irreducible.

The statement of Theorem 2.5.1 is non-trivial even when the set of forbidden structures \mathcal{F} is empty. One may use the class \mathcal{F} to impose conditions like irreflexivity and symmetry on the class of structures under consideration. In a model theoretic context, if one wishes to consider the class of finite ordered graphs, then the language \mathcal{L} will consist of a binary relation, and the class of \mathcal{L} -structures is quite broad; thus one needs the restrictions in \mathcal{F} to impose even rudimentary conditions on the structures. One can also do something more substantial: forbid n -cliques for some n . In this way, one picks up a much more delicate Ramsey theoretic result.

The proof of Theorem 2.5.1 involves a double application of the partite method alluded to above: a first step gives the result when \mathcal{F} is empty, and a second step reduces the general case

to this case. This reduction process is considerably generalized by the main result of [HN16], which gives a more flexible approach to the derivation of new Ramsey theoretic results from old ones. This result is very well-suited to a model theoretic setting, and will easily cover the cases of interest to us.

This result involves a very flexible notion of *local finiteness* which we will explain following the statement. In the next section, we will consider a more straightforward notion of *finite constraint* (Definition 2.6.2), which will suffice for our applications.

Theorem 2.5.2. [HN16, Theorem 2.1] *Let \mathcal{L} be a relational language, and let \mathcal{A} be a class of finite irreducible \mathcal{L} -structures, and let \mathcal{A}^* be a hereditary class contained in \mathcal{A} . Suppose the following hold.*

- *The class \mathcal{A} has the Ramsey property.*
- *The class \mathcal{A}^* is a strong amalgamation class.*
- *The class \mathcal{A}^* is locally finite relative to \mathcal{A} .*

Then the class \mathcal{A}^ has the Ramsey property.*

Here it would perhaps be more natural to use the term “locally finitely constrained” in place of “locally finite,” in view of the terminology introduced in the next section, but we follow [HN16] on this point. The main technical definitions run as follows.

Definition 2.5.2. [HN16, Definitions 2.2–2.4] *Let \mathcal{L} be a relational language, \mathcal{A} a class of finite \mathcal{L} -structures, and \mathcal{A}^* a subclass of \mathcal{A} .*

1. *Let A_0 be a finite \mathcal{L} -structure and A a structure in \mathcal{A} containing A_0 as a partial substructure (that is, the relations on A_0 are contained in the relations on A). Then A is a strong \mathcal{A} -completion of A_0 if every irreducible substructure of A_0 is an induced substructure of A .*
2. *The class \mathcal{A}^* is locally finite relative to \mathcal{A} if for every structure X in \mathcal{A} , there is a finite number n such that any structure A in \mathcal{A} which satisfies the following conditions has a strong \mathcal{A} -completion.*
 - *(Local embedding condition) Every substructure of A on n vertices has a strong \mathcal{A} -completion.*

- (Global X -homomorphism condition) *There is a homomorphism from A to X which is injective on every irreducible substructure of A .*

We have rephrased and condensed the definitions from [HN16], concentrating on the points required for our applications later. We will apply this in the context of a metrically homogeneous graph Γ of finite diameter $[\delta]$ expanded by a generic linear ordering, viewing the age of Γ as a subclass of the class of all edge-labeled graphs with labels in $[\delta]$. After due attention to the interpretation of the formalism of [HN16] in this context, our focus will be on the verification of the local finiteness condition of Theorem 2.5.2.

We will use this to conclude that $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$, when augmented with a linear order, is Ramsey. Explicitly, the augmented class will be the following.

Definition 2.5.3. *We define the augmented class $\mathcal{A}_{K_1, K_2, C, C', S}^{\delta, <}$ to consist of the finite metric spaces of $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$ each of which now are augmented with a linear order $<$. We define $(\Gamma, <)$ to be the corresponding Fraïssé limit.*

This is an example of an expansion via a *generic linear order*.

Definition 2.5.4. *Let \mathcal{A} be a strong amalgamation class. Then its ordered expansion \mathcal{A}^\leq is an amalgamation class as well. We say that \leq is a generic linear order of the Fraïssé limit Γ of \mathcal{A} if it is a linear order on Γ so that the expansion (Γ, \leq) is isomorphic to the Fraïssé limit of \mathcal{A}^\leq .*

2.6 Partial Structures and Finite Constraint

When dealing with amalgamation classes via the Fraïssé theory, we will first pass to a larger class of partial substructures of the Fraïssé limit and analyze this class combinatorially. The present section establishes a formalism for passing back and forth between the point of view of Fraïssé theory and this more combinatorial setting.

In order to obtain our finiteness results, we require the following terminology.

Definition 2.6.1. *Let \mathcal{L} be a fixed relational language. All structures considered here are \mathcal{L} -structures.*

1. *Let A, B be relational structures with $A \subseteq B$ as sets. Then A is a partial substructure of B if every relation holding in A holds also in B (i.e., the identity map is a homomorphism).*

2. Let \mathcal{A} be a class of structures. Then $\tilde{\mathcal{A}}$ denotes the class of structures which are isomorphic to partial substructures of structures in \mathcal{A} .
3. $\tilde{\mathcal{A}}^c$ denotes the complement of $\tilde{\mathcal{A}}$ in the class of all \mathcal{L} -structures. Structures in $\tilde{\mathcal{A}}^c$ will be called the forbidden partial structures associated with the class \mathcal{A} (e.g., in the context of graph theory, one would speak of forbidden subgraphs).

Definition 2.6.2. Let \mathcal{L} be a relational language and let \mathcal{A} be a class of finite \mathcal{L} -structures.

We say that \mathcal{A} is finitely constrained if there exists some finite set \mathcal{F} of finite “forbidden” \mathcal{L} -structures such that the structures in $\tilde{\mathcal{A}}$ are precisely the finite \mathcal{L} -structures not containing any structure in \mathcal{F} as a partial substructure.

It might seem more natural to call the class $\tilde{\mathcal{A}}$ finitely constrained in this context, but as the property will be applied to give information about the original class \mathcal{A} , we have chosen this less natural terminology. It should be noted that the classes \mathcal{A} of interest are generally “finitely constrained” in a very different sense: namely, they are defined by finitely many forbidden *induced substructures*. It turns out that the notion of finite constraint introduced in Definition 2.6.2 is a non-trivial and powerful condition on the class \mathcal{A} . We illustrate with a pair of examples (cf. Sections 3.1.1, 3.1.2).

Example 2.6.1. Let \mathcal{A}^δ be the class of all finite $[\delta]$ -valued metric spaces, and let \mathcal{U}^δ be the class of all finite $[\delta]$ -valued ultrametric spaces.

Then \mathcal{A}^δ is finitely constrained and \mathcal{U}^δ is not finitely constrained.

In this example, both classes \mathcal{A}^δ and \mathcal{U}^δ are determined by finitely many forbidden substructures, specifically the triangles violating the triangle inequality or the ultrametric inequality, respectively. But the class of $[\delta]$ -metric spaces is also finitely constrained in our stronger sense: it suffices to avoid violations of the triangle inequality along cycles of length at most δ (Definition 3.1.5), whereas in the case of \mathcal{U}^δ there are arbitrarily long cycles violating the ultrametric inequality as soon as $\delta \geq 2$.

We need to develop systematic methods for relating properties of amalgamation classes \mathcal{A} with properties of the associated class $\tilde{\mathcal{A}}$. The first essential point is to see how freeness with respect to homomorphisms arises if one begins with a suitable amalgamation class of structures (Lemma 2.6.1 below). This relies on the following notions.

Definition 2.6.3. Let \mathcal{A} be a class of finite structures. A structure $A \in \mathcal{A}$ is relationally maximal in \mathcal{A} if it is an induced structure of any structure $B \in \mathcal{A}$ which contains it as a partial substructure. The class \mathcal{A} is relationally complete if all of its elements are relationally maximal in \mathcal{A} .

A relation in n variables is irreflexive if it contains no n -tuple with at least two equal entries. A relational \mathcal{L} -structure will be called irreflexive if all of the relations in the language \mathcal{L} other than equality are irreflexive.

Example 2.6.2. The class of graphs in the usual language of graph theory is not relationally complete. The class of tournaments is relationally complete in its natural language.

The class of $[\delta]$ -metric spaces is relationally complete when construed in the usual way as the class of edge-labeled graphs with labels from $[\delta]$.

Lemma 2.6.1. Let \mathcal{A} be a relationally complete amalgamation class of finite irreflexive structures in a relational language \mathcal{L} . Then the following are equivalent.

1. \mathcal{A} is a strong amalgamation class.
2. \mathcal{A}^c is closed under homomorphism.

Proof. We will use irreflexivity in the forward direction and relational completeness in the reverse direction.

(1 \Rightarrow 2)

Let $A \in \widetilde{\mathcal{A}}^c$ be a forbidden partial structure and $h: A \rightarrow A'$ a homomorphism. We claim that A' is forbidden as well.

Any homomorphism can be obtained by composing injective homomorphisms with canonical homomorphisms that identify exactly two vertices a, b of A and take only the induced relations, so we may suppose that h is either injective or canonical in this sense.

If h is injective, the claim is clear. We may suppose therefore that h identifies the points a, b of A to a single point c , and let A' be the resulting homomorphic image of A .

Now suppose toward a contradiction that A' is not forbidden, and thus extends to a structure B in \mathcal{A} . Let A_0 be the structure induced by A on $A \setminus \{a, b\}$. Take two copies B_1, B_2 of B with intersection A and consider the amalgamation diagram $A_0 \rightarrow B_1, B_2$. By strong amalgamation

and closure under isomorphism, we may take a structure $C \in \mathcal{A}$ containing $B_1 \cup B_2$ (the structure formed by taking the union of the sets and the union of the relations).

We claim that C is an extension of the image of A under the map sending a, b to the images of c in B_1, B_2 respectively.

For any relation not involving both a and b this holds since $A' \subseteq B \simeq B_1, B_2$. On the other hand, there can be no relation on A holding between a and b , since the equation $a = b$ does not hold and any other relation would give rise to a violation of irreflexivity in A' . Thus C extends the image of A , and as A is not in $\widetilde{\mathcal{A}}$, we have arrived at a contradiction.

(2 \Rightarrow 1)

Consider an amalgamation problem

$$A_0 \rightarrow A_1, A_2.$$

Let A be the free join of A_1, A_2 over A_0 , formed by taking copies of A_1, A_2 disjoint over A and forming the union of the underlying sets and the union of the corresponding relations. As \mathcal{A} is an amalgamation class, there is some amalgam B of A_1, A_2 over A_0 , and by construction the natural map from A to B will then be a homomorphism.

Since $B \in \mathcal{A}$, it follows from our hypothesis that $A \in \widetilde{\mathcal{A}}$ and thus A extends to some $C \in \mathcal{A}$. In particular, B_1, B_2 are partial structures in C and by relational completeness are induced substructures of C . Thus C is a strong amalgam of A_1, A_2 over A . \square

We extend the previous lemma to take into account the finiteness conditions on the associated constraints.

Lemma 2.6.2. *Let \mathcal{A} be a relationally complete strong amalgamation class of finite irreflexive structures in a relational language \mathcal{L} . Then the following are equivalent.*

- *The class \mathcal{A} is finitely constrained.*
- *The associated class $\widetilde{\mathcal{A}}$ of partial structures is the class of finite homomorphically \mathcal{T} -free structures for some finite set \mathcal{T} of finite \mathcal{L} -structures.*

Proof. (1 \Rightarrow 2)

By assumption, there is a finite set $\mathcal{T} \subseteq \mathcal{A}^c$ of finite \mathcal{L} -structures such that any structure not in $\tilde{\mathcal{A}}$ contains some structure in \mathcal{T} as a partial substructure. By Lemma 2.6.1, any \mathcal{L} -structure containing a homomorphic image of a structure in \mathcal{T} is in $\tilde{\mathcal{A}}^c$. Thus $\tilde{\mathcal{A}}$ is the class of homomorphically \mathcal{T} -free structures.

(2 \Rightarrow 1)

Conversely, if \mathcal{T} is a finite set of \mathcal{L} -structures and $\tilde{\mathcal{A}}$ is the class of homomorphically \mathcal{T} -free structures, then let \mathcal{F} be the closure of \mathcal{T} under homomorphic image. Then \mathcal{F} is again finite, and provides a finite set of constraints for partial \mathcal{A} -structures, in the sense of Definition 2.6.2. Then the structures in $\tilde{\mathcal{A}}$ are precisely the finite \mathcal{L} -structures not containing any structure in \mathcal{F} as partial substructure, so \mathcal{A} is finitely constrained. \square

Lemma 2.6.2 allows for the direct application of the results of [HL99].

Lemma 2.6.3. *Let \mathcal{A} be a finitely constrained relationally complete strong amalgamation class of finite irreflexive structures in a relational language \mathcal{L} . Let \mathcal{C} be the class of structures C such that every finite partial substructure of C lies in $\tilde{\mathcal{A}}$.*

Then \mathcal{C} has the weak EPPA.

Proof. By Lemma 2.6.2, there is a finite set \mathcal{T} of finite \mathcal{L} -structures such that $\tilde{\mathcal{A}}$ is the class of finite homomorphically \mathcal{T} -free structures.

Then the class \mathcal{C} is the set of homomorphically \mathcal{T} -free structures. By Theorem 2.3.3, the class \mathcal{C} has the weak EPPA. \square

In order to transfer information about the APA or the EPPA between a class of partial substructures and a class of induced substructures, we will need to have a *canonical completion process* in the following sense.

Definition 2.6.4. *Let \mathcal{L} be a fixed language and suppose that ζ is a function (or algorithm) which takes \mathcal{L} -structures to \mathcal{L} -structures.*

1. *We say that ζ is a completion process if and only if the following hold, with \mathcal{A} the range of ζ .*

- *For A in the domain of ζ , A is a partial substructure of $\zeta(A)$.*

- The domain of ζ contains $\tilde{\mathcal{A}}$.
- On \mathcal{A} , ζ is the identity.

2. A completion process ζ is called *canonical* if its domain is closed under isomorphism, and for any isomorphism

$$\alpha : A \rightarrow B$$

between structures in the domain of ζ , the map α is also an isomorphism between $\zeta(A)$ and $\zeta(B)$.

A canonical completion process provides a *stationary* independence relation in the sense of Tent and Ziegler [TZ12]. This amounts to a canonical completion process for amalgamation diagrams, and amalgamation diagrams are a particular instance of partial structures.

A canonical completion process gives us the APA, in the following sense.

Lemma 2.6.4. *Let \mathcal{A} be a strong amalgamation class of structures, and suppose that there is a canonical completion process for partial \mathcal{A} -structures. Then \mathcal{A} has the amalgamation property for automorphisms (APA).*

Proof. We take an amalgamation diagram $A_0 \rightarrow A_1, A_2$ in \mathcal{A} , where A_0 may be empty. We may suppose that the embeddings from A_0 to A_1 and A_2 are inclusions and that $A_1 \cap A_2 = A_0$.

As \mathcal{A} is a strong amalgamation class, the union $A_1 \cup A_2$ is a partial \mathcal{A} -structure. Therefore it has a canonical completion B in \mathcal{A} . We will check that B serves to amalgamate automorphisms, in the sense of Definition 2.3.10.

Fix automorphisms α_1, α_2 of A_1, A_2 respectively which leave A_0 invariant and agree on A_0 . The union $\alpha = \alpha_1 \cup \alpha_2$ is then an automorphism of the partial structure $A_1 \cup A_2$. By canonicity (invariance under isomorphisms), it is also an automorphism of the canonical completion of B .

The lemma follows. □

A canonical completion process also allows transfer of information about EPPA between classes of partial substructures and classes of induced substructures.

Lemma 2.6.5. *Let \mathcal{A} be an amalgamation class of finite structures, and suppose that there is a canonical completion process for finite partial \mathcal{A} -structures. Let \mathcal{C} be the class of structures*

C such that every finite partial substructure of C belongs to $\tilde{\mathcal{A}}$. If \mathcal{C} has the weak EPPA, then \mathcal{A} has the EPPA.

Proof. Let A be a structure in \mathcal{A} , and let $(\alpha_1, \dots, \alpha_n)$ be a list of all partial automorphisms of A .

The Fraïssé limit Γ of \mathcal{A} belongs to the class \mathcal{C} , so application of the weak EPPA in \mathcal{C} , taking $C = \Gamma$ in the definition, gives a structure B in $\tilde{\mathcal{A}}$ containing A , such that all partial automorphisms of A extend to automorphisms of B . Now let C be the canonical completion of B in \mathcal{A} . Then any automorphism of B acts on C as an automorphism as well. Thus all partial automorphisms of A extend to C and the EPPA holds in \mathcal{A} . \square

We may now prove a useful general criterion for the EPPA.

Proposition 2.6.1. *Let \mathcal{A} be a finitely constrained relationally complete strong amalgamation class of finite irreflexive structures in a relational language \mathcal{L} . Suppose that there is a canonical completion process for finite partial \mathcal{A} -structures.*

Then \mathcal{A} has the EPPA.

Proof. Let \mathcal{C} be the class of structures C such that every finite partial substructure of C lies in $\tilde{\mathcal{A}}$. By Lemma 2.6.3, the class \mathcal{C} has the weak EPPA. By Lemma 2.6.5, the class \mathcal{A} has the EPPA. \square

At one point in our finiteness argument, the following will be useful.

Lemma 2.6.6. *The set of finite sequences over a finite alphabet is well-quasi-ordered by the subsequence relation.*

2.7 Twisted Isomorphisms

In Chapter 4, we will study the *twisted isomorphisms* between metrically homogeneous graphs and, in particular, the *twisted automorphism group*.

Definition 2.7.1. *Let \mathcal{L} be a relational language. A twisted isomorphism between two \mathcal{L} -structures is an isomorphism up to a permutation of the language.*

The following more explicit definition has some practical utility.

Definition 2.7.2. *Let \mathcal{L} be a relational language, and let σ be a permutation of the relational symbols of \mathcal{L} which respects the number of variables in the relations.*

1. *If Γ is an \mathcal{L} -structure, then Γ^σ denotes the structure with the same universe, and the same set of relations, but with*

$$(R^\sigma)^{\Gamma^\sigma} = R^\Gamma \quad (R \in \mathcal{L}).$$

In other words, in Γ^σ the symbol R^σ stands for the relation denoted by R in Γ .

2. *A σ -isomorphism between two \mathcal{L} -structures Γ_1 and Γ_2 is an isomorphism of Γ_1 with Γ_2^σ .*
3. *A twisted isomorphism between two \mathcal{L} -structures is a σ -isomorphism for some permutation σ of the language, as above.*

The permutation σ is called the associated twist.

A twisted automorphism of a structure Γ is a twisted isomorphism of Γ with itself.

Observe that as Γ and Γ^σ carry the same set of relations, they have the same automorphism groups. Thus twisted isomorphisms preserve substantial amounts of information; in fact, all information not tied to the choice of language.

We mention some typical and familiar examples.

Example 2.7.1.

1. *Let F be a field and consider F -vector spaces V_1, V_2 viewed in the functional language with the binary operation $+$ and the unary multiplication operators $(\mu_a : a \in F)$, or in the equivalent relational language, where the definitions above may be applied directly.*

A twisted isomorphism between V_1, V_2 will be a isomorphism of vector spaces up to an isomorphism of the base field (unless V_1, V_2 are trivial, in which case any permutation of F is allowed). This holds since V_2^σ must satisfy the vector space axioms for the corresponding operations, namely: $\mu_{ab}^\sigma = \mu_a^\sigma \mu_b^\sigma$ and $\mu_{(a+b)}^\sigma = \mu_a^\sigma + \mu_b^\sigma$, and if V_2 is non-trivial, the operations determine the corresponding field elements.

2. *Let Γ_1 be a graph and Γ_2 its graph complement. If we work in a language with symbols E^+ and E^- for the edge and non-edge relations, then the identity map is a twisted isomorphism of Γ_1 with Γ_2 , where the associated twist σ swaps the two symbols.*

3. Similarly, we may consider twisted automorphisms of a graph Γ , which include isomorphisms with the graph complement.

As the last example indicates, the notion of twisted isomorphism is excessively dependent on the choice of language, and in the interesting case one chooses the language to allow the largest possible group of twisted isomorphisms. It may not be immediately clear that there is such a largest choice, but as we will see, the canonical language in the sense of Section 2.1 is appropriate here.

For our present purposes, there is no loss of generality in considering structures with the same underlying set, since any twisted isomorphism is the composition of an ordinary isomorphism with a twisted isomorphism which does not change the underlying set.

Lemma 2.7.1. *Let \mathcal{L} be a relational language, and let Γ_1, Γ_2 be \mathcal{L} -structures with the same underlying set Ω .*

1. *If $f \in \text{Sym}(\Omega)$ is a twisted isomorphism of Γ_1 with Γ_2 , then f conjugates $\text{Aut}(\Gamma_1)$ to $\text{Aut}(\Gamma_2)$. In particular, twisted automorphisms of Γ_1 normalize $\text{Aut}(\Gamma_1)$.*
2. *Conversely, if \mathcal{L} is interpreted as the canonical language for Γ_1 and for Γ_2 , then a permutation $f \in \text{Sym}(\Omega)$ which conjugates $\text{Aut}(\Gamma_1)$ to $\text{Aut}(\Gamma_2)$ is a twisted isomorphism.*

Proof.

1. Let $h \in \text{Aut}(\Gamma_1)$, let R be a relation symbol of \mathcal{L} . Let R_1 be the relation on Γ_1 and R_2 be the relation on Γ_2 . We claim that fhf^{-1} , acting on the left, preserves R_2 .

Let σ be the twist associated with f . Then R^σ denotes the relation R_2 in Γ_2^σ . So by assumption, $R_1 = f^{-1}[R_2]$ is the relation denoted by R on Γ_1 . Hence $hf^{-1}[R^\sigma] = R_1$ as well, and $fhf^{-1}[R_2]$ is again R_2 .

2. Now suppose that the symbols of \mathcal{L} denote the orbits of $\text{Aut}(\Gamma_i)$ on Γ_i . If f conjugates $\text{Aut}(\Gamma_1)$ to $\text{Aut}(\Gamma_2)$ then f takes the orbits of $\text{Aut}(\Gamma_1)$ to the orbits of $\text{Aut}(\Gamma_2)$ and hence induces a permutation σ of the language \mathcal{L} , and a σ -isomorphism of Γ_1 with Γ_2 . □

Now we examine the case of twisted automorphisms.

Lemma 2.7.2. *Let \mathcal{L} be a relational language and Γ an \mathcal{L} -structure with underlying set Ω .*

Then the twisted automorphisms of Γ form a subgroup of $\text{Aut}^*(\Gamma)$ of $\text{Sym}(\Omega)$ is contained in the normalizer

$$N_{\text{Sym}(\Omega)}(\text{Aut}(\Gamma)).$$

Moreover, the permutations of \mathcal{L} associated with the twisted automorphisms of Γ form a subgroup $\text{Out}(\Gamma)$ of $\text{Sym}(\mathcal{L})$.

If distinct relation symbols in \mathcal{L} represent distinct relations on Γ , then there is a canonical surjective homomorphism

$$\text{Aut}^*(\Gamma) \rightarrow \text{Out}(\Gamma)$$

with kernel $\text{Aut}(\Gamma)$.

If \mathcal{L} is the canonical language for Γ , then

$$\text{Aut}^*(\Gamma) = N_{\text{Sym}(\Omega)}(\text{Aut}(\Gamma))$$

$$\text{Out}(\Gamma) \simeq N_{\text{Sym}(\Omega)}(\text{Aut}(\Gamma)) / \text{Aut}(\Gamma)$$

Proof. This is largely formal. Since $(\Gamma^\sigma)^\tau = \Gamma^{\sigma\tau}$ for any two $\sigma, \tau \in \text{Sym}(\Omega)$, it follows that the twisted automorphisms and the twists each form groups.

If the symbols of \mathcal{L} represent distinct relations on Γ , then clearly the map from twisted automorphisms to associated twists is well-defined and is a homomorphism. In particular, the kernel $\text{Aut}(\Gamma)$ is normal in $\text{Aut}^*(\Gamma)$ in this case.

In the case in which \mathcal{L} has more than one symbol representing a given relation, we may eliminate redundancies; this may increase the twisted automorphism group but will not decrease it, so $\text{Aut}(\Gamma)$ must be normal in $\text{Aut}^*(\Gamma)$ in general.

Coming down to the case in which \mathcal{L} is the canonical language for Γ , and applying the previous lemma, we see that $\text{Aut}^*(\Gamma)$ is the full normalizer $N_{\text{Sym}(\Omega)}(\text{Aut}(\Gamma))$.

Thus everything is proved. □

We prefer to work with finite languages and finite permutation groups where possible. In particular, we have already chosen a language for metrically homogeneous graphs of diameter δ . The following result allows this.

Lemma 2.7.3. *Let k be fixed, and let Γ be a homogeneous structure whose language is the restriction of the canonical language to relations in k variables. Then*

$$\text{Aut}^*(\Gamma) = N_{\text{Sym}(\Omega)}(\text{Aut}(\Gamma)).$$

Proof. The proof is the same as in the case of the full canonical language. The inclusion $\text{Aut}^*(\Gamma) \subseteq N_{\text{Sym}(\Omega)}(\text{Aut}(\Gamma))$ holds generally, and the reverse inclusion holds since the normalizer of $\text{Aut}(\Gamma)$ permutes its orbits on k -tuples. \square

One must be slightly careful in applying this result. If one starts with a homogeneous structure Γ then the language will be traded in for a more symmetric language \mathcal{L}_k , but k must be chosen large enough so that the resulting automorphism group is still $\text{Aut}(\Gamma)$. In other words, we require Γ to be homogeneous with respect to a language of bounded complexity.

Our customary language for metrically homogeneous graphs is in fact the language \mathcal{L}_2 (even when the diameter δ is infinite). A similar analysis would show that this is satisfactory not only for the consideration of twisted automorphisms, but also for the consideration of twisted isomorphisms between pairs of structures.

Chapter 3

Topological Results

The present chapter deals with the properties of the automorphism group of a metrically homogeneous graph, when that group is viewed as a Polish group. This involves on the one hand descriptive set theoretic or topological properties of the group, and notably the question as to the extent to which the purely algebraic structure determines the topological structure, and on the other hand the dynamical properties of the group in terms of continuous actions on compact topological spaces. In technical terms, the former is handled using the method of ample generics (discussed in Section 2.3), while the latter is reduced to structural Ramsey theory (Section 2.5). Remarkably, the main tool in both cases is the finiteness property discussed in Section 3.1, and an associated *completion process*. We obtain this property for a broad class of known metrically homogeneous graphs of generic type. Subsequent work [ABH⁺17] extends this combinatorial result to all known primitive metrically homogeneous graphs of generic type, and therefore the related topological and dynamical results can be obtained at the same level of generality.

In Section 3.1 we deal with the purely combinatorial issues. That is, we show that the amalgamation classes associated with certain primitive metrically homogeneous graphs of generic type are finitely constrained, in the sense of Definition 2.6.2, via a completion process which has an additional *canonicity* property which will be exploited further. In Section 3.2, we exploit both finite constraint and the associated canonical completion process as discussed in Section 2.3, and in Section 3.3, we proceed similarly to derive dynamical properties of the automorphism group via connections with structural Ramsey theory.

3.1 Finite constraint

In this section, we will show that the classes of finite metric spaces associated with the known primitive metrically homogeneous graphs of generic type whose associated numerical parameters satisfy some additional restrictions are finitely constrained (Proposition 2). This makes use of a fairly delicate completion process.

As a model for the construction, we first treat two similar problems for the classes of finite metric spaces defined by less intricate conditions. Namely, in Section 3.1.1, we deal with the class of all $[\delta]$ -metric spaces for fixed finite δ , which is an instance of our general problem. Then in Section 3.1.2, we consider the broader class of S -metric spaces for S an arbitrary finite subset of the positive real numbers. This is not an instance of our general problem, but it is a case of substantial independent interest, in which the corresponding issues are more transparent.

Sauer characterized the finite sets S for which the S -metric spaces form an amalgamation class by a condition which may be expressed as the associative law in an associated semigroup defined on the set S , while Hubička and Nešetřil [HN16, §4.2.2] characterized the finite sets S for which the class is a finitely constrained amalgamation class by a further condition on this semigroup, viewed now as an ordered semigroup. In Section 3.1.2, we characterize the finite sets S for which finite constraint holds in terms of the same operation; however, as not all such sets correspond to amalgamation classes, we must work with this operation also in a non-associative context. Remarkably, a characterization very similar to that of [HN16] results, even in the absence of associativity.

After these preparations, we return to the case of metrically homogeneous graphs in Section 3.1.3 and show finite constraint under the assumptions of Proposition 2. Later analysis has shown that there is also an associated *partially ordered* semigroup, on which the relevant completion process can be viewed as a “shortest path” completion (as in Section 3.1.2), but this lies outside the scope of our discussion.

3.1.1 $[\delta]$ -metric spaces

In this section we prove the following, to serve as a general template for all of our work on finite constraint.

Proposition 3.1.1. *The class of $[\delta]$ -metric spaces, for a fixed finite δ , is finitely constrained.*

Our general approach is as follows. One may associate with any weighted graph the metric space given by the associated *path metric*. We may call this the *path completion* of the given metric space. But this is an abuse of technology, as the given weights may not be preserved in the path completion (notably, if the original graph was a complete graph but not a metric space). At the same time, the passage from the weighted graph to the path completion is a reasonable completion process in the following sense: if there is any extension of the given weighted graph to a metric space, then the path completion is one such extension.

With minor modifications, this completion process can be adjusted to one applicable to $[\delta]$ -weighted graphs, which produces a $[\delta]$ -metric extension wherever there is one. This may be done by making a direct and ad hoc adjustment to the definition, or by modifying the operation of addition to restrict it to $[\delta]$ and then extending the notion of path metric to allow the use of a modified addition.

Once we have settled on a completion process which serves our purposes, we know that the partial $[\delta]$ -metric spaces are those whose weights are preserved by the completion process. The remaining step in the proof of finite constraint is to determine the minimal obstructions to this last condition, and to show that only finitely many occur. The same method will apply in the following section to the class of S -metric spaces, and then, with more substantial adjustments, to certain primitive metrically homogeneous structures of generic type.

Definition 3.1.1. *Let $\mathcal{G} = (G, w)$ be an edge-labeled graph with edge weights in $\mathbb{R}_{>0}$.*

For γ a path in \mathcal{G} , the weight $w(\gamma)$ is the sum of the weights along γ .

The path metric d_p on \mathcal{G} is given by

$$d_p(u, v) = \min(w(\gamma) : \gamma \text{ is a path from } u \text{ to } v),$$

for $u, v \in G$.

If no such path exists, then the path distance is set to ∞ .

For $\delta > 0$, the δ -restricted path metric d_p^δ is defined by

$$d_p^\delta = \min(d_p(u, v), \delta).$$

We set $\bar{\mathcal{G}} = (G, d_p^\delta)$.

The following is well-known and easily checked.

Lemma 3.1.1.

1. Let Γ be an $\mathbb{R}_{>0}$ -edge-labeled graph. Then the path metric on Γ is a generalized metric with values in $\mathbb{R}_{>0} \cup \{\infty\}$.
2. The infimum of two generalized metrics is a generalized metric.

Corollary 3.1.1. For any $[\delta]$ -edge-labeled graph $\mathcal{G} = (G, w)$, the edge-labeled graph $\bar{\mathcal{G}} = (G, d_p^\delta)$ is a $[\delta]$ -metric space.

We are now ready to prove the following.

Lemma 3.1.2. Let \mathcal{G} be an $[\delta]$ -edge-labeled graph. Then the following are equivalent.

1. \mathcal{G} is a partial $[\delta]$ -metric space.
2. \mathcal{G} is a labeled subgraph of $\bar{\mathcal{G}}$.
3. \mathcal{G} contains no cycles (u_1, u_2, \dots, u_n) such that $w(u_1, u_n) > w(u_1, \dots, u_n)$.

We take a moment to distinguish between the first two items. The first item says that \mathcal{G} can be extended to some $[\delta]$ -metric space, and the second item says that the above defined completion will extend \mathcal{G} , in that it will respect the original edge weights of \mathcal{G} , and indeed yield a $[\delta]$ -metric space, though this last point was already shown in Corollary 3.1.1. Their equivalence means that whether or not \mathcal{G} can be extended to a $[\delta]$ -metric space is determined by whether or not the defined completion will indeed extend \mathcal{G} .

Proof.

(1 \longrightarrow 3)

Any metric space will satisfy (3), and any partial subspace will therefore also satisfy (3).

(3 \longrightarrow 2)

In other words, if we have no violations of the generalized triangle inequality, then the completion process given in Definition 3.1.1 will maintain the weights originally present in $\mathcal{G} = (G, w)$.

Item 3 ensures that for every edge (u, v) in \mathcal{G} and every path γ from u to v , we have $w(u, v) \leq w(\gamma)$. Therefore $w(u, v) = d_p(u, v)$. Of course, $w(u, v) \leq \delta$, and thus $w(u, v) = d_p^\delta(u, v)$.

(2 \longrightarrow 1)

This follows immediately from the definition of a partial $[\delta]$ -metric space. \square

Proof of Proposition 3.1.1. It remains to show that there are finitely many cycles which violate the generalized triangle inequality in (3).

That is, we must show that there are finitely many cycles (u_1, \dots, u_n) where $w(u_1, u_n) > w(u_1, \dots, u_n)$. This follows, as there are finitely many choices for $n, w(u_1, \dots, u_n), w(u_1, u_n)$ such that

$$n \leq w(u_1, \dots, u_n) < w(u_1, u_n) \leq \delta.$$

\square

3.1.2 S -metric spaces

Our main result is the following.

Proposition 3.1.2. . For a finite set $S \subseteq \mathbb{R}_{>0}$, the following are equivalent.

1. For $a, b \in S$ with $a < \max(S)$, there exists $s \in S$ such that $a < s \leq a + b$.
2. S -metric spaces are finitely constrained.

We begin by giving a completion process for completing an S -edge-labeled graph to an S -metric space. After checking that the completion process indeed yields an S -metric space, we show the sufficiency of the restriction on S , that is, (1 \rightarrow 2), after which we show its necessity.

Definition 3.1.2. Let $\mathcal{G} = (G, w)$ be an edge-labeled graph with edge weights in a finite set S .

Define

$$a \oplus_S b = \max\{s \in S : s \leq a + b\}$$

for $a, b \in S$.

We record some formal properties of the operation \oplus_S .

Lemma 3.1.3. The operation \oplus_S has the following properties.

- Commutativity.
- $\max(a, b) \leq a \oplus_S b \leq a + b$.

- *Monotonicity: if $a_1 \leq a_2$ and $b_1 \leq b_2$ then $a_1 \oplus_S b_1 \leq a_2 \oplus_S b_2$.*

A *walk* is a sequence of vertices and edges $(v_0, e_1, v_1, \dots, e_n, v_n)$ with each term incident with the next. In our context, this may be represented either as a sequence of vertices joined by edges or a sequence of edges with common vertices.

If P_1, P_2 are walks in \mathcal{G} with the terminal vertex of P_1 equal to the initial vertex of P_2 , we write $P_1 + P_2$ for the combined walk, which is simply the concatenation when P_1, P_2 are represented as sequences of edges. We are interested primarily in *paths*, where by definition the vertices do not repeat, except possibly the first and the last. When the context is clear, we sometimes only write the ordered vertices of a path.

Define the *length* $|P|$ of a walk as the number of edges occurring in the walk.

Definition 3.1.3. Define the S -weight $w_S(P)$ of a walk P inductively by

$$w_S(P) = \begin{cases} 0 & \text{if } |P| = 0 \\ w(e) & \text{if } |P| = 1 \\ \inf(w_S(P_1) \oplus w_S(P_2) \mid \\ \quad P = P_1 + P_2 \quad |P_1|, |P_2| > 0) & \text{otherwise} \end{cases}$$

Finally, we define the path metric d_S in \mathcal{G} by

$$d(x, y) = \min(w_S(P) \mid P \text{ a path from } x \text{ to } y).$$

If no such path exists, then we define $d(x, y) = \max(S)$.

Lemma 3.1.4. Let $\mathcal{G} = (G, w)$ be an S -edge-labeled graph, let Q be a walk in \mathcal{G} , and let P be a walk contained in Q , with the edges ordered as in Q . Then

$$w_S(P) \leq w_S(Q).$$

Proof. If Q contains a single edge, this is clear.

Otherwise, we argue inductively. Write $Q = Q_1 + Q_2$ with

$$w_S(Q) = w_S(Q_1) \oplus_S w_S(Q_2).$$

Then P decomposes as $P_1 + P_2$ with P_1, P_2 contained in Q_1, Q_2 respectively, in order. Then

$$w_S(P) \leq w_S(P_1) \oplus_S w_S(P_2) \leq w_S(Q_1) \oplus_S w_S(Q_2) = w_S(Q)$$

by induction and monotonicity. \square

The condition that the ordering on P be induced by the ordering on Q is necessary.

Example 3.1.1. Let S be the set $\{2, 3, 4, 5, 7, 9\}$ and let \mathcal{G} be a triangle with edges of lengths $(2, 3, 4)$. Then \mathcal{G} contains a walk P_1 with weights $(2, 3, 4)$ and another walk P_2 with weights $(3, 2, 4)$. The weight of P_1 is $\min(9, 9) = 9$ and the weight of P_2 is $\min(9, 7) = 7$.

Remark 3.1.1. Every walk W between two vertices u, v contains a path P from u to v with the edges of P inheriting the order from W . In fact, P may be taken to consist of an initial segment of W followed by a terminal segment of W . By Lemma 3.1.4, we find $w_S(P) \leq w_S(W)$ and thus if we define the same path metric using walks rather than paths, we obtain the same metric.

Lemma 3.1.5. Let $S \subseteq \mathbb{R}_{>0}$ be finite and let $\mathcal{G} = (G, w)$ be an S -edge-labeled graph. Then the path metric d_S is in fact a metric on G .

Proof. This is clear using walks rather than paths to define the metric: given vertices u, v, w and walks from u to v and from v to w , their sum is a walk from u to w . \square

Now we move on to show the sufficiency of the restriction on S .

Definition 3.1.4. Let $S \subseteq \mathbb{R}_{>0}$. A non-metric S -cycle is an S -edge-labeled cycle consisting of a path P and one additional edge e such that

$$w(e) > w_S(P).$$

Lemma 3.1.6. Let $S \subseteq \mathbb{R}_{>0}$. Then an S -edge-labeled graph (G, w) embeds in an S -metric space if and only if it contains no non-metric S -cycle.

Proof. We show first that no non-metric S -cycle can be a substructure of an S -metric space (G, d) . In other words, we claim that the S -metric inequality implies the generalized triangle inequality

$$d(u, v) \leq w_S(P) \text{ for } P \text{ a path from } u \text{ to } v.$$

This is immediate by induction on the length of P , using the ordinary triangle inequality. It is clear for $|P| \leq 1$. For $|P| > 1$, we write $P = P_1 + P_2$ with $|P_1|, |P_2| > 0$ and $w_S(P) = w_S(P_1) \oplus w_S(P_2)$, and we let w be the terminal vertex of P_1 . Since $d(u, v) \in S$ and $d(u, v) \leq d(u, w) + d(w, v)$, we then have by induction and monotonicity

$$d(u, v) \leq d(u, w) \oplus d(w, v) \leq w_S(P_1) \oplus w_S(P_2) = w_S(P)$$

as required.

Now we show that any S -edge-labeled graph not containing a non-metric S -cycle as a substructure can be extended to an S -metric space. We use the path metric d_S with respect to S , which certainly gives G the structure of an S -metric space. It remains to show that d_S extends the given weight function w .

Suppose $u, v \in G$ and $w(u, v)$ is defined. By definition, $d_S(u, v) \leq w(u, v)$. If $d_S(u, v) < w(u, v)$, then there is a path P from u to v with $w_S(P) < w(u, v)$. Letting Q be the path (v, u) , $P + Q$ is then a non-metric S -cycle, and we have a contradiction. \square

This now allows us to deduce the following.

Proposition 3.1.3. *Let $S \subseteq \mathbb{R}_{>0}$ be a finite set and suppose that for every $a, b \in S$ with $a < \max S$ we have $a < a \oplus_S b$. Then the class of S -metric spaces is finitely constrained.*

Proof. In view of the preceding lemma, it suffices to show that the set of non-metric cycles is finite. Since S is finite, it suffices to bound the lengths of these cycles.

Let the elements of S be listed in increasing order as (s_1, s_2, \dots, s_k) . We may show by induction that any walk P of length $n \geq 2^i$ has weight at least s_i for $i \leq k$, using the relation $w_S(P_1) \oplus w_S(P_2) > w_S(P_1), w_S(P_2)$ which holds for $|P_1|, |P_2| > 0$ and $w_S(P_1), w_S(P_2) < s_k$, by assumption on S .

For a cycle $\gamma = (u_0, \dots, u_n)$ in \mathcal{F} , letting P be the path (u_0, \dots, u_{n-1}) we have

$$w_S(P) < w(u_{n-1}, u_n) \leq \max S$$

and hence the length $n - 2$ is at most 2^{k-1} .

Thus n is bounded, and the proposition follows. \square

Now we show the necessity of the restriction on S .

Proposition 3.1.4. *Let $S \subseteq \mathbb{R}_{>0}$ be finite, and let $\mathcal{G} = (G, w)$ be an S -edge-labeled graph. Suppose there are $a, b \in S$ with $a < \max S$ such that $a \oplus_S b = a$. Then the class of finite S -metric spaces is not finitely constrained.*

Proof. Observe that $b \leq a$.

Consider the set \mathcal{F}_b of all cycles consisting of a path P consisting of a single repeated edge weight b and one additional edge of weight $\max(S)$.

By our assumption on a , by monotonicity, and by induction on the length of P , we find

$$w_S(P) \leq a < \max S.$$

Thus these cycles are non-metric. As there are infinitely many such cycles, if the class of finite S -metric spaces were finitely constrained then there would be one constraint embedding into infinitely many of the cycles in \mathcal{F}_b . In particular, some proper substructure of a cycle in \mathcal{F}_b would be forbidden. But a proper substructure of a cycle contains no cycle, and therefore contains no non-metric cycle. So by Lemma 3.1.6, no such proper substructure can be forbidden. \square

Thus we may finally deduce Proposition 3.1.2.

Proof of Proposition 3.1.2. This follows immediately from Propositions 3.1.3 and 3.1.4. \square

3.1.3 Metrically Homogeneous Graphs

Here we develop a completion algorithm for extending a partial substructure of a given metrically homogeneous graph Γ to an induced substructure of Γ . More precisely, this process associates a complete $[\delta]$ -edge-labeled graph to any $[\delta]$ -edge-labeled graph. The aim of this process is to assign an element of $\mathcal{A} = \text{age}(\Gamma)$ extending the given $[\delta]$ -edge-labeled graph, whenever one exists. For a broad class of metrically homogeneous graphs Γ , this succeeds. Analysis of this procedure shows that with few exceptions, associated to imprimitive cases, the minimal obstructions to success form a finite set. Thus whenever the completion process works as desired, the associated class of finite structures will be finitely constrained. Our main result is as follows.

Proposition 2. *For any primitive metrically homogeneous graph $\Gamma_{K_1, K_2, C, C', S}^\delta$ of generic type for which $C > 2\delta + \max(K_1, \delta/2)$, $C' = C + 1$, and $K_2 \geq \delta - 1$, the class $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$ is finitely constrained.*

Explicitly, by *finitely constrained*, we mean there exists some finite set of edge-labeled graphs \mathcal{F} such that the exclusion of this set is equivalent to having an extension to the relevant class of metric spaces. If an edge-labeled graph $\mathcal{G} = (G, w)$ can be completed to a metric space in $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$, then we refer to \mathcal{G} as a *partial \mathcal{A} -metric space*.

The assumption that a metrically homogeneous graph of generic type is primitive is equivalent to the assumption that $K_1 < \infty$ (not bipartite) and $C > 2\delta + 1$ (not antipodal). Our numerical assumptions actually imply these conditions, but we wish to stress that we are dealing here with a portion of the primitive case. Dealing with the imprimitive case generally requires more attention and a change of language. The bipartite analog of these results would require a change of language but little change in the arguments.

As noted above, a variation of this completion process gives a more general result applying to any primitive metrically homogeneous graph of generic type [ABH⁺17].

Proposition 2 has strong consequences for the dynamical and Ramsey theoretic properties of the associated amalgamation classes and automorphism groups, as seen in Sections 3.2 and 3.3. In addition, we show that the completion process has a certain *canonicity* property which provides additional information about the automorphism groups of these structures (see Lemma 3.2.1).

We give two proofs of Proposition 2. The first proof shows the constraint set is finite, but gives little information about it. This proof is based mainly on an analysis of the completion process defined and analyzed in Section 3.1.3.3. The second proof gives an entirely explicit description of the constraint set, which is the subject of Section 3.1.3.1.

Our less explicit proof was our original argument, and the more general proof from [HKK17] was based on a similar quantitative analysis of a different completion process.

As per our discussion in Chapter 1, we continue the blanket assumption that $\delta \geq 3$.

3.1.3.1 The set of constraints \mathcal{F}

In this section, we introduce and study the constraint discussed in the introduction in connection with the more explicit form of Proposition 2. We first check that under very general conditions, this set will be finite (the exceptions being associated with the imprimitive cases). Then we will show that it is necessary to forbid \mathcal{F} under isometric embedding in order for a $[\delta]$ -edge-labeled graph \mathcal{G} to be able to be extended to a space in $\mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$. Sufficiency will be shown in Section 3.1.3.5 as part of Corollary 3.1.27.1, which comes after the proof of Proposition 2 in a more qualitative form, which relies only on the material of Section 3.1.3.3 and Corollary 3.1.11.1 to Lemma 3.1.11 below.

Apart from the elementary Corollary 3.1.11.1 below, none of the material in this subsection will be needed until the end of Section 3.1.3, as it comes in only to the proof of the sharper proof of Proposition 3.1.1 alluded to above. For Proposition 3.1.1 as stated, it suffices to begin with the completion process discussed in Section 3.1.3.3.

Definition 3.1.5. *Given finite numerical parameters δ, K_1, K_2, C , where $K_2 \geq \delta - 1$, and a family \mathcal{S} of $(1, \delta)$ -spaces, we define the set $\mathcal{F} = \mathcal{F}(\delta, K_1, K_2, C, \mathcal{S})$ of forbidden configurations to be the union of the following sets of $[\delta]$ -edge-labeled graphs:*

- *The minimal spaces in \mathcal{S} , and the triangle type $(1, \delta, \delta)$ if $K_2 = \delta - 1$;*
- *Cycles of odd perimeter less than $2K_1$;*
- *n -Cycles where*

$$w(e_1) + w(e_2) + \cdots + w(e_{2\ell+1}) > \ell(C - 1) + w(e_{2\ell+2}) + \cdots + w(e_n) \quad (3.1)$$

where the n edges are not necessarily ordered so that e_i and e_{i+1} are adjacent.

For brevity, we refer to the collection of cycles that satisfy Inequality (3.1) as \mathcal{F}_0 .

Note that when $\ell = 0$, the cycles which obey Inequality (3.1) are the non-metric cycles, in that they contradict the generalized form of the triangle inequality (and thus any completion would itself violate the triangle inequality).

We take up the question of finiteness of the constraint set \mathcal{F} , which will be settled by Lemma 3.1.10. We deal first with \mathcal{S} , which is a set of forbidden $(1, \delta)$ -spaces.

Lemma 3.1.7. *For $\delta \geq 3$, any antichain of finite $(1, \delta)$ -spaces is finite.*

Proof. We follow an argument sketched in [ACM16, page 10]: If $\delta \geq 3$, then $(1, \delta)$ -spaces consist of collections of cliques, with distinct cliques being at distance δ away from each other. These spaces therefore can be described up to isomorphism by multisets of integers, where each integer represents the size of a clique.

The embeddability relation between $(1, \delta)$ -spaces then corresponds to the following relation on multisets: $A \leq B$ if and only if there is a function $f: A \rightarrow B$ with $a \leq f(a)$; here we allow a function defined on a multiset to take different values at different occurrences of the same element. There is a similar relation on ordered sequences, namely: $\mathbf{a} \leq \mathbf{b}$ if and only if there is an order-preserving function ϕ on the indices such that $a_i \leq b_{\phi(i)}$ for all i . Higman's Lemma (Lemma 2.6.6) says that there is no infinite antichain of finite sequences of integers under this relation, which implies that there is no infinite antichain of finite multisets of integers under the corresponding relation. The lemma follows. \square

Since the minimal elements of any partially ordered set form an antichain, we have the following.

Corollary 3.1.7.1. *If \mathcal{S} is a set of $(1, \delta)$ -spaces then the set of minimal elements of \mathcal{S} is finite, up to isomorphism.*

Lemma 3.1.8. *Let K_1 and δ be finite. Then in a $[\delta]$ -edge-labeled graph $\mathcal{G} = (G, w)$, there can be only finitely many cycles (u_1, u_2, \dots, u_k) up to isomorphism whose perimeter is odd and less than $2K_1$.*

Proof. Indeed, there are up to isomorphism only finitely many cycles with perimeter less than $2K_1$. This is because each label is an integer at least 1. \square

Lemma 3.1.9. *When $C > 2\delta + 1$, there are finitely many forbidden configurations in \mathcal{F}_0 (as in Definition 3.1.5).*

Proof. Recall that \mathcal{F}_0 consists of finite cycles γ where

$$w(e_1) + \dots + w(e_{2\ell+1}) > \ell(C - 1) + w(e_{2\ell+2}) + \dots + w(e_n). \quad (3.2)$$

We argue first that ℓ is bounded.

Since each $w(e_i)$ is at most δ , there will be no possible such cycles when

$$\delta(2\ell + 1) \leq \ell(C - 1).$$

This is equivalent to the following restriction on ℓ :

$$\ell \leq \frac{\delta}{C - 2\delta - 1}.$$

Note that this denominator is well-defined, since $C > 2\delta + 1$.

As ℓ is bounded and the left-side of Inequality (3.2) contains $2\ell + 1$ terms, each bounded by δ , the left-hand side is bounded. Hence the number of terms on the right is also bounded. Thus there are finitely many possible configurations of this type. \square

Hence we have the following.

Lemma 3.1.10. *Given finite numerical parameters δ, K_1, K_2, C with $C > 2\delta + 1$ and $K_2 \geq \delta - 1$, and a set \mathcal{S} of $(1, \delta)$ -spaces, the associated family \mathcal{F} (Definition 3.1.5) is finite.*

Proof. Recall that \mathcal{F} is the union of three sets. The first such set is the set of minimal spaces in \mathcal{S} together one more element if $K_2 = \delta$. Thus by Corollary 3.1.7.1 this first set is finite.

The second set consists of cycles of odd perimeter less than $2K_1$, and we showed in Lemma 3.1.8 that this set is finite.

Finally, the third set is the set \mathcal{F}_0 . We showed in Lemma 3.1.9 that this set is finite.

Thus \mathcal{F} is itself finite. \square

3.1.3.2 Necessity

We will show in Proposition 3.1.5 that under mild conditions, the structures in $\mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$ do not contain any configurations in the corresponding constraint set \mathcal{F} . Thus, to characterize embeddability for this class, it is *necessary* to forbid the configurations in \mathcal{F} . The converse (sufficiency) will be proved under more limited conditions.

This result follows from the more precise Lemma 3.1.27, to be proven in Section 3.1.3.5, which reduces it in large part to an inequality established by a direct argument in the proof

of Proposition 2, via the proof of Lemma 3.1.27. As it plays no direct role in anything done subsequently, the proof given here may be omitted. However, Corollary 3.1.11.1 to Lemma 3.1.11 will itself be used in the proof of Proposition 2.

The more delicate Lemma 3.1.12 will not be used subsequently, but the calculation used in its proof recurs in the proof of Lemma 3.1.27.

Proposition 3.1.5. *Let $\mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$ be the amalgamation class of finite $[\delta]$ -metric spaces associated to admissible parameters*

$$(\delta, K_1, K_2, C, C', \mathcal{S})$$

with δ, K_1 , and C finite, and $C > 2\delta + 1$, $C' = C + 1$. Then no configuration in $\mathcal{F} = \mathcal{F}(\delta, K_1, K_2, C, \mathcal{S})$ embeds into a member of $\mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$.

We first show the necessity of forbidding cycles of odd perimeter less than $2K_1$.

We will use the following definition.

Definition 3.1.6. *We say that an edge-labeled graph $\mathcal{G} = (G, w)$ satisfies the generalized triangle inequality if for any finite collection of vertices (u_1, u_2, \dots, u_k) in G , we have*

$$w(u_1, u_k) \leq w(u_1, u_2) + w(u_2, u_3) + \dots + w(u_{k-1}, u_k)$$

whenever all these distances are defined.

We call cycles which violate the generalized triangle inequality *non-metric cycles*.

Lemma 3.1.11. *Let (u_1, u_2, \dots, u_k) be the vertices of a cycle γ in the $[\delta]$ -edge-labeled graph $\mathcal{G} = (G, w)$. Assume moreover that the perimeter of γ is odd and less than $2K_1$.*

Then any proper $[\delta]$ -edge-labeled extension of γ which satisfies the generalized triangle inequality will contain a cycle on strictly fewer vertices which also has odd perimeter less than $2K_1$.

Proof. Assume we have such an extension $\hat{\gamma}$ of γ . Let u_i, u_j be non-adjacent vertices of γ for which the weight $\hat{w}(u_i, u_j)$ is defined in $\hat{\gamma}$. Then each of the cycles $(u_i, u_{i+1}, \dots, u_j)$ and $(u_j, u_{j+1}, \dots, u_i)$ has perimeter less than $2K_1$ by the generalized triangle inequality. As exactly one of the paths $(u_i, u_{i+1}, \dots, u_j)$ and $(u_j, u_{j+1}, \dots, u_i)$ has odd length, we also have that exactly

one of the cycles $(u_i, u_{i+1}, \dots, u_j)$, $(u_j, u_{j+1}, \dots, u_i)$ has odd perimeter. Thus, we have our desired result. \square

Corollary 3.1.11.1. *Let K_1 be finite. Then any integer-valued metric space A which contains a cycle of odd perimeter less than $2K_1$, also contains a triangle of odd perimeter less than $2K_1$.*

Proof. Take a cycle of minimal odd perimeter embedding in A . By Lemma 3.1.11, this cycle is complete, and is therefore a triangle. \square

Lemma 3.1.12. *Let δ and C be finite, with $C > 2\delta + 1$. Then no configuration satisfying Inequality (3.1) of Definition 3.1.5 embeds into a $[\delta]$ -valued metric space (A, d) in which every triangle has perimeter strictly less than C .*

Proof. We suppose toward a contradiction that A contains a cycle γ whose edges may be indexed so that for some ℓ we have

$$w(e_1) + \dots + w(e_{2\ell+1}) > \ell(C - 1) + w(e_{2\ell+2}) + \dots + w(e_n). \quad (3.3)$$

We proceed by induction on ℓ . If $\ell = 0$, then (3.3) violates the generalized triangle inequality. So $\ell \geq 1$.

View γ as an oriented cycle, and each e_i as an oriented edge (u_i, v_i) . We may suppose that $e_1, e_2, \dots, e_{2\ell+1}$ have been enumerated so as to come in cyclic order around the cycle. Let γ_i be the path in γ joining e_i to e_{i+1} (computing modulo $2\ell + 1$). We note that γ_i consists solely of edges from $\{e_{2\ell+1}, \dots, e_n\}$.

See Figure 3.1 below for clarity.

By the triangle inequality, we have

$$d(v_{2\ell}, v_{2\ell+1}) \geq w(e_{2\ell+1}) - \sum_{e \in \gamma_{2\ell}} w(e).$$

So if we replace the path from $v_{2\ell}$ to $v_{2\ell+1}$ in γ by a single edge $e'_{2\ell+1} = (v_{2\ell}, v_{2\ell+1})$ to get a new cycle γ' , then we get a corresponding inequality to (3.3)

$$w(e_1) + w(e_2) + \dots + w(e_\ell) + d(v_{2\ell}, v_{2\ell+1}) \geq \ell(C - 1) + \sum_{\substack{e \in \gamma_i \\ i \in \{1, \dots, 2\ell+1\} \setminus 2\ell}} w(e). \quad (3.4)$$




Figure 3.1: Lemma 3.1.12

This inequality has fewer terms on the right than that of Inequality (3.3). Note that $e_{2\ell}$ and $e'_{2\ell+1}$ are adjacent edges. Then applying the perimeter bound to the triangle with edges $e_{2\ell}$, $e'_{2\ell+1}$, and $e'' = (u_{2\ell}, v_{2\ell+1})$ we have

$$w(e_{2\ell}) + w(e'_{2\ell+1}) \leq (C - 1) - w(e'').$$

So we may replace the terms $w(e_{2\ell}) + w(e'_{2\ell+1})$ on the left side of Inequality (3.4) by $C - 1 - w(e'')$, and then, after cancelling $C - 1$ and moving $w(e)$ to the right, we have the inequality corresponding to the cycle γ'' which is obtained from γ' by replacing $e_{2\ell}$ and $e'_{2\ell+1}$ by e'' .

Then ℓ is reduced, and we conclude by induction. \square

We may now deduce Proposition 3.1.5.

Proof of Proposition 3.1.5. Lemma 3.1.12 and Corollary 3.1.11.1 show that the second and third families of configurations (those relating to K_1 and those referred to as \mathcal{F}_0) are forbidden.

Let A belong to $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$. By definition, no configuration in \mathcal{S} embeds into A , and if $K_2 = \delta - 1$, the same applies to triangles of type $(1, \delta, \delta)$. By Corollary 3.1.11.1 and Lemma 3.1.12 none of the other configurations in \mathcal{F} embeds into A .

The result follows. \square

3.1.3.3 A completion process

In the present section, we will define what we mean by a ‘‘candidate configuration’’ for completion to a structure in the class $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$; define a completion process for such candidates;

and show that under suitable numerical conditions, candidates for completion may in fact be completed by our process.

The notion of candidate for completion, given in Definition 3.1.8, involves omitting some of the more natural constraints discussed in the previous section, as well as satisfying an additional condition which is a prerequisite for the sensible application of our completion process, which relies on the preliminary definitions of Section 3.1.3.3 below.

The main result of this section will then be the following.

Proposition 3.1.6. *Let $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$ be the amalgamation class of finite $[\delta]$ -metric spaces associated to admissible parameters*

$$(\delta, K_1, K_2, C, C', S)$$

with δ, K_1 , and C finite, satisfying the following conditions.

- $C > 2\delta + 1$;
- $C' = C + 1$;
- $K_2 \geq \delta - 1$.

If a $[\delta]$ -edge-labeled graph $\mathcal{G} = (G, w)$ is a candidate configuration for completion to $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$, then the completion $\overline{\mathcal{G}}$ of \mathcal{G} using the completion algorithm in Definition 3.1.9 yields a metric space in $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$.

Towards the completion process

As preparation for the completion process given in Definition 3.1.9 below, we now introduce two weight functions, ρ^- and ρ^+ , associated with any parameters δ, C , and any $[\delta]$ -edge-labeled graph \mathcal{G} . These functions may be defined informally as natural lower and upper bounds for the distance in any extension of \mathcal{G} to a $[\delta]$ -metric space with perimeter bound C . In particular, our completion process will attempt to assign a weight to the edge (u, v) which lies in the interval $[\rho^-(u, v), \rho^+(u, v)]$. Of course, if $\rho^-(u, v) > \rho^+(u, v)$, it will not succeed, but it will assign a value anyway.

Definition 3.1.7. *Let $\mathcal{G} = (G, w)$ be a $[\delta]$ -edge-labeled graph and C a numerical parameter. We*

define weight functions $\rho_i^\pm = \rho_{i,\mathcal{G},C}^\pm$ for the complete graph on G by induction on i as follows:

$$\rho_0^-(u, v) = \begin{cases} w(u, v) & \text{if defined} \\ 1 & \text{otherwise} \end{cases} \quad \rho_0^+(u, v) = \begin{cases} w(u, v) & \text{if defined} \\ \delta & \text{otherwise} \end{cases}$$

$$\rho_{i+1}^-(u, v) = \max(\rho_i^-(u, v), \max_w(\rho_i^-(u, w) - \rho_i^+(v, w)), \max_w(\rho_i^-(v, w) - \rho_i^+(u, w))),$$

$$\rho_{i+1}^+(u, v) = \min(\rho_i^+(u, v), \min_w(\rho_i^+(u, w) + \rho_i^+(v, w)),$$

$$\min_w(C - 1 - (\rho_i^-(u, w) + \rho_i^-(v, w)))).$$

Notice that the sequences of functions (ρ_i^-) , (ρ_i^+) are monotonic (non-decreasing and non-increasing, respectively), integer-valued, and uniformly bounded above and below, and hence eventually constant (pointwise). So we may define limit functions

$$\rho^-(u, v) = \lim_{i \rightarrow \infty} \rho_i^-(u, v)$$

and

$$\rho^+(u, v) = \lim_{i \rightarrow \infty} \rho_i^+(u, v).$$

We show first that once the parameter δ is fixed, the limit values are attained in a uniformly bounded number of iterations, specifically: $2\delta - 1$.

Lemma 3.1.13. *Let the numerical parameters δ and C be fixed, with $C > 2\delta + 1$.*

Then for any $[\delta]$ -edge-labeled graph $\mathcal{G} = (G, w)$, we have

$$\rho^- = \rho_{2\delta-1}^- \quad \rho^+ = \rho_{2\delta-1}^+.$$

Proof of Lemma 3.1.13. We make use of the following inductive definitions, which keep track of the places where the assigned values of our weight functions change.

$$R_0^- = \{\rho_0^-(u, v) \mid u, v \in G, u \neq v\}$$

$$R_0^+ = \{\rho_0^+(u, v) \mid u, v \in G, u \neq v\}$$

$$R_{i+1}^- = \{\rho_{i+1}^-(u, v) \mid u, v \in G, u \neq v, \rho_{i+1}^-(u, v) > \rho_i^-(u, v)\}$$

$$R_{i+1}^+ = \{\rho_{i+1}^+(u, v) \mid u, v \in G, u \neq v, \rho_{i+1}^+(u, v) < \rho_i^+(u, v)\}$$

We note that Lemma 3.1.13 is shown if we can find a single $i \leq 2\delta$ where $R_i^- = \emptyset$ and likewise $R_i^+ = \emptyset$. Indeed, if that were to happen, then both $\rho_i^- = \rho_{i-1}^-$ and $\rho_i^+ = \rho_{i-1}^+$, so all subsequent ρ_j^-, ρ_j^+ agree with $\rho_{i-1}^-, \rho_{i-1}^+$ respectively.

We shift then to showing the following claim.

Claim 3.1.13.1. *For any i , we have:*

$$\begin{aligned} \sup R_{2i}^- &\leq \delta - i & \inf R_{2i}^+ &\geq i + 1 \\ \sup R_{2i+1}^- &\leq \delta - i - 1 & \inf R_{2i+1}^+ &\geq i + 1 \end{aligned}$$

Once we have shown this claim, it then follows that $R_{2\delta}^\pm$ is empty. Indeed, we would have that $\sup R_{2\delta}^- \leq \delta - \delta = 0$, and $\inf R_{2\delta}^+ \geq \delta + 1$. Thus, once we prove our claim, we will also be done with our proof.

Proof of Claim 3.1.13.1. We proceed by induction over i .

It follows immediately that $\sup R_0^- \leq \delta$ because our partial space is of bounded diameter δ , and that $\inf R_0^+ \geq 1$, since we are only considering distinct points $u, v \in G$. We also get immediately that $\inf R_1^+ \geq 1$, since $C - 2\delta - 1 \geq 1$.

We now assume for $k \leq i$ that:

$$\begin{aligned} \sup R_{2k}^- &\leq \delta - k \\ \inf R_{2k}^+ &\geq k + 1 \\ \inf R_{2k+1}^+ &\geq k + 1. \end{aligned}$$

We wish to show that it then follows that $\sup R_{2k+1}^- \geq \delta - k - 1$, and that the three inequalities above are valid for $k = i + 1$ as well.

Subclaim. $\sup R_{2k+1}^- \leq \delta - k - 1$.

By the inductive definition of ρ_i^- , if $\rho_{2k+1}^-(u, v)$ lies in R_i^- , then it has the form

$$\rho_{2k}^-(u_1, w) - \rho_{2k}^+(u_2, w)$$

with u_1, u_2 equal to u, v in some order. Furthermore, either the first term is in R_{2k}^- or the second term is in R_{2k}^+ , as otherwise ρ_{2k}^- would already have reached this value.

Accordingly,

$$\begin{aligned}\rho_{2k+1}^-(u, v) &\leq \max(\sup R_{2k}^- - 1, \delta - \inf R_{2k}^+) \\ &\leq \max(\delta - k - 1, \delta - k - 1) = \delta - k - 1,\end{aligned}$$

by the inductive hypothesis, and the subclaim is proved.

Similarly, if $\rho_{2k+2}^-(u, v) \in R_{2k+1}^-$, then we find by the same calculation

$$\rho_{2k+2}^-(u, v) \leq \max(\sup R_{2k+1}^-, \delta - \inf R_{2k+1}^+) = \delta - k - 1$$

and thus $\sup R_{2k+1}^- \leq \delta - k - 1$, which is one of the desired inequalities for $k = i + 1$.

It remains to deal with the two inequalities concerning R_{2k+2}^+ and R_{2k+3}^+ :

$$\inf R_{2k+2}^+, \inf R_{2k+3}^+ \geq k + 2.$$

There are two ways a value $\rho_{2k+2}^+(u, v)$ could get into R_{2k+2}^+ . Either

$$\rho_{2k+2}^+(u, v) = \rho_{2k+1}^+(u, w) + \rho_{2k+1}^+(v, w),$$

with one of the entries on the right in R_{2k+1}^+ , or

$$\rho_{2k+2}^+(u, v) = C - (\rho_{2k+1}^-(u, w) + \rho_{2k+1}^-(v, w)) - 1,$$

with one of the two entries on the right in R_{2k+1}^- .

In the first case, we find at once that $\rho_{2k+2}^+(u, v) \geq k + 2$.

In the second case, we find

$$\rho_{2k+2}^+(u, v) \leq C - (\delta - k - 1 + \delta) - 1 = C - (2\delta - k) \geq k + 2$$

since $C > 2\delta + 1$.

This takes care of the inequality for R_{2k+2}^+ , and the argument for R_{2k+3}^+ is almost exactly the same.

This proves Claim [3.1.13.1](#). □

By our earlier remarks, the lemma follows from the claim. □

The following result will prove useful.

Lemma 3.1.14. *Let δ and C be finite, with $C > 2\delta + \delta/2$, and let $\mathcal{G} = (G, w)$ be a $[\delta]$ -edge-labeled graph. Then*

$$\rho^+(x, y) \leq \rho^+(x, w) + \rho^+(y, w) \quad \text{and} \quad \rho^-(x, y) \leq \rho^-(x, w) + \rho^-(y, w)$$

for any $x, y, w \in A$. Moreover,

$$\rho^+(x_1, x_n) \leq \sum_{i=1}^{n-1} \rho^+(x_i, x_{i+1}) \quad (3.5)$$

for any $x_1, \dots, x_n \in G$.

Proof. By definition of ρ_i^\pm , we have

$$\rho_{i+1}^+(x, y) \leq \rho_i^+(x, w) + \rho_i^+(y, w) \quad \rho_{i+1}^-(x, w) \geq \rho_i^-(x, y) - \rho_i^+(y, w).$$

Taking the limit over i gives the first two inequalities.

Inequality (3.5) follows by induction. \square

Lemma 3.1.15. *With δ and C fixed, there is a finite set of configurations \mathcal{F}_1 , such that a $[\delta]$ -edge-labeled graph $\mathcal{G} = (G, w)$ satisfies the condition $\rho^- \leq \rho^+$ everywhere if and only if \mathcal{G} omits all configurations in \mathcal{F}_1 .*

Proof. More precisely, we claim the following.

Claim. For any edge e in \mathcal{G} , there are two configurations γ^-, γ^+ contained in \mathcal{G} , and containing the vertices of e , with the following properties.

- In any $[\delta]$ -edge-labeled graph \mathcal{G}' containing γ^- , we have $\rho_{\mathcal{G}'}^-(e) \geq \rho_{\mathcal{G}}^-(e)$.
- In any $[\delta]$ -edge-labeled graph \mathcal{G}'' containing γ^+ , we have $\rho_{\mathcal{G}''}^+(e) \leq \rho_{\mathcal{G}}^+(e)$.
- γ^- and γ^+ each contain at most $2^{2\delta-1}$ vertices.

As we have noted in the course of the previous proof, a lower bound for $\rho_i^-(u, v)$ or an upper bound for $\rho_i^+(u, v)$ must have a witness consisting of a third vertex w together with a corresponding upper or lower bound for $\rho^\pm(u, w)$ and $\rho^\pm(v, w)$.

Since the value of $\rho^\pm(u, v)$ is given by $\rho_{2\delta-1}^\pm$, tracing backward we can replace the given edge by two edges with a common third vertex, then by at most four edges, and so forth, together

with estimates on $\rho_{2\delta-1-i}^\pm$ for these edges, for $i = 0, 1, \dots, 2\delta - 1$, concluding when we reach the original edges of \mathcal{G} , which then make up the desired configurations.

The claim follows.

Now the lemma is immediate: the configurations of the form $\gamma^- \cup \gamma^+$ as above, ranging over all possible $[\delta]$ -edge-labeled graphs \mathcal{G} , provide the required set \mathcal{F}_1 . This set is finite, since the number of vertices occurring is bounded by $2 \cdot 2^{2\delta-1} = 2^{2\delta}$, and there are only finitely many such configurations for a fixed δ . \square

A sharper analysis, Lemma 3.1.27, will show that the set \mathcal{F}_1 (which is a set of circuits) can be replaced by the set \mathcal{F}_0 of cycles (and even that \mathcal{F}_0 is the set of minimal elements of \mathcal{F}_1).

Lemma 3.1.16. *Let δ and C be finite. Let $\mathcal{G} = (G, w)$ be a $[\delta]$ -edge-labeled graph such that $\rho^- \leq \rho^+$ on G .*

Then the weights ρ^- and ρ^+ extend the weight w given on G .

Proof. By construction, we have

$$\rho^+(u, v) \leq \rho_0^+(u, v) = w(u, v) \qquad \rho^-(u, v) \geq \rho_0^-(u, v) = w(u, v)$$

and so

$$\rho^+(u, v) \leq w(u, v) \leq \rho^-(u, v).$$

As $\rho^- \leq \rho^+$, we conclude $\rho^+(u, v) = w(u, v) = \rho^-(u, v)$. \square

Corollary 3.1.16.1. *Let δ and C be finite. Let $\mathcal{G} = (G, w)$ be a $[\delta]$ -edge-labeled graph satisfying*

$$\rho^- \leq \rho^+.$$

Then \mathcal{G} contains no non-metric cycle.

Proof. By Lemma 3.1.14 (3.5), (G, ρ^+) contains no non-metric cycle. By the previous lemma, ρ^+ agrees with w when the latter is defined. \square

We will make copious use of the following definition.

Definition 3.1.8. A finite $[\delta]$ -edge-labeled graph $\mathcal{G} = (G, w)$ will be called a candidate configuration for completion to $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$, a class of $[\delta]$ -metric spaces associated to admissible parameters

$$(\delta, K_1, K_2, C, C', S)$$

with δ, K_1 , and C finite, and $C > 2\delta + 1$ if

1. \mathcal{G} satisfies $\rho^- \leq \rho^+$ everywhere;
2. None of the spaces in S embed into \mathcal{G} ;
3. If $K_2 = \delta - 1$, then the triangle type $(1, \delta, \delta)$ also does not embed into \mathcal{G} .
4. No cycles of odd perimeter less than $2K_1$ embed into S .

We note that candidate configurations will not contain any non-metric cycles (Corollary 3.1.16.1).

Remark 3.1.2. By Lemma 3.1.15, a configuration which does not embed any of the configurations in \mathcal{F} is a candidate configuration.

Furthermore, we have the following.

Lemma 3.1.17. Let $\Gamma_{K_1, K_2, C, C', S}^\delta$ be a primitive metrically homogeneous graph of generic type for which $C > 2\delta + \max(K_1, \delta/2)$, $C' = C + 1$, and $K_2 \geq \delta - 1$.

Then configurations which embed into an element of $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$ are candidate configurations for completion to $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$.

Proof. Since the set of candidate configurations is closed under taking substructures, it suffices to check that the structures in $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$ are candidate configurations.

Let (A, d) be in $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$. By definition, A contains no forbidden Henson constraint or forbidden triangle, so clauses (2, 3) of Definition 3.1.8 are satisfied.

From the inductive definition of ρ^\pm and the condition $C' = C + 1$ it follows at once that

$$\rho_k^- \leq d \leq \rho_k^+$$

for all k . (The condition $C' = C + 1$ means that no triangles of perimeter greater than $C - 1$ can occur.)

Finally, Corollary 3.1.11.1 to Lemma 3.1.11 shows that no cycle of odd perimeter less than $2K_1$ can occur.

This proves the lemma. \square

Our main task will be to show, conversely, that any candidate configuration does embed in a member of $\mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$ when the parameters satisfy some additional constraints.

We are now ready to give our completion algorithm.

The completion algorithm

Definition 3.1.9. Let $\mathcal{G} = (G, w)$ be a candidate configuration for completion to the class of finite metric spaces $\mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$ which correspond to a primitive metrically homogeneous graph Γ of generic type. Then the completion algorithm for \mathcal{G} is to:

1. Assign the value of ρ^+ to all unlabeled edges (u, v) for which $\rho^+(u, v) < C - 2\delta - 1$.

We write this extended edge-labeled graph as $\hat{\mathcal{G}} = (G, \hat{w})$.

2. Recalculate the values ρ^-, ρ^+ in $\hat{\mathcal{G}}$. We denote these new values by $\hat{\rho}^-$ and $\hat{\rho}^+$.
3. Assign the value of $\max(\hat{\rho}^-, (C - 2\delta - 1)')$ to the remaining unlabeled edges, where $(C - 2\delta - 1)' = \min(C - 2\delta - 1, \delta - 1)$, unless \mathcal{S} is empty, in which case $(C - 2\delta - 1)' = C - 2\delta - 1$.

Thus we begin to fill in edges with the first step, in which we use as many ρ^+ values as possible.

The completion process pauses here in step 2, as we recalculate ρ^- and ρ^+ .

The second step of the completion algorithm is necessary, as the ρ^- and ρ^+ values may have changed, but we note that this recalculation only happens the one time. This raises the question as to whether the new weight functions $\hat{\rho}^-$ and $\hat{\rho}^+$ also satisfy the desirable inequality $\hat{\rho}^- \leq \hat{\rho}^+$. Indeed, it is possible that after the first step of the completion algorithm, the ρ^- terms have grown, that is that $\hat{\rho}^- > \rho^-$. We will show that under suitable restrictions on the parameters, the ρ^+ values will not shrink, and that despite potential ρ^- growth, the desired inequality is maintained. In fact, the configuration produced by Step 1 will again be a candidate configuration (Lemma 3.1.22).

We then recommence the completion process in step 3, now using $\hat{\rho}^-$ and $(C - 2\delta - 1)'$.

We introduce the parameter $(C - 2\delta - 1)'$ in order to ensure that the completion $\bar{\mathcal{G}}$ satisfies the Henson constraints (as will be seen in Lemmas 3.1.19 and 3.1.25).

We note that the assumption $K_2 \geq (\delta - 1)$ ensures that the only triangle type that might be forbidden based on the K_2 value is $(1, \delta, \delta)$. Since this a $(1, \delta)$ -space, arguments for its exclusion can be grouped together with those for \mathcal{S} .

3.1.3.4 The completion of a candidate configuration

In this section, we will show the following.

Proposition 3.1.6. *Let $\mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$ be the amalgamation class of finite $[\delta]$ -metric spaces associated to admissible parameters*

$$(\delta, K_1, K_2, C, C', \mathcal{S})$$

with δ, K_1 , and C finite, satisfying the following conditions.

- $C > 2\delta + \max(K_1, \delta/2)$;
- $C' = C + 1$;
- $K_2 \geq \delta - 1$.

If a $[\delta]$ -edge-labeled graph $\mathcal{G} = (G, w)$ is a candidate configuration for completion to $\mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$, then the completion $\bar{\mathcal{G}}$ of \mathcal{G} using the completion algorithm in Definition 3.1.9 yields a metric space in $\mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$.

Our process in proving Proposition 3.1.6 is to show the following.

- a. After Step 1, the extension $\hat{\mathcal{G}}$ satisfies conditions (2–4) of Definition 3.1.8.
- b. The inequality $\hat{\rho}^- \leq \hat{\rho}^+$ holds.
- c. The final step of the completion algorithm (Definition 3.1.9) produces a $[\delta]$ -metric space $\bar{\mathcal{G}}$ which again satisfies conditions (2–4) of Definition 3.1.8, as well as the perimeter bound corresponding to C, C' , and is therefore in $\mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$.

Under the hypotheses of Proposition 3.1.6, all three points can be carried through. We will be more precise about the hypotheses actually used at various stages of the analysis.

We take up the first point. We begin by showing that no K_1 violations have been introduced.

Lemma 3.1.18. *Let $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$ be the amalgamation class of finite $[\delta]$ -metric spaces associated to admissible parameters*

$$(\delta, K_1, K_2, C, C', S)$$

with δ, K_1 , and C finite, and $C > 2\delta + \max(K_1, \delta/2)$. Let $\mathcal{G} = (G, w)$ be a candidate configuration for completion to $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$.

Suppose $\hat{\mathcal{G}} = (G, \hat{w})$ is the extension of \mathcal{G} obtained after Step 1 of the completion algorithm. Then $\hat{\mathcal{G}}$ does not contain any cycles of odd perimeter less than $2K_1$.

Proof. Note that since \mathcal{G} is a candidate configuration, it does not itself contain any cycles of odd perimeter less than $2K_1$. Thus we restrict our attention to a cycle γ in (G, \hat{w}) which contains newly labeled edges (u, v) with labels $\rho^+(u, v)$.

When $\rho^+(u, v) \geq K_1$, we know that any cycle of odd perimeter which contains the edge (u, v) must have perimeter at least $2K_1$, since by Lemma 3.1.16 the edges in $\hat{\mathcal{G}}$ are in fact labeled with ρ^+ , and by Lemma 3.1.14 (3.5), ρ^+ satisfies the generalized triangle inequality.

We assume then that any newly labeled edge (u, v) occurring in γ satisfies

$$\rho^+(u, v) < K_1.$$

Recall that for $j \geq 0$, $\rho_j^+(u, v)$ is the minimum of the following three quantities:

$$\begin{aligned} & \rho_{j-1}^+(u, v), \\ & \min_w (\rho_{j-1}^+(u, w) + \rho_{j-1}^+(v, w)), \\ & \min_w (C - (\rho_{j-1}^-(u, w) + \rho_{j-1}^-(v, w)) - 1). \end{aligned}$$

Let j be minimal such that $\rho^+(u, v) = \rho_j^+(u, v)$. If $j = 0$, then as (u, v) is a new edge, we have that $\rho^+(u, v) = \delta \geq K_1$, a contradiction. So $j > 0$.

Now $\rho_j^+(u, v)$ cannot be of the form $(C - (\rho_{j-1}^-(u, w) + \rho_{j-1}^-(v, w)) - 1)$, since

$$(C - (\rho_{j-1}^-(u, w) + \rho_{j-1}^-(v, w)) - 1) \geq (2\delta + K_1 + 1) - 2\delta - 1 = K_1.$$

Thus we may assume that

$$\rho_j^+(u, v) = \rho_{j-1}^+(u, w) + \rho_{j-1}^+(v, w).$$

Similarly, if (u, w) is a new edge, then the term $\rho_{j-1}^+(u, w)$ may be replaced by $\rho_k^+(u, w') + \rho_k^+(w', v)$, for some $k < j - 1$. Continuing in this fashion, one ultimately expands $\rho^+(u, v)$ into the sum along a walk (u_1, \dots, u_n) from u to v with edges in \mathcal{G} :

$$\rho^+(u, v) = \sum_{i < n} w(u_i, u_{i+1}).$$

Replacing each edge (u, v) of γ not in \mathcal{G} by such a walk, one obtains a circuit γ' in \mathcal{G} with the same perimeter as γ . But then this is a union of cycles with total perimeter the perimeter of γ . Thus if γ has odd perimeter less than $2K_1$, then at least one of these cycles has odd perimeter less than $2K_1$.

Hence extending to $\hat{\mathcal{G}}$ cannot have introduced any cycles of odd perimeter less than $2K_1$. \square

Lemma 3.1.19. *Let $\mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$ be the amalgamation class of finite $[\delta]$ -metric spaces associated to admissible parameters*

$$(\delta, K_1, K_2, C, C', \mathcal{S})$$

with δ, K_1 , and C finite, and $C > 2\delta + K_1$. Let $\mathcal{G} = (G, w)$ be a candidate configuration for completion to $\mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$.

Suppose $\hat{\mathcal{G}} = (G, \hat{w})$ is the extension of \mathcal{G} obtained after Step 1 of the completion algorithm. Then the space $\hat{\mathcal{G}}$ will not contain any isometric copies of the forbidden Henson spaces in \mathcal{S} , or the triangle type $(1, \delta, \delta)$, in the case that $K_2 = \delta - 1$.

Proof. We deal first with \mathcal{S} .

Since \mathcal{G} is a candidate configuration, none of the Henson spaces in \mathcal{S} embed into \mathcal{G} . Thus we must ensure that extending by ρ^+ will not introduce any such Henson spaces. As these spaces are $(1, \delta)$ -spaces, we consider when a newly assigned value $\rho^+(u, v)$ could introduce the value 1 or δ .

As $C \leq 3\delta + 1$ and $\rho^+(u, v) < C - 2\delta - 1 \leq \delta$, we do not have that $\rho^+(u, v) = \delta$.

In particular as the edge (u, v) is not in \mathcal{G} and $\rho^+(u, v) \neq \delta$, the minimal j for which $\rho^+(u, v) = \rho_j^+(u, v)$ is nonzero, and thus $\rho^+(u, v)$ has one of the two forms

$$\begin{aligned} \rho_{j-1}^+(u, w) + \rho_{j-1}^+(w, v) &\geq 2 \\ C - (\rho_{j-1}^-(u, w) + \rho_{j-1}^-(w, v)) - 1 &\geq C - 2\delta - 1. \end{aligned}$$

So if $\rho^+(u, v) = 1$ then $C = 2\delta + 2$. Then by definition of admissibility (Definition 2.2.10), we have that \mathcal{S} is empty, and no violation is possible.

Thus, either ρ^+ does not introduce the values 1 or δ (and therefore cannot introduce any violations of Henson constraints), or there are no Henson constraints.

Similarly, if $K_2 = \delta - 1$ and a triangle of type $(1, \delta, \delta)$ appears in $\hat{\mathcal{G}}$, then the corresponding value $\rho^+(u, v)$ must be 1, and $C = 2\delta + 2$. However it follows from the definition of admissibility that for $C > 2\delta + K_1$ and $K_2 = \delta - 1$, we must have that $C > 2\delta + 2$ (Table 2.1). \square

Lemma 3.1.20. *Let $\mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$ be the amalgamation class of finite $[\delta]$ -metric spaces associated to admissible parameters*

$$(\delta, K_1, K_2, C, C', \mathcal{S})$$

with δ, K_1 , and C finite.

Let $\mathcal{G} = (G, w)$ be a candidate configuration for completion to $\mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$. Suppose that (u, v) is a non-edge in \mathcal{G} with $\rho^+(u, v) < C - 2\delta - 1$. Let $\check{\mathcal{G}} = (G, \check{w})$ be the extension of \mathcal{G} in which the edge (u, v) is added, and is labeled with $\rho^+(u, v)$. Then the functions $\check{\rho}^\pm$ associated with $\check{\mathcal{G}}$ satisfy

$$\check{\rho}^+ = \rho^+ \tag{3.6}$$

$$\rho^-(x, y) \leq \check{\rho}^-(x, y) \leq \max(\rho^-(x, y), C - 2\delta - 1) \tag{3.7}$$

for all $x, y \in G$.

Proof. We will prove the following two statements for all k , by induction on k :

$$\rho^+(x, y) \leq \check{\rho}_k^+(x, y) \leq \rho_k^+(x, y) \tag{3.8}$$

$$\rho_k^-(x, y) \leq \check{\rho}_k^-(x, y) \leq \max(\rho^-(x, y), C - 2\delta - 2). \tag{3.9}$$

Then passing to the limit proves the lemma.

The base case is

$$k = 0.$$

Then $\check{\rho}_0^\pm(x, y) = \rho_0^\pm(x, y)$ unless $(x, y) = (u, v)$ (in some order). When $(x, y) \neq (u, v)$, then (3.8) is clear for $k = 0$, as is (3.9), since in this case $\rho_0^-(u, v) = r \leq C - 2\delta - 2$.

On the other hand, (3.8) is equally clear for $k = 0$ when applied to the edge (u, v) in $\check{\mathcal{G}}$, and (3.9) for $k = 0$, since $\check{\rho}_0^-(u, v) = r$.

This disposes of the case $k = 0$. Now we proceed inductively, passing from $(k - 1)$ to k .

We begin with (3.8) for k .

By definition, $\check{\rho}_k^+(x, y)$ is the minimum of terms of the following forms:

- a. $\check{\rho}_{k-1}^+(x, y)$
- b. $\check{\rho}_{k-1}^+(x, z) + \check{\rho}_{k-1}^+(y, z)$
- c. $C - (\check{\rho}_{k-1}^-(x, z) + \check{\rho}_{k-1}^-(y, z)) - 1$.

Applying the upper and lower bounds from (3.8, 3.9), respectively, for $k - 1$, shows that each of the corresponding terms in the definition of $\rho_k^+(x, y)$ dominates a term shown here, so $\check{\rho}_k^+(x, y) \leq \rho_k^+(x, y)$.

We must also show that each of the terms of types (a-c) is dominated by $\rho^+(x, y)$. Applying induction together with Lemma 3.1.14 (3.5) gives the required estimate for terms of types (a, b).

Now we come to the critical case (c), where the term has the form

$$C - (\check{\rho}_{k-1}^-(x, z) + \check{\rho}_{k-1}^-(y, z)) - 1.$$

According to our inductive hypothesis (3.9) for $k - 1$, one of the following bounds applies:

$$\begin{aligned} \check{\rho}_{k-1}^-(x, z) &\leq \rho^-(x, z) \\ \check{\rho}_{k-1}^-(x, z) &\leq C - 2\delta - 2 \end{aligned}$$

with a similar bound applying to $\check{\rho}_{k-1}^-(y, z)$.

If we have both $\check{\rho}_{k-1}^-(x, z) \leq \rho^-(x, z)$ and $\check{\rho}_{k-1}^-(y, z) \leq \rho^-(y, z)$, then these bounds together yield

$$\begin{aligned} C - (\check{\rho}_{k-1}^-(x, z) + \check{\rho}_{k-1}^-(y, z)) - 1 &\geq C - (\rho^-(x, z) + \rho^-(y, z)) - 1 \\ &\geq \rho^+(x, y) \end{aligned}$$

as required.

So we may suppose that a bound of the second type applies to $\check{\rho}_{k-1}^-(x, z)$,

$$\check{\rho}_{k-1}^-(x, z) \leq C - 2\delta - 2.$$

In this case we have the following

$$\begin{aligned} C - (\check{\rho}_{k-1}^-(x, z) + \check{\rho}_{k-1}^-(y, z)) - 1 &\geq C - [(C - 2\delta - 2) + \delta] - 1 = \delta + 1 \\ &\geq \rho^+(x, y). \end{aligned}$$

Thus we have the first inequality for (3.8) for all cases and (3.8) follows for k .

Now we deal with (3.9) for k , namely:

$$\rho_k^-(x, y) \leq \check{\rho}_k^-(x, y) \leq \max(\rho^-(x, y), C - 2\delta - 2).$$

By definition, $\check{\rho}_{k-1}^-(x, y)$ is the maximum of terms of the forms

$$(a) \check{\rho}_{k-1}^-(x, y)$$

$$(b) \check{\rho}_{k-1}^-(x, z) - \check{\rho}_{k-1}^+(y, z)$$

$$(b') \check{\rho}_{k-1}^-(y, z) - \check{\rho}_{k-1}^+(x, z).$$

Now $\rho_k^-(x, y)$ is given by a similar formula, and each of the terms of types (a, b, b') dominates the corresponding term without the check, by our induction hypothesis. So the first inequality in (3.9) is immediate. We turn to the second inequality in (3.9).

The term of type (a) is covered by our induction hypothesis

$$\check{\rho}_{k-1}^-(x, y) \leq \max(\rho^-(x, y), C - 2\delta - 2).$$

So it will suffice now to show that terms of type (b) are bounded above by $\max(\rho^-(x, y), C - 2\delta - 2)$. These terms have the following form:

$$\check{\rho}_{k-1}^-(x, z) - \check{\rho}_{k-1}^+(y, z).$$

By induction, we have one of the following:

$$\check{\rho}_{k-1}^-(x, z) \leq \rho^-(x, z)$$

$$\check{\rho}_{k-1}^-(x, z) \leq C - 2\delta - 2.$$

In the second case, the whole term (b) is bounded by $C - 2\delta - 2$. In the first case, we bound it as follows,

$$\check{\rho}_{k-1}^-(x, z) - \check{\rho}_{k-1}^+(y, z) \leq \rho^-(x, z) - \rho^+(y, z) \leq \rho^-(x, y),$$

to conclude.

This completes the inductive proof of (3.9), and Lemma 3.1.20 follows. \square

Lemma 3.1.21. Let $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$ be the amalgamation class of finite $[\delta]$ -metric spaces associated to admissible parameters

$$(\delta, K_1, K_2, C, C', S)$$

with δ, K_1 , and C finite.

Let $\mathcal{G} = (G, w)$ be a candidate configuration for completion to $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$. Suppose that (u, v) is a non-edge in \mathcal{G} with $\rho^+(u, v) < C - 2\delta - 1$. Let $\check{\mathcal{G}} = (G, \check{w})$ be the extension of \mathcal{G} in which the edge (u, v) is added, and is labeled with $\rho^+(u, v)$. Then the functions $\check{\rho}^\pm$ associated with $\check{\mathcal{G}}$ satisfy

$$\check{\rho}^- \leq \check{\rho}^+. \quad (3.10)$$

Proof. We first strengthen the upper bound for $\check{\rho}^-$ from Lemma 3.1.20.

Claim 3.1.21.1. For all $x, y \in G$ we have

$$\begin{aligned} \check{\rho}^-(x, y) \leq \max(\rho^-(x, y), r - (\rho^+(x, u) + \rho^+(y, v)), \\ r - (\rho^+(x, v) + \rho^+(y, u))) \end{aligned} \quad (3.11)$$

where we adopt the convention that $\rho^+(u, u) = \rho^+(v, v) = 0$, when x or y is equal to u or v in

Proof of Claim 3.1.21.1. We prove inductively

$$\begin{aligned} \check{\rho}_k^-(x, y) \leq \max(\rho^-(x, y), r - \rho^+(u, x) - \rho^+(v, y), \\ r - \rho^+(v, x) - \rho^+(u, y)) \end{aligned} \quad (3.12)$$

for $k \geq 0$. The claim will then follow when passing to the limit.

For the base case, $k = 0$, we have $\check{\rho}_0^-(x, y) = \rho_0^-(x, y)$, and (3.12) is clear, unless $(x, y) = (u, v)$. In the latter case, $\check{\rho}^-(x, y) = r$, and as this term occurs on the right in (3.12), again (3.12) is clear.

Now we proceed inductively, passing from $(k - 1)$ to k .

Once again, $\check{\rho}_k^-(x, y)$ is among the three following forms:

$$\begin{aligned} \check{\rho}_{k-1}^-(x, y) \\ \check{\rho}_{k-1}^-(x, z) - \check{\rho}_{k-1}^+(y, z) \\ \check{\rho}_{k-1}^-(y, z) - \check{\rho}_{k-1}^+(x, z). \end{aligned}$$

If $\check{\rho}_k^-(x, y) = \check{\rho}_{k-1}^-(x, y)$, then our inductive hypothesis yields our desired result.

We assume then without loss of generality that

$$\begin{aligned}\check{\rho}_k^-(x, y) &= \check{\rho}_{k-1}^-(x, z) - \check{\rho}_{k-1}^+(y, z) \\ &= \check{\rho}_{k-1}^-(x, z) - \rho_{k-1}^+(y, z).\end{aligned}\tag{3.13}$$

By induction, we have one of the following.

- a. $\check{\rho}_{k-1}^-(x, z) \leq \rho^-(x, z)$;
- b. $\check{\rho}_{k-1}^-(x, z) \leq r - (\rho^+(x, u) + \rho^+(z, v))$; or
- c. $\check{\rho}_{k-1}^-(x, z) \leq r - (\rho^+(x, v) + \rho^+(z, u))$.

In case (a), we have

$$\check{\rho}_k^-(x, y) \leq \rho^-(x, z) - \rho_{k-1}^+(y, z) \leq \rho^-(x, z)$$

and (3.12) holds.

In case (b), we have

$$\begin{aligned}\check{\rho}_k^-(x, y) &= \check{\rho}_{k-1}^-(x, z) - \rho_{k-1}^+(y, z) \\ &\leq r - (\rho^+(x, u) + \rho^+(z, v) + \rho^+(z, y))\end{aligned}$$

All that remains to be shown is that $r - \rho^+(u, x) - \rho^+(v, y) \leq \rho^+(x, y)$. This follows immediately from Lemma 3.1.14 (3.5).

The treatment of case (c) is similar.

This proves Claim 3.1.21.1.

We now prove the lemma. We have $\check{\rho}^+ = \rho^+$ by Lemma 3.1.20, so by Claim 3.1.21.1, it suffices to show that

$$\max(\rho^-(x, y), r - (\rho^+(x, u) + \rho^+(y, v)), r - (\rho^+(x, v) + \rho^+(y, u))) \leq \rho^+(x, y).$$

Now $\rho^- \leq \rho^+$ by hypothesis, and it follows from Lemma 3.1.14 (3.5) that

$$\begin{aligned}r - (\rho^+(x, u) + \rho^+(y, v)) &\leq \rho^+(x, y) \\ r - (\rho^+(x, v) + \rho^+(y, u)) &\leq \rho^+(x, y).\end{aligned}$$

The lemma follows. □

Corollary 3.1.21.1. *Suppose $\mathcal{G} = (G, w)$ is a candidate configuration for completion to $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$, where δ, K_1 , and C are finite. Then after the first and second steps of the completion algorithm have been applied to \mathcal{G} , yielding $\hat{\mathcal{G}} = (G, \hat{w})$, we will have that*

$$\hat{\rho}^-(x, y) \leq \hat{\rho}^+(x, y),$$

for every $x, y \in G$.

Proof. Let (u, v) and (x, y) be two edges of \mathcal{G} satisfying $\rho^+(u, v), \rho^+(x, y) < C - 2\delta - 1$. After adding the edge (u, v) with the label $\rho^+(u, v)$ to \mathcal{G} to form $\check{\mathcal{G}}$, by Lemma 3.1.21 we have $\check{\rho}^+ = \rho^+$ and therefore $\check{\rho}^+(x, y) < C - 2\delta - 1$ also in $\check{\mathcal{G}}$. Therefore the application of Lemma 3.1.21 can be iterated until all (finitely many) such pairs from \mathcal{G} have been dealt with. \square

We have thus shown the following.

Lemma 3.1.22. *Let $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$ be the amalgamation class of finite $[\delta]$ -metric spaces associated to admissible parameters*

$$(\delta, K_1, K_2, C, C', S)$$

with δ, K_1 , and C finite, and $C > 2\delta + \max(K_1, \delta/2)$, $C' = C + 1$. Let $\mathcal{G} = (G, w)$ be a $[\delta]$ -edge-labeled graph which is a candidate configuration for completion to $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$.

Then applying the first step of the completion algorithm to \mathcal{G} will extend \mathcal{G} to another candidate configuration $\hat{\mathcal{G}}$ for completion to $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$.

Furthermore, the new configuration $\hat{\mathcal{G}}$ satisfies

$$\tilde{\rho}^+(x, y) \geq C - 2\delta - 1$$

for all pairs (x, y) which are not edges of $\check{\mathcal{G}}$.

Proof. The first point follows immediately from Corollary 3.1.21.1 and Lemmas 3.1.19 and 3.1.18.

For the second point, since $\hat{\rho}^+ = \rho^+$, any edge with $\tilde{\rho}^+(u, v) < C - 2\delta - 1$ has already been added to $\hat{\mathcal{G}}$. \square

We proceed to show that the entirety of the completion algorithm yields a desired metric space in $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$.

Lemma 3.1.23. Let $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$ be the amalgamation class of finite $[\delta]$ -metric spaces associated to admissible parameters

$$(\delta, K_1, K_2, C, C', S)$$

with δ, K_1 , and C finite, and $C > 2\delta + \max(K_1, \delta/2)$. Let $\mathcal{G} = (G, w)$ be a candidate configuration for completion to $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$.

Suppose that for every non-edge (u, v) of \mathcal{G} , we have

$$\rho^+(u, v) \geq C - 2\delta - 1.$$

Then assigning weights

$$\bar{w} = \max(\rho^-(u, v), (C - 2\delta - 1)')$$

to the remaining non-edges (u, v) yields a completion $\bar{\mathcal{G}}$ which satisfies the triangle inequality.

Proof. Since \mathcal{G} is a candidate configuration for completion to $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$, it does not contain any triangles violating the triangle inequality.

We check the following inequalities:

1. $w(x, y) \leq w(x, z) + \bar{w}(y, z)$;
2. $\bar{w}(x, y) \leq w(x, z) + w(y, z)$;
3. $w(x, y) \leq \bar{w}(x, z) + \bar{w}(y, z)$;
4. $\bar{w}(x, y) \leq w(x, z) + \bar{w}(y, z)$;
5. $\bar{w}(x, y) \leq \bar{w}(x, z) + \bar{w}(y, z)$,

for $x, y, z \in G$, where we write \bar{w} only when the pair in question is a non-edge in G .

We observe first that $\rho^- \leq \bar{w} \leq \rho^+$ by our definition and the hypotheses. Therefore it will suffice to prove any of the inequalities with ρ^+ substituted for \bar{w} wherever it occurs on the left side, and with ρ^- substituted for \bar{w} wherever it occurs on the right side. This transforms inequalities (1, 2, 4) into

$$(1') \quad w(x, y) \leq w(x, z) + \rho^-(y, z);$$

$$(2') \quad \rho^+(x, y) \leq w(x, z) + w(y, z);$$

$$(4') \quad \rho^+(x, y) \leq w(x, z) + \rho^-(y, z).$$

Note that in this form, inequalities (1', 2') are instances of (4'), which is an instance of an inequality given in Lemma 3.1.14.

This leaves us with

$$(3) \quad w(x, y) \leq \bar{w}(x, z) + \bar{w}(y, z);$$

$$(5) \quad \bar{w}(x, y) \leq \bar{w}(x, z) + \bar{w}(y, z).$$

Here each term on the right is at least $(C - 2\delta - 1)'$; if one of these terms is $\delta - 1$, then the inequality becomes trivial, and otherwise both are at least $C - 2\delta - 1 \geq \delta/2$, and again the inequality becomes trivial. \square

Lemma 3.1.24. *Let $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$ be the amalgamation class of finite $[\delta]$ -metric spaces associated to admissible parameters*

$$(\delta, K_1, K_2, C, C', S)$$

with δ, K_1 , and C finite, and $C > 2\delta + \max(K_1, \delta/2)$. Let $\mathcal{G} = (G, w)$ be a candidate configuration for completion to $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$.

Suppose that for every non-edge (u, v) of \mathcal{G} , we have

$$\rho^+(u, v) \geq C - 2\delta - 1.$$

Then assigning weights

$$\bar{w} = \max(\rho^-(u, v), (C - 2\delta - 1)')$$

to the remaining non-edges (u, v) yields a completion $\bar{\mathcal{G}}$ which does not isometrically embed any triangles of odd perimeter less than $2K_1$.

Proof. By assumption, \mathcal{G} contains no triangles of odd perimeter less than $2K_1$, so any such triangle in $\bar{\mathcal{G}}$ must contain at least one new edge with weight at least $(C - 2\delta - 1)' = C - 2\delta - 1$ or $\delta - 1$.

Suppose first that $(C - 2\delta - 1)' = C - 2\delta - 1$. By our assumptions on C , this is at least K_1 . But the triangle inequality holds in $\bar{\mathcal{G}}$ by Lemma 3.1.23, so any triangle with an edge of weight at least K_1 has perimeter at least $2K_1$.

This leaves the case $(C - 2\delta - 1)' \neq C - 2\delta - 1$, a very special case in which we have

$$(C - 2\delta - 1)' = \delta - 1 \qquad \mathcal{S} \neq \emptyset.$$

If $K_1 = \delta$ then $\mathcal{S} = \emptyset$, by the definition of admissibility (Table 2.1). So here we have $(C - 2\delta - 1)' \geq K_1$, and we conclude as above. \square

Lemma 3.1.25. *Let $\mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$ be the amalgamation class of finite $[\delta]$ -metric spaces associated to admissible parameters*

$$(\delta, K_1, K_2, C, C', \mathcal{S})$$

with δ, K_1 , and C finite, and $C > 2\delta + \max(K_1, \delta/2)$. Let $\mathcal{G} = (G, w)$ be a candidate configuration for completion to $\mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$.

Suppose that for every non-edge (u, v) of \mathcal{G} , we have

$$\rho^+(u, v) \geq C - 2\delta - 1.$$

Then assigning weights

$$\bar{w} = \max(\rho^-(u, v), (C - 2\delta - 1)')$$

to the remaining non-edges (u, v) yields a completion $\bar{\mathcal{G}}$ which does not contain any violations of the Henson constraints, or the triangle type $(1, \delta, \delta)$, in the case that $K_2 = \delta - 1$.

Proof. We first address \mathcal{S} .

We first show that if (u, v) is a non-edge of \mathcal{G} and $\rho^-(u, v) > (C - 2\delta - 1)'$, then $\rho^-(u, v) \notin \{1, \delta\}$.

Since $\rho^-(u, v) > (C - 2\delta - 1)' > 0$, we have that $\rho^-(u, v) > 1$.

In particular, $\rho^-(u, v) \neq \rho_0^-(u, v)$ and thus

$$\rho^-(u, v) = \rho_j^-(u, v) = \rho_{j-1}^-(u, w) - \rho_{j-1}^+(w, v)$$

for some j and some $w \notin \{u, v\}$. Hence $\rho^-(u, v) \leq \delta - 1$.

Therefore $\rho^-(u, v)$, when $\rho^-(u, v) > (C - 2\delta - 1)'$, cannot introduce any values of 1 or δ and thus using the value ρ^- will not yield any violations of Henson constraints.

Now we show that the value $(C - 2\delta - 1)'$ does not lie in $\{1, \delta\}$ either. Certainly the value $\delta - 1$ is not in $\{1, \delta\}$, since $\delta \geq 3$.

Suppose then that $(C - 2\delta - 1)' = C - 2\delta - 1$ and that this value lies in $\{1, \delta\}$. If $C - 2\delta - 1 = 1$, then $C = 2\delta + 2$. By the definition of admissibility (Table 2.1), this implies that $\mathcal{S} = \emptyset$.

If $C - 2\delta - 1 = \delta$ and $\mathcal{S} \neq \emptyset$, then step 3 of the completion algorithm would use the value $(C - 2\delta - 1)' = \delta - 1$ instead of $C - 2\delta - 1$.

Thus, extending by $(C - 2\delta - 1)'$ would also not result in any violations of Henson constraints.

Similarly, if $K_2 = \delta - 1$ and a triangle of type $(1, \delta, \delta)$ appears in $\bar{\mathcal{G}}$, then the corresponding value $\bar{w}(u, v)$ must be 1, and $C = 2\delta + 2$. However it follows from the definition of admissibility that for $C > 2\delta + K_1$ and $K_2 = \delta - 1$, we must have that $C > 2\delta + 2$ (again from Table 2.1). \square

Lemma 3.1.26. *Let $\mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$ be the amalgamation class of finite $[\delta]$ -metric spaces associated to admissible parameters*

$$(\delta, K_1, K_2, C, C', \mathcal{S})$$

with δ, K_1 , and C finite, and $C > 2\delta + \max(K_1, \delta/2)$. Let $\mathcal{G} = (G, w)$ be a candidate configuration for completion to $\mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$.

Suppose that for every non-edge (u, v) of \mathcal{G} , we have

$$\rho^+(u, v) \geq C - 2\delta - 1.$$

Then assigning weights

$$\bar{w} = \max(\rho^-(u, v), (C - 2\delta - 1)')$$

to the remaining non-edges (u, v) yields a completion $\bar{\mathcal{G}}$ which does not contain any triangles of perimeter larger than $C - 1$.

Proof. As \mathcal{G} is a candidate configuration, it satisfies $\rho^- \leq \rho^+$ and therefore contains no triangles of perimeter greater than C . So any triangle in $\bar{\mathcal{G}}$ must contain at least one newly labeled edge.

If a newly labeled edge is labeled with the value $(C - 2\delta - 1)' \leq C - 2\delta - 1$, then any triangle containing this edge would have perimeter at most $C - 1$.

Thus we turn our attention to triangles for which any newly labeled edges were labeled with ρ^- .

Since $w = \rho^+$ on all edges in \mathcal{G} , and $\rho^- \leq \rho^+$, it suffices to prove

$$\rho^+(x, y) + \rho^+(x, z) + \rho^+(y, z) < C$$

which follows at once from the inductive definition of the weights ρ_i^+ . \square

We therefore have the following.

Proof of Proposition 3.1.6. By Lemma 3.1.22, the extended graph $\hat{\mathcal{G}}$ obtained after Step 1 of the completion algorithm is again a candidate configuration and satisfies the condition

$$\tilde{\rho}^+(x, y) \geq C - 2\delta - 1$$

for all pairs (x, y) which are not edges of $\hat{\mathcal{G}}$. Lemmas 3.1.23–3.1.26 show that the further extension $\bar{\mathcal{G}}$ constructed in Step 3 belongs to $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$. \square

3.1.3.5 Proof of the main theorem

We may now prove Proposition 2.

Proposition 2. *For any primitive metrically homogeneous graph $\Gamma_{K_1, K_2, C, C', S}^\delta$ of generic type for which $C > 2\delta + \max(K_1, \delta/2)$, $C' = C + 1$, and $K_2 \geq \delta - 1$, the class $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$ is finitely constrained.*

Proof.

Claim 3.1.0.1. *The candidate configurations for $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$ are precisely those configurations which embed into an element of $\mathcal{A}_{K_1, K_2, C, C', S}^\delta$.*

The claim comes easily now: inclusion in one direction is given by Proposition 3.1.6, using the completion process. Inclusion in the reverse direction is given by Lemma 3.1.17.

By Corollary 3.1.7.1 to Lemma 3.1.7 Lemmas 3.1.8 and 3.1.15, the candidate configurations are characterized by a finite set of forbidden configurations. The theorem follows. \square

We now give a sharper version of Proposition 2, with an explicit set of forbidden configurations. The argument above was our original proof, and the more general form given in [ABH⁺17] is also based on a certain completion procedure terminating in uniformly bounded time.

The following analysis is similar to that given more generally in [HKK17].

Lemma 3.1.27. *Let δ and C be finite, with $C > 2\delta + 1$.*

Then a finite $[\delta]$ -edge-labeled graph $\mathcal{G} = (G, w)$ satisfies the condition

$$\rho^- \leq \rho^+$$

if and only if \mathcal{G} does not contain any of the configurations in the set \mathcal{F}_0 .

Proof.

Claim 3.1.27.1. *If \mathcal{G} contains a cycle whose edges can be indexed so that*

$$\sum_{i \leq 2\ell+1} \rho^-(e_i) > \ell(C-1) + \sum_{i \geq 2\ell+2} \rho^+(e_i)$$

then $\rho^-(x, y) > \rho^+(x, y)$ for some vertices x, y in G .

We proceed by induction on ℓ .

For $\ell = 0$, the inequality becomes

$$\rho^-(e_1) > \sum_{i \geq 2} \rho^+(e_i).$$

As the edges form a cycle, Lemma 3.1.14 (3.5) applies and gives $\rho^-(e_1) > \rho^+(e_1)$. The claim follows.

Now we proceed inductively from $\ell - 1$ to ℓ .

Take the edges $e_1, \dots, e_{2\ell+1}$ to be in cyclic order according to some orientation of the cycle, with vertices (u_i, v_i) labeled consistently with respect to the same orientation, and let γ_i be the path from v_i to u_{i+1} (take modulo $2\ell + 1$).

Then repeated use of the inequalities in Lemma 3.1.14 yields

$$\rho^-(v_{2\ell}, u_1) \geq \rho^-(e_{2\ell+1}) - \sum_{e \in \gamma_{2\ell} \cup \gamma_{2\ell+1}} \rho^+(e).$$

Then the definition of ρ^+ gives

$$\rho^+(v_{2\ell}, u_1) \leq (C-1) - (\rho^-(e_{2\ell}) + \rho^-(e_{2\ell+1})) - \sum_{e \in \gamma_{2\ell} \cup \gamma_{2\ell+1}} \rho^+(e)$$

and then

$$\rho^+(v_{2\ell-1}, u_1) \leq (C-1) - (\rho^-(e_{2\ell}) + \rho^-(e_{2\ell+1})) - \sum_{e \in \gamma_{2\ell-1} \cup \gamma_{2\ell} \cup \gamma_{2\ell+1}} \rho^+(e).$$

Thus if we shorten the cycle by replacing the segment from $v_{2\ell-1}$ to u_1 by the edge $(v_{2\ell-1}, u_1)$, the corresponding inequality will continue to hold with ℓ replaced by $\ell-1$, and we may conclude by induction.

This proves the claim.

Claim 3.1.27.2. *If $\rho^- \leq \rho^+$ in \mathcal{G} , then \mathcal{G} contains no cycle in \mathcal{F}_0 .*

Supposing the contrary, since $\rho^- \leq \rho^+$, it follows that ρ^\pm must agree with w on the cycle, by Lemma 3.1.16. But this contradicts the first claim. Thus \mathcal{G} does not contain any cycle in \mathcal{F}_0 .

Claim 3.1.27.3. *If there is a cycle in \mathcal{G} whose edges can be enumerated so that*

$$\sum_{i \leq 2\ell+1} \rho^-(e_i) > \ell(C-1) + \sum_{i \geq 2\ell+2} \rho^+(e_i)$$

then \mathcal{G} contains one of the cycles in \mathcal{F}_0 .

For large k , this may be written as

$$\sum_{i \leq 2\ell+1} \rho_k^-(e_i) > \ell(C-1) + \sum_{i \geq 2\ell+2} \rho_k^+(e_i). \quad (\star)$$

Take k minimal so that there is a cycle satisfying (\star) . We show first that $k = 0$.

Supposing the contrary, replace each edge e for which a term $\rho_k^\pm(e)$ occurs in (\star) with a different value from the corresponding term $\rho_{k-1}^\pm(e)$ by a pair of edges e', e'' joining the endpoints of e and satisfying one of the relations

$$\begin{aligned} \rho_k^+(e) &= \rho_{k-1}^+(e') + \rho_{k-1}^+(e'') \\ \rho_k^+(e) &= (C-1) - (\rho_{k-1}^-(e') + \rho_{k-1}^-(e'')) \\ \rho_k^-(e) &= \rho_{k-1}^-(e') - \rho_{k-1}^+(e''). \end{aligned}$$

This produces a circuit whose edges can be enumerated so as to satisfy the same inequality, with $k-1$ in place of k , and with an increased value of ℓ if the second formula has been used. (Since terms with ρ^+ go on the right and terms with ρ^- go on the left, it is clear how the edges should be indexed, and one then checks that the inequality is preserved.)

This circuit then breaks up into cycles. For those with an even number of edges occurring on the left side of the inequality (\star) for $k-1$, we have the reverse inequality

$$\sum'_e \rho^-(e) \leq 2\ell\delta \leq \ell(C-1) + \sum''_e \rho^+(e)$$

with the sums restricted to edges in the given cycle.

It follows that one of the cycles with an odd number of edges occurring on the left side of (\star) must satisfy the inequality (\star) for $k - 1$, applied to that cycle. But this then contradicts the minimality of k .

So $k = 0$. We examine (\star) in this case.

$$\sum_{i \leq 2\ell+1} \rho_0^-(e_i) > \ell(C - 1) + \sum_{i \geq 2\ell+2} \rho_0^+(e_i).$$

If one of the terms on the left is 1 then we find

$$2\ell\delta + 1 \geq \sum_{i \leq 2\ell+1} \rho_0^-(e_i) > \ell(C - 1) \geq \ell(2\delta + 1)$$

hence $\ell = 0$.

But then the inequality becomes

$$1 > \sum_{i \geq 2} \rho^+(e_i)$$

and even a degenerate cycle of length two would have a term on the right. So all of the terms on the left side of our inequality are of the form $w(e)$.

If one of the terms on the right side is δ , then our inequality implies

$$(2\ell + 1)\delta > \ell(C - 1) + \delta \geq \ell(2\delta + 1) + \delta = (2\ell + 1)\delta + \ell,$$

a contradiction. So the terms on the right side are also of the form $w(e)$ and our cycle lies in \mathcal{F}_0 .

Claim 3.1.27.4. *If there is a pair (x, y) in G with $\rho^-(x, y) > \rho^+(x, y)$ then \mathcal{G} contains one of the cycles in \mathcal{F}_0 .*

We may consider the degenerate cycle with two oriented edges $e_1 = (x, y)$ and $e_2 = (y, x)$. We then have

$$\rho^-(e_1) > \rho^+(e_2)$$

and the previous claim applies even to this degenerate case, with $\ell = 0$, giving the desired result.

Now our second and fourth claims give the lemma. □

We may now deduce the following.

Corollary 3.1.27.1. *Let $\Gamma = \Gamma_{K_1, K_2, C, C', \mathcal{S}}^\delta$ be a primitive metrically homogeneous graph of generic type for which $C > 2\delta + \max(K_1, \delta/2)$, $C' = C + 1$, and $K_2 \geq \delta - 1$. Then a finite $[\delta]$ -edge-labeled graph embeds into Γ if and only if it does not contain any of the cycles in the finite set \mathcal{F} .*

Proof. By the proof of Proposition 2, the finite $[\delta]$ -edge-labeled graphs which embed in Γ are the candidate configurations for completion to $\mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$. By definition and by Lemma 3.1.27, these are the configurations which do not contain any cycle from \mathcal{F} .

Finally, \mathcal{F} is finite by Lemma 3.1.10. □

3.2 Ample generics and the topological group $\text{Aut}(\Gamma)$

Both this section and the next are concerned with $\text{Aut}(\Gamma)$ as a Polish group.

In Section 3.1, we showed that the known primitive metrically homogeneous graphs of generic type whose associated numerical parameters satisfy some additional restrictions are *finitely constrained* in the sense of Definition 2.6.2. At the same time, we developed a *completion process* for systematically completing a partial substructure to an induced substructure.

In this section, we show that the results of Section 3.1 have powerful consequences for the behavior of the automorphism group as a Polish group. Using the formulation developed in Section 2.6, we combine the results of Herwig-Lascar and Kechris-Rosendal, discussed in Section 2.3, with the results of Section 3.1, to show that the automorphism group has ample generics, and then deduce a plethora of topological consequences. In the following section, we will derive dynamical properties of the automorphism group as a topological group by applying the result of Section 3.1 in conjunction with the combinatorial results of Hubička-Nešetřil and the general theory of Kechris-Pestov-Todorčević, discussed in Section 2.4.

The main result of this section is the following.

Theorem 1. *Let $(\delta, K_1, K_2, C, C', \mathcal{S})$ be an admissible parameter sequence with δ and K_1 finite, for which $C > 2\delta + \max(K_1, \delta/2)$, $K_2 \geq \delta - 1$ and $C' = C + 1$, and let $\Gamma = \Gamma_{K_1, K_2, C, C', \mathcal{S}}^\delta$ be the associated primitive metrically homogeneous graph of generic type.*

Then the Polish group $\text{Aut}(\Gamma)$ has ample generics, and therefore has the following properties.

- *The small index property and automatic continuity;*
- *uncountable cofinality;*
- *the fixed point properties (FA) and (FH).*

As we mentioned in the introduction, the small index property and automatic continuity both address ways in which the structure of the automorphism group as an abstract group determines its topology. Uncountable cofinality and properties (FA) and (FH) conversely state that properties of the abstract group can be derived from topological properties.

In the setting of Theorem 1, we will derive the existence of ample generics from the properties EPPA and APA discussed in Sections 2.3 and 2.6. The key to this is to show that the completion process of Section 3.1.3 is a canonical completion process in the sense of Definition 2.6.4.

Lemma 3.2.1. *Under the hypotheses of Theorem 1, the completion process given in Definition 3.1.9 is a canonical completion process for finite partial $\mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$ -structures.*

Proof. We recall the procedure in general terms.

1. Calculate weights ρ_i^\pm inductively and take limits to get weights ρ^\pm .
2. Add edges with weight ρ^+ whenever $\rho^+ < C - 2\delta - 1$.
3. Recalculate ρ_i^\pm and ρ^\pm , calling the result $\hat{\rho}^\pm$.
4. Assign weights $\max(\rho^-, (C - 2\delta - 1)')$ to the remaining edges, where $(C - 2\delta - 1)'$ is an explicitly given constant.

Since the weights ρ_i^\pm are defined inductively in terms of the original weights and the previous weights ρ_j^\pm ($j < i$), at each of these four stages these weights are preserved by isomorphisms at the level of the initial structures. Therefore the completion process is canonical. \square

Lemma 3.2.2. *Let $(\delta, K_1, K_2, C, C', \mathcal{S})$ be an admissible parameter sequence with δ and K_1 finite, for which*

$$C > 2\delta + \max(K_1, \delta/2) \qquad K_2 \geq \delta - 1 \qquad C' = C + 1.$$

Let $\mathcal{A} = \mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$ be the associated amalgamation class of finite $[\delta]$ -valued metric spaces. Then \mathcal{A} has the amalgamation property for automorphisms (APA) and the extension property for partial automorphisms (EPPA).

Proof. The class \mathcal{A} is a strong amalgamation class (Observation 1). By Lemma 3.2.1, there is a canonical completion process for finite partial \mathcal{A} -structures. By Lemma 2.6.4, the class \mathcal{A} has the APA.

By Proposition 2, the class \mathcal{A} is finitely constrained. As we are treating metrically homogeneous graphs of diameter δ as edge-labeled complete graphs with labels in $[\delta]$, the class is relationally complete and the structures are irreflexive. By Proposition 2.6.1, the class \mathcal{A} has the EPPA. \square

We may now prove Theorem 1.

Theorem 1. *Let $(\delta, K_1, K_2, C, C', \mathcal{S})$ be an admissible parameter sequence with δ and K_1 finite, for which $C > 2\delta + \max(K_1, \delta/2)$, $K_2 \geq \delta - 1$ and $C' = C + 1$, and let $\Gamma = \Gamma_{K_1, K_2, C, C', \mathcal{S}}^\delta$ be the associated primitive metrically homogeneous graph of generic type.*

Then the Polish group $\text{Aut}(\Gamma)$ has ample generics, and therefore has the following properties.

- *The small index property and automatic continuity;*
- *uncountable cofinality;*
- *the fixed point properties (FA) and (FH).*

Proof. By Lemma 3.2.2, the age of Γ has the APA and the EPPA. By Theorem 2.3.3, Γ has ample generics. By Theorem 2.3.1, we have that Γ has the small index property and automatic continuity. By Theorem 2.3.2, G has uncountable cofinality and properties (FA) and (FH). \square

3.3 The Ramsey property and extreme amenability

In this section, we use the finiteness result of Section 3.1.3 to derive results on the dynamical properties of the automorphism groups of the metrically homogeneous graphs Γ to which Proposition 2 applies. We first consider the automorphism group of the generic expansion of Γ by a linear order, and then return to the full automorphism group of the underlying graph.

Our main result is as follows.

Theorem 2. *Let $(\delta, K_1, K_2, C, C', \mathcal{S})$ be an admissible parameter sequence with δ and K_1 finite, for which $C > 2\delta + \max(K_1, \delta/2)$, $K_2 \geq \delta - 1$ and $C' = C + 1$, and let $\Gamma = \Gamma_{K_1, K_2, C, C', \mathcal{S}}^\delta$ be the*

associated primitive metrically homogeneous graph of generic type.

Then the universal minimal flow of $\text{Aut}(\Gamma)$ is metrizable.

Furthermore, the universal minimal flow of $\text{Aut}(\Gamma)$ is the space $L(\Gamma)$ of all linear orderings of Γ .

With Γ as in Theorem 2, we first consider the automorphism group of its generic expansion by a linear order $(\Gamma, <)$.

Proposition 1. *Let $(\delta, K_1, K_2, C, C', \mathcal{S})$ be an admissible parameter sequence with δ and K_1 finite, for which $C > 2\delta + \max(K_1, \delta/2)$, $K_2 \geq \delta - 1$ and $C' = C + 1$. Let $\Gamma = \Gamma_{K_1, K_2, C, C', \mathcal{S}}^\delta$ be the associated primitive metrically homogeneous graph of generic type, and let $(\Gamma, <)$ be the generic expansion of Γ by a linear order.*

Then $\text{Aut}(\Gamma, <)$ is extremely amenable.

Proof. Let \mathcal{A} be the age of Γ and let $\mathcal{A}^<$ be the age of $(\Gamma, <)$. By Theorem 2.4.2, it suffices to show that $\mathcal{A}^<$ has the Ramsey property. Note that for our choices of parameters, the classes \mathcal{A} and $\mathcal{A}^<$ are strong amalgamation classes.

Let \mathcal{L}_0 be the language consisting of binary relations R_i for $i \in [\delta]$ and let \mathcal{L} be the expansion of \mathcal{L}_0 by a binary relation \leq . Let \mathcal{F}_0 be the class of irreducible \mathcal{L}_0 -structures on at most two points in which at least one of the following occurs:

- an edge relation is asymmetric;
- an edge relation is reflexive;
- more than one edge relation holds on a pair of points.

Let \mathcal{F} be the class of ordered \mathcal{L} -structures whose reducts to the language \mathcal{L}_0 lie in \mathcal{F}_0 .

Let \mathcal{R} be the class of finite ordered \mathcal{L} -structures which forbid all elements of \mathcal{F} under embedding. In other words, \mathcal{R} consists of \mathcal{L} -structures for which \leq is a total order and the reduct to the language \mathcal{L}_0 is a $[\delta]$ -edge-labeled graph.

By Theorem 2.5.1, the class \mathcal{R} is a Ramsey class.

By Theorem 2.5.2, it now suffices to check that the subclass \mathcal{A} is locally finite.

By Proposition 2, the class \mathcal{A} is finitely constrained. Let \mathcal{F} be a suitable finite set of constraints and let n be the maximal cardinality of an element of \mathcal{F} . It suffices to show that any

finite \mathcal{L} -structure A with the following two properties can be completed to an element in $\mathcal{A}^<$.

1. Any induced substructure of \mathcal{A} on n elements completes to an element of $\mathcal{A}^<$;
2. There is a homomorphism h from A to a structure \bar{A} in \mathcal{R} which is injective on each irreducible substructure of A .

The first condition implies that the reduct of A to the language of $[\delta]$ -metric spaces is \mathcal{F} -free, and therefore has a completion in \mathcal{A} .

Consider A and \bar{A} as directed graphs with edge relations given by \leq , the interpretation of \leq , respectively. Since h is a homomorphism, it takes cycles in A to cycles in \bar{A} . As \bar{A} is acyclic, we have that A is also acyclic. Therefore the relation \leq on A extends to a linear order, as required. \square

We now turn to the topological dynamical properties of the full automorphism group of our metrically homogeneous graphs. We show that the universal minimal flow is metrizable, and even identify it explicitly.

Theorem 2. *Let $(\delta, K_1, K_2, C, C', \mathcal{S})$ be an admissible parameter sequence with δ and K_1 finite, for which $C > 2\delta + \max(K_1, \delta/2)$, $K_2 \geq \delta - 1$ and $C' = C + 1$, and let $\Gamma = \Gamma_{K_1, K_2, C, C', \mathcal{S}}^\delta$ be the associated primitive metrically homogeneous graph of generic type.*

Then the universal minimal flow of $\text{Aut}(\Gamma)$ is metrizable.

Furthermore, the universal minimal flow of $\text{Aut}(\Gamma)$ is the space $L(\Gamma)$ of all linear orderings of Γ .

Proof. This will follow from the two parts of [KPT05, Theorem 7.5] (Theorem 2.4.3), with $\mathcal{K}_0 = \mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$, and $\mathcal{K} = \mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^{\delta, <}$.

By Theorem 2.5.1, we have that $\mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^{\delta, <}$ is Ramsey, and therefore metrizability follows from part 1 from the aforementioned theorem. The explicit identification of the universal minimal flow follows from part 2, once the ordering property is known.

The ordering property follows from Theorem 2.4.4, since Γ is homogeneous for a symmetric binary language (see the comment following Theorem 2.4.4). \square

Chapter 4

Algebraic Results

We now take a very different approach to the study of metrically homogeneous graphs from the topological approach of Chapter 3. We consider two topics of a more algebraic character, relying ultimately on the same type of Fraïssé theoretic information as in Chapter 3.

The first of these topics is the classification of *twisted isomorphisms* of metrically homogeneous graphs (Definition 2.7.2), and in particular, the twisted automorphism groups and associated groups of twists acting on the language, which we identify with the set $[\delta]$. As noted in the introduction, we find the same twists in our context (Theorem 3) as were found by Bannai and Bannai, and by Gardiner, in [BB80, Gar80], and we classify the metrically homogeneous graphs for which non-trivial twists occur, to the extent the current state of the classification theory allows. In particular, the group of outer automorphisms turns out to be either trivial or cyclic of order 2 in the case of metrically homogeneous graphs of generic type, known or otherwise. This is proved in Sections 4.1 and 4.2.

Our second topic is the question as to when the algebra of the age of a *known* metrically homogeneous graph of generic type is a polynomial algebra. Using a criterion given in [Cam97], we find that this is usually the case, with some extreme cases still unsettled (Theorem 4). Here a very special case of canonical completion of partial structures is required, namely the canonical completion of a disjoint union.

4.1 Twists

The following lemmas will be widely applicable in our argument.

Lemma 4.1.1. *Let $\delta \leq \infty$ be fixed, and let σ be a permutation of $[\delta]$ for which there is some metrically homogeneous graph Γ twistable by σ . Let $k = \sigma^{-1}(1)$. Then $\sigma(ik) = i$ for all i satisfying $ik \leq \delta$.*

Proof. We proceed by induction. Our base case $i = 1$ holds by definition, so we fix some i with $1 < i \leq \delta/k$ and assume for all $j < i$ that $\sigma(jk) = j$. A geodesic of type $(k, (i-1)k, ik)$ must be realized in Γ and therefore its image under σ , the triangle type $(1, i-1, \sigma(ik))$, must be realized in Γ^σ . By the triangle inequality we have $|\sigma(ik) - (i-1)| \leq 1$. As $\sigma(jk) = j$ for all $j < i$, the only remaining option then is $\sigma(ik) = i$. \square

Note that from this we may deduce that if $\sigma(1) = 1$, then $\sigma(i) = i$ for all $i \leq \delta$. This point also follows from Fact 2.2.1 as the metric d in Γ^σ must be the graph metric. Thus if σ is a non-trivial permutation of the language for which there is some metrically homogeneous graph Γ twistable by σ , then $\sigma(1) > 1$.

We also deduce that the diameter of Γ must be finite:

Lemma 4.1.2. *Let Γ be a metrically homogeneous graph and let σ be a non-trivial permutation of the language of a metrically homogeneous graph Γ such that Γ^σ is also a metrically homogeneous graph. Then the diameter δ of Γ is finite.*

Proof. Suppose $\delta = \infty$ and let $k = \sigma^{-1}(1)$. By Lemma 4.1.1, we have

$$\sigma[k\mathbb{N}] = \mathbb{N}$$

and hence $k\mathbb{N} = \mathbb{N}$. Thus $k = 1$, but then σ is trivial. \square

We will make frequent use of the following fact.

Fact 2.2.1 from 2.2 may be phrased usefully for our purposes as follows.

Fact 4.1.1. *Let Γ be a metrically homogeneous graph and σ a permutation of the language of Γ . Then the following are equivalent.*

- Γ^σ is a metrically homogeneous graph
- Γ^σ is a metric space, and contains all triangles of type $(1, k, k+1)$ for k less than the diameter of Γ .

4.1.1 Non-generic type

Here we work towards proving the following main result:

Proposition 4.1.1. *Let Γ be a metrically homogeneous graph of non-generic type or diameter $\delta \leq 2$, and σ a non-trivial permutation of the language such that Γ^σ is metrically homogeneous. Then Γ^σ is also of non-generic type or diameter $\delta \leq 2$, and one of the following applies.*

- Γ has diameter 2 and is not complete multipartite, with σ the transposition (12).
- Γ is finite, antipodal of diameter 3, not bipartite, with σ the transposition (12).
- Γ is an n -cycle C_n , $n \geq 7$; $\delta = \lfloor n/2 \rfloor$: σ is given by multiplication by $\pm k \pmod{n}$ for some k with $(k, n) = 1$, with $\pm i$ identified.

We prove this result by first showing the following three claims.

Claim 4.1.1.1. *For $\delta \leq 2$ the metrically homogeneous graphs with non-trivial twists are the homogeneous graphs which are neither disjoint unions of complete graphs nor complements of disjoint unions of complete graphs.*

Claim 4.1.1.2. *Let Γ be an n -cycle with $n \geq 3$. Then the permutations σ of the language of Γ for which Γ^σ is a metrically homogeneous graph are given by the group $U_n/(\pm 1)$ acting on the set of distances, where U_n is the group of units modulo n . More specifically, we define the action μ_k on the set of distances to be multiplication by k modulo n where $k \in U_n$.*

Claim 4.1.1.3. *Let Γ be a metrically homogeneous antipodal graph of diameter 3. If Γ is not bipartite, then $\sigma = (12)$ is the unique non-trivial twist for which Γ^σ is metrically homogeneous. If Γ is bipartite, then no such permutation exists.*

Proof of Claim 4.1.1.1. The metrically homogeneous graphs of diameter at most 2 are precisely the connected homogeneous graphs. For diameter $\delta = 2$, twisting via $\sigma = (12)$ is equivalent to taking the complement of the graph. Since the complement of a homogeneous graph is a homogeneous graph, in order to ensure that both Γ and Γ^σ are metrically homogeneous graphs, we need to show that both Γ and Γ^σ are connected homogeneous graphs.

The only disconnected homogeneous graphs of diameter at most 2 are those which are disjoint unions of complete graphs (Fact 2.1.2). Thus, the twistable metrically homogeneous graphs of diameter $\delta \leq 2$ are the homogeneous graphs which are neither disjoint unions of complete graphs nor are they the complements of disjoint unions of complete graphs.

Proof of Claim 4.1.1.2. If $k \in U_n$, then μ_k is an automorphism of the group $\mathbb{Z}/n\mathbb{Z}$. The group $\mathbb{Z}/n\mathbb{Z}$ can be endowed with the graph structure of C_n by defining two elements a, b to be adjacent if $a - b = \pm 1$. The action of μ_k on $\mathbb{Z}/n\mathbb{Z}$ yields another graph where now there is an edge between a and b if $a - b = \pm k$. This graph is isomorphic to C_n since k is a generator of $\mathbb{Z}/n\mathbb{Z}$. Hence the action of μ_k not only sends C_n to some metrically homogeneous graph, rather it sends C_n to C_n .

Now suppose C_n^σ is a metrically homogeneous graph and $k = \sigma^{-1}(1)$. As Γ is connected, all points of Γ are connected by paths with successive distances equal to k , and thus all distances occurring in Γ^σ are divisible by $\gcd(k, n)$. But one of these distances is 1, so $k \in U_n^*$. Recalling from Lemma 4.1.1 that $\sigma(1)$ determines σ , we have $\sigma = \mu_k$.

Proof of Claim 4.1.1.3. Suppose Γ is antipodal of diameter 3 and that Γ^σ is metrically homogeneous (in particular, connected) with σ a non-trivial twist. Note that Γ^σ is also of diameter 3. Write d^σ for the path metric on Γ^σ ; $d^\sigma(x, y) = d(x, y)^\sigma$.

Let $k = \sigma(3)$. Then $d^\sigma(x, y) = k$ defines a pairing on Γ^σ ; that is, for each vertex v there is a unique v' with $d^\sigma(x, y) = k$.

If $k = 1$, then since Γ^σ is connected, it consists of just two vertices, and cannot have diameter 3.

If $k = 2$, then $(\Gamma^\sigma)_2$ consists of a unique vertex v . But every vertex of $(\Gamma^\sigma)_1 \cup (\Gamma^\sigma)_2$ has a neighbor in $(\Gamma^\sigma)_2$, so v is adjacent to all points of Γ^σ other than the chosen basepoint. But then no vertex lies at distance 3 from v , a contradiction.

So σ must fix the value 3, and being non-trivial, must be the permutation (12). Since Γ^σ is connected, Γ is connected with respect to the edge relation $d(x, y) = 2$. But then Γ cannot be bipartite.

Now, conversely, suppose that Γ is antipodal, metrically homogeneous, of diameter 3, and not bipartite, and that $\sigma = (12)$. We claim that Γ^σ is a metrically homogeneous graph. By Fact 4.1.1, it suffices to show that Γ^σ contains geodesic triangles of types (1, 1, 2) and (1, 2, 3), and does not contain triples of type (1, 1, 3).

As Γ contains geodesic triangles of type (1, 2, 3), so does Γ^σ .

As Γ is not bipartite, it contains some triangle of odd perimeter, and as Γ is antipodal, the

perimeter is bounded by $2\delta = 6$. So the possible triangle types are $(1, 1, 1)$ or $(2, 2, 1)$. By a given triangle of type $(1, 1, 1)$, after replacing a point by an antipodal point and applying the antipodal law of Fact 2.2.8, we get a triangle of type $(2, 2, 1)$. So in any case, Γ contains a triangle of type $(2, 2, 1)$, and thus Γ^σ contains a triangle of type $(1, 1, 2)$.

This concludes the proof of the claim.

Proof of Proposition 4.1.1. The case of diameter at most 2 is covered by Claim 4.1.1.1, bearing in mind that Γ^σ has the same diameter as Γ . So we suppose $\delta \geq 3$ and Γ is of non-generic type. By the classification of non-generic type, given as Fact 2.2.2 in Section 2.2, Γ is then finite or one of the tree-like graphs $T_{m,n}$. When Γ is finite, Fact 2.2.5 describes the possibilities, and Claims 4.1.1.2 and 4.1.1.3 deal with those possibilities. Of course in this case Γ^σ is also finite and thus not of generic type. Finally, the case of $T_{m,n}$ does not arise since Lemma 4.1.2 tells us that the diameter must be finite. \square

4.1.2 Generic type

We work towards the following result.

Proposition 4.1.2. *Let σ be a non-trivial permutation of the language of a metrically homogeneous graph Γ of generic type where Γ^σ is itself a metrically homogeneous graph. Then σ is one of $\rho, \rho^{-1}, \tau_0, \tau_1$, which are defined as follows:*

$$\rho(i) = \begin{cases} 2i & i \leq \delta/2 \\ 2(\delta - i) + 1 & i > \delta/2 \end{cases} \quad \rho^{-1}(i) = \begin{cases} i/2 & i \text{ even} \\ \delta - \frac{i-1}{2} & i \text{ odd} \end{cases}$$

and for $\epsilon = 0$ or 1 , τ_ϵ is the involution defined by

$$\tau_\epsilon(i) = \begin{cases} (\delta + \epsilon) - i & \text{for } \min(i, (\delta + \epsilon) - i) \text{ odd} \\ i & \text{otherwise.} \end{cases}$$

We begin by showing the following:

Proposition 4.1.3. *Let σ be a non-trivial permutation of the language of a metrically homogeneous graph Γ of generic type which maps Γ to another metrically homogeneous graph.*

Then δ is finite, and $\sigma(1) \in \{2, \delta - 1, \delta\}$.

Proof. Let $\sigma(1) = k$, and assume towards a contradiction that $2 < k < \delta - 1$. By Proposition 2.2.1 of Section 2.2, the triangle types $(k, k, 2)$ and $(k, k, 4)$ must be realized in Γ^σ . Thus their inverse images under σ , namely $(1, 1, \sigma^{-1}(2))$ and $(1, 1, \sigma^{-1}(4))$, satisfy the triangle inequality. This implies that $\sigma^{-1}(\{2, 4\}) = \{1, 2\}$. Hence

$$\sigma(1) = 4 \qquad \sigma(2) = 2 \qquad \delta \geq 6,$$

as $\sigma(1) < \delta - 1$.

We will now argue that this implies that $\delta = 6$, and a contradiction will follow.

If $\delta \geq 7$, then again by Proposition 2.2.1, the triangle types $(k, k, 2)$, $(k, k, 4)$ and $(k, k, 6)$ must all be realized in Γ^σ . However, there are only two possible values i for which the triple $(1, 1, i)$ will satisfy the triangle inequality. We therefore have a contradiction in this case.

Now suppose that $\delta = 6$. Since the triangle type $(2, 4, 6)$ is of geodesic type, it must be realized in Γ^σ (Observation 2.2.1), and therefore $\sigma^{-1}(2, 4, 6) = (1, 2, \sigma^{-1}(6))$ must be realized in Γ . This implies then that $\sigma^{-1}(6) \leq 3$. The only option then is that

$$\sigma(3) = 6.$$

This leaves $\sigma(4) \in \{1, 3, 5\}$. The geodesic types $(1, 3, 4)$ and $(2, 2, 4)$ are realized in Γ , and therefore their images $(4, 6, \sigma(4))$ and $(2, 2, \sigma(4))$ are realized in Γ^σ . This implies that $2 \leq \sigma(4) \leq 4$. Thus,

$$\sigma(4) = 3.$$

Finally, we examine $\sigma^{-1}(1) \in \{5, 6\}$. The geodesic type $(1, 3, 4)$ being realized in Γ^σ implies that $(1, 4, \sigma^{-1}(1))$ is realized in Γ , and therefore $\sigma^{-1}(1) \leq 5$. This gives us that $\sigma(5) = 1$, and thus $\sigma = (14365)$. However, this permutation would send the geodesic type $(2, 3, 5)$ to $(1, 2, 6)$, and hence is not a suitable twist.

We have therefore indeed shown that $\sigma(1) \in \{2, \delta - 1, \delta\}$. □

Proposition 4.1.4. *Let σ be a permutation of the language of a metrically homogeneous graph Γ of generic type which sends Γ to another metrically homogeneous graph, satisfying $\sigma(1) = 2$. Then either $\sigma = \rho$ or $\delta = 3$ and σ is the transposition (12) .*

Proof. By Lemma 4.1.1, $\sigma(i) = 2i$ for all $i \leq \delta/2$. So the image of σ on $(\delta/2, \delta]$ is the set I of odd numbers in the interval $[1, \delta]$.

Consider the geodesic type $(1, i, i + 1)$ for any $i < \delta$, which σ maps to $(2, \sigma(i), \sigma(i + 1))$. Since this triple must satisfy the triangle inequality, we know that

$$|\sigma(i) - \sigma(i + 1)| \leq 2. \quad (4.1)$$

If $i > \delta/2$, then both $\sigma(i)$ and $\sigma(i + 1)$ are odd, and therefore $|\sigma(i) - \sigma(i + 1)| = 2$. Thus the values $\sigma(i)$ for $i > \delta/2$ give either an increasing or a decreasing enumeration of I . In the latter case $\sigma = \rho$. So we suppose that $\sigma(i)$ enumerates I in increasing order for $i > \delta/2$.

Let $k = \lfloor \delta/2 \rfloor + 1$. Then we have in particular that $\sigma(k) = 1$. The image of the geodesic triangle type $(1, k - 1, k)$ under σ is $(2, 2k - 2, 1)$ and the triangle inequality gives $2\lfloor \delta/2 \rfloor \leq 3$, hence $\delta \leq 3$. But then, as $\sigma(1) = 2$, we either have $\sigma = \rho$ or $\sigma = (12)$ with $\delta = 3$. \square

Corollary 4.1.1. *Let σ be a permutation of the language of a metrically homogeneous graph Γ of generic type for which Γ^σ is a metrically homogeneous graph, with $\sigma(2) = 1$. Then either $\sigma = \rho^{-1}$, or σ is the transposition (12) and $\delta = 3$.*

Proof. The graph Γ^σ is twistable by σ^{-1} and is of non-generic type by Proposition 4.1.1. So Proposition 4.1.4 applies to σ^{-1} and Γ^σ , giving the result. \square

We finally show the following:

Proposition 4.1.5. *Let σ be a permutation of the language of a metrically homogeneous graph Γ of generic type with diameter $\delta \geq 3$ such that Γ^σ is itself a metrically homogeneous graph. Assume in addition that $\sigma^{-1}(1) > 2$, that $\sigma(1) \geq \delta - 1$, and $\sigma(1) > 2$. Then either $\sigma = \tau_\epsilon$, with $\epsilon = \sigma(1) - (\delta - 1) \in \{0, 1\}$, or $\sigma = \rho^{-1}$ and $\delta = 3$.*

The proof will be inductive. The base of the induction depends in part on the following.

Lemma 4.1.3. *Let σ be a permutation of the language of a metrically homogeneous graph Γ of generic type with diameter $\delta \geq 3$ such that Γ^σ is itself a metrically homogeneous graph. Assume moreover that $\sigma(1) \geq \delta - 1$, $\sigma(1) > 2$, and $\sigma^{-1}(1) > 2$. Then $\sigma(2) = 2$.*

Proof. Suppose first that $\sigma(1) = \delta - 1$. Since by Proposition 2.2.1 the triangle type $(\delta - 1, \delta - 1, 2)$ is realized in Γ^σ , its inverse image $(1, 1, \sigma^{-1}(2))$ must satisfy the triangle inequality, meaning that $\sigma^{-1}(2) \leq 2$. Since by assumption $\sigma(1) > 2$, we have that $\sigma(2) = 2$.

Now assume that $\sigma(1) = \delta$. Then $(\Gamma^\sigma)_\delta = (\Gamma_1)^\sigma$. There are at most two distances realized in Γ_1 , and hence the same applies to $(\Gamma_1)^\sigma$; namely, at most $\sigma(1)$ and $\sigma(2)$ occur. Hence the same applies to $(\Gamma^\sigma)_\delta$; thus the only two distances which may occur in $(\Gamma^\sigma)_\delta$ are $\sigma(1) = \delta$ and $\sigma(2)$.

Using Fact 2.2.16, we know that each vertex in $\Gamma_{\delta-1}$ has two neighbors in Γ_δ . The distance i between these two points is either 1 or 2. So $i \neq \sigma(1)$. We therefore have that $\sigma(2) = i \leq 2$. Since $\sigma^{-1}(1) > 2$, we have our desired result: $\sigma(2) = 2$. \square

Lemma 4.1.4. *Let σ be a permutation of the language of a metrically homogeneous graph Γ of generic type such that Γ^σ is metrically homogeneous and suppose that $\sigma(2) = 2$. Then for $3 \leq k \leq \delta$, we have that*

$$|\sigma(k) - \sigma(k - 2)| \leq 2.$$

and

$$|\sigma^{-1}(k) - \sigma^{-1}(k - 2)| \leq 2.$$

Proof. Apply the triangle inequality to the image under σ or σ^{-1} of the geodesic type $(2, k - 2, k)$. Our assumption that $\sigma(2) = 2$ then yields our desired result. \square

We now proceed with the proof of Proposition 4.1.5.

Proof of Proposition 4.1.5. We initially assert the following claim.

Claim 4.1.5.1. *Let k be even and at most δ . Assume moreover that $k \leq (\delta + \epsilon)/2$ or $(\delta + \epsilon)$ is even. Then*

$$\sigma^{-1}(\tau_\epsilon(k)) \leq k \qquad \sigma^{-1}(\tau_\epsilon(k - 1)) \leq k - 1$$

Proof of Claim 4.1.5.1. Note that for the values specified, $\tau_\epsilon(k) = k$. Thus we show that $\sigma^{-1}(k) \leq k$.

We proceed by induction.

For $k = 2$, we have from Lemma 4.1.3 that $\sigma(2) = 2$. Moreover, by assumption, $\sigma^{-1}(\delta + \epsilon - 1) = 1$. Thus our base case holds. We assume then that $k > 2$ and for all even $j < k$ that $\sigma^{-1}(j) \leq j$ and $\sigma^{-1}(\tau_\epsilon(j - 1)) \leq j - 1$.

By Lemma 4.1.4, we know that $|\sigma^{-1}(k) - \sigma^{-1}(k - 2)| \leq 2$. Since by assumption $\sigma^{-1}(k - 2) \leq k - 2$, we have that $\sigma^{-1}(k) \leq k$.

We consider now $\sigma^{-1}(\tau_\epsilon(k - 1))$. Note that for even j with $j \leq (\delta + \epsilon)/2$ or $(\delta + \epsilon)$ even, we have $\tau_\epsilon(j + 1) = \tau_\epsilon(j - 1) - 2$, and thus $\sigma^{-1}(\tau_\epsilon(k - 1)) = \sigma^{-1}(\tau_\epsilon(k - 3) - 2)$. Again using Lemma 4.1.4, we get that

$$|\sigma^{-1}(\tau_\epsilon(k - 1)) - \sigma^{-1}(\tau_\epsilon(k - 3))| \leq 2.$$

Since by induction $\sigma^{-1}(\tau_\epsilon(k - 3)) \leq k - 3$, we indeed have that $\sigma^{-1}(\tau_\epsilon(k - 1)) \leq k - 1$.

We now move on to the next claim:

Claim 4.1.5.2. *Suppose that k is even and $2 \leq k \leq \delta$. Assume that $k \leq (\delta + \epsilon)/2$ or $\delta + \epsilon$ is even. Then $\sigma(k) = \tau_\epsilon(k)$ and $\sigma(k - 1) = \tau_\epsilon(k - 1)$.*

Proof of Claim 4.1.5.2. We begin by noting that for $i \leq (\delta + \epsilon)/2$ or for $\delta + \epsilon$ even, the permutation τ_ϵ is as follows.

$$\tau_\epsilon(i) = \begin{cases} (\delta + \epsilon) - i & i \text{ odd} \\ i & i \text{ even} \end{cases}$$

We proceed by induction. We know by assumption and by Lemma 4.1.3 that $\sigma(2) = 2 = \tau_\epsilon(2)$ and $\sigma(1) = \delta + \epsilon - 1 = \tau_\epsilon(1)$. Thus we assume that $k > 2$ and for j even, $j \leq k - 2$, that $\sigma(j) = \tau_\epsilon(j) = j$ and $\sigma(j - 1) = \tau_\epsilon(j - 1) = \delta + \epsilon - i$.

For the assumed values of k and $\delta + \epsilon$, we have from Claim 4.1.5.1 that $\sigma^{-1}(\tau_\epsilon(k)) \leq k$ and $\sigma^{-1}(\tau_\epsilon(k - 1)) \leq k - 1$.

Since $\tau_\epsilon(k - 1) \neq \tau_\epsilon(i)$ for any $i < k - 1$, it is also the case that $\sigma^{-1}(\tau_\epsilon(k - 1)) \neq \sigma^{-1}(\tau_\epsilon(i))$ for any $i < k - 1$. Thus by our inductive hypothesis, $\sigma^{-1}(\tau_\epsilon(k - 1))$ cannot equal any of $\{\sigma^{-1}(\sigma(1)), \sigma^{-1}(\sigma(2)), \dots, \sigma^{-1}(\sigma(k - 2))\} = \{1, 2, \dots, k - 2\}$. Therefore the only possible value that remains is $\sigma^{-1}(\tau_\epsilon(k - 1)) = k - 1$. By a similar argument, we obtain that $\sigma^{-1}(\tau_\epsilon(k)) = k$.

We note here that one possible value of $\sigma(i)$ has not been explicitly determined from Claim 4.1.5.2 when $\delta + \epsilon$ is even, namely $\sigma(\delta)$ when δ is odd. Of course this is easily resolved. Claim 4.1.5.2 tells us that for $\delta + \epsilon$ even, $\sigma(i) = \tau_\epsilon(i)$ for all $i < \delta$. Hence $\sigma(\delta) = \tau_\epsilon(\delta)$.

It remains to consider the case when

$$\delta + \epsilon \text{ is odd.}$$

We maintain this assumption for the rest of the proof of the proposition.

Claim 4.1.5.3. $\sigma = \tau_\epsilon$.

Proof of Claim 4.1.5.3. By assumption $\sigma^{-1}(1) > 2$, and therefore by Proposition 4.1.3, $\sigma^{-1}(1) \in \{\delta - 1, \delta\}$. We write then $\sigma^{-1}(1) = \delta + \epsilon' - 1$ with $\epsilon' \in \{0, 1\}$. If $\epsilon' \neq \epsilon$, then $\delta + \epsilon'$ is even, and by from Claim 4.1.5.2, we have that $\sigma^{-1} = \tau_{\epsilon'}$. Since $\tau_{\epsilon'}^{-1} = \tau_{\epsilon'}$, we would have that $\sigma = \tau_{\epsilon'}$, and we would therefore still obtain that $\epsilon' = \epsilon$.

Hence we may apply Claims 4.1.5.1 and 4.1.5.2 to both σ and σ' , yielding the following for $k \leq \delta/2$.

$$\begin{aligned} \sigma(k) &= k \text{ if } k \text{ is even} \\ \sigma(k) &= \delta + \epsilon - k \text{ if } k \text{ is odd} \\ \sigma(\delta + \epsilon - k) &= k \text{ if } k \text{ is odd} \end{aligned}$$

As $\delta + \epsilon$ is odd, it remains to determine σ on the set

$$A = \{i \mid (\delta + \epsilon)/2 < i \leq \delta \text{ and } i \text{ is odd}\}.$$

We already know that $\sigma[A] = A$. Moreover, since all the elements in A are odd, we deduce from Lemma 4.1.4 that σ either fixes or reverses A .

If σ fixes A , then $\sigma = \tau_\epsilon$. Thus we assume towards a contradiction that σ reverses A and that $|A| \geq 2$, so $\delta \geq 5$.

We consider first the case when $\epsilon = 1$. Since we are also assuming that $\delta + \epsilon$ is odd, we have that δ is even and $\max A = \delta - 1$. Under our assumptions then, σ would send the geodesic

type $(1, \delta - 1, \delta)$ to $(1, \min A, \delta)$, which must therefore satisfy the triangle inequality. Therefore $\min A \geq \delta - 1$ which would imply that $|A| = 1$, which is a contradiction.

Now suppose that $\epsilon = 0$. Then δ is odd and $\max A = \delta$. Under these assumptions, σ maps the geodesic type $(2, \min A - 2, \min A)$ to the triangle type $(2, \delta - \min A + 2, \delta)$. Once again we obtain from the triangle inequality a restriction:

$$\delta \leq \delta - \min A + 4.$$

Since $\min A$ is odd, we find that $\min A \leq 3$. However $\min A > \delta/2$, so then $\delta \leq 5$.

Thus $\delta = 5$, and $\sigma = \sigma^{-1} = (14)(35)(2)$. This permutation sends the geodesic type $(2, 2, 4)$ to the triangle type $(2, 2, 1)$ and the forbidden triple $(1, 1, 4)$ to the triple $(1, 4, 4)$. Therefore the triangle type $(2, 2, 1)$ is realized in Γ and the triangle type $(4, 4, 1)$ is not realized in Γ . So $K_1 \leq 2$ and Γ_4 contains no edge. By Fact 2.2.13 we find that Γ is antipodal. Thus Γ does not realize the triangle type $(5, 5, 2)$ and hence Γ^σ does not realize the triangle type $(3, 3, 2)$. This contradicts Proposition 2.2.1 of Section 2.2. \square

Proof of Proposition 4.1.2. This follows directly from Propositions 4.1.3, 4.1.4, Corollary 4.1.1, and Proposition 4.1.5. \square

4.2 Twistable graphs

We work towards showing the following main result:

Proposition 4.2.1. *Let σ be one of the permutations ρ, ρ^{-1}, τ_0 , or τ_1 , with $\delta \geq 3$. Then the metrically homogeneous graphs Γ of generic type whose images Γ^σ are also metrically homogeneous are precisely those with the numerical parameters K_1, K_2, C, C' as in Table 4.1 below.*

Note that we do not assume that Γ is of known type. If Γ is of known type, then its isomorphism type will be determined by its numerical parameters together with a set of Henson constraints. But by Proposition 4.2.1, twistability depends only on the value of the numerical parameters.

The values of the parameters given in the tables correspond to the realization or omission of certain triangle types which are either of the form $(k, k, 1)$ or of some fixed perimeter, and will be proved in that form.

σ	δ	K_1	K_2	C	C'
ρ	$\geq 3, < \infty$	1	δ	$2\delta + 2$	$2\delta + 3$
ρ^{-1}	$\geq 3, < \infty$	δ	δ	$3\delta + 1$	$3\delta + 2$
τ_ϵ	$\geq 3, < \infty$	$\lfloor \frac{\delta+\epsilon}{2} \rfloor$	$\lceil \frac{\delta+\epsilon}{2} \rceil$	$2(\delta + \epsilon) + 1$	$2(\delta + \epsilon) + 2$
τ_ϵ	$\geq 3, \equiv \epsilon \pmod{2}$	∞	0	$2\delta + 1$	$2(\delta + \epsilon) + 2$
<i>Exceptional Cases</i>					
$\sigma = \tau_1, \delta = 3 \text{ or } 4$					
τ_1	3	1	2	10	11
	3	1	2	9	10
	3	2	2	10	11
	4	1	3	11	14
	4	1	3	11	12
	4	2	3	11	14

Table 4.1: Twistable Metrically Homogeneous Graphs

We break up our analysis into two subsections: one addressing the necessity of these parameter values for twistability, and another addressing the sufficiency of these parameter values for twistability.

4.2.1 Necessity of the restrictions on the parameters

In this section we prove the following.

Proposition 4.2.2. *Let σ be one of the permutations ρ, ρ^{-1}, τ_0 or τ_1 , with $\delta \geq 3$. Then the metrically homogeneous graphs of generic type whose images Γ^σ are also metrically homogeneous graphs must have numerical parameters among those shown in Table 4.1.*

The numerical parameters were defined in Definition 2.2.6, Section 2.2.

We consider twists individually in the following order: $\sigma = \rho, \rho^{-1}, \tau_\epsilon$.

Lemma 4.2.1. *Let Γ be a metrically homogeneous graph of generic type such that Γ^ρ is also a metrically homogeneous graph. Then the associated numerical parameters K_1, K_2, C, C' for Γ must be $1, \delta, 2\delta + 2, 2\delta + 3$ respectively.*

Proof. The triangle types $(2, 2, 2)$ and $(1, 1, 2)$ must both be realized in Γ^ρ , by the definition of generic type. Therefore their inverse images $(1, 1, 1)$ and $(\delta, \delta, 1)$ must both be realized in Γ .

Thus we already know that

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

Consider a distance k realized in Γ_δ , that is, the triangle type (δ, δ, k) is realized in Γ . Under ρ , this is mapped to $(1, 1, \rho(k))$, so $\rho(k) \leq 2$. Thus $k = 1$ or δ . Since Γ_δ is connected (Fact 2.2.12) and $\delta \geq 3$, the distance δ cannot occur. Thus, Γ_δ has diameter at most 1. As the distance 1 occurs in Γ_δ , the diameter equals 1. It follows from Lemma 2.2.3 that $C = 2\delta + 2$ and $C' = 2\delta + 3$. \square

Lemma 4.2.2. *Let Γ be a metrically homogeneous graph of generic type such that $\Gamma^{\rho^{-1}}$ is also a metrically homogeneous graph. Then the associated numerical parameters K_1, K_2, C, C' for Γ must be $\delta, \delta, 3\delta + 1, 3\delta + 2$ respectively.*

Proof. For ease of notation, we write $\tilde{\Gamma} = \Gamma^{\rho^{-1}}$.

The graph $\tilde{\Gamma}$ is twistable by ρ , and is of generic type by Proposition 4.1.1, so its numerical parameters are given by Lemma 4.2.1. We denote them by $\tilde{K}_1, \tilde{K}_2, \tilde{C}, \tilde{C}'$.

We begin by showing the following.

Claim 4.2.2.1. $K_1 = K_2 = \delta$.

Proof of Claim 4.2.2.1. By Corollary 2.2.2, in order to prove the claim, it suffices to find triangle types of all odd perimeters less than $2\delta + 1$ which are not realized in Γ , as well as a triangle type of perimeter $2\delta + 1$ which is realized in Γ . We first find forbidden triangles of perimeter $2k - 1$ for $2 \leq k \leq \delta$.

As $\tilde{C} = 2\delta + 2$ and $\tilde{C}' = 2\delta + 3$ (or by the proof of that fact), we have that $\tilde{\Gamma}_\delta$ has diameter 1. Since the triangle type (δ, δ, δ) is not realized in $\tilde{\Gamma}$, the triangle type $(1, 1, 1)$ is not realized in Γ .

Thus we turn our attention to k satisfying $3 \leq k \leq \delta$ and consider the triangle type $(2, 2\lfloor k/2 \rfloor - 2, 2\lfloor k/2 \rfloor - 1)$. This triangle type has perimeter $2k - 1$ and is mapped under ρ^{-1} to $(1, \lfloor k/2 \rfloor - 1, \delta - \lfloor k/2 \rfloor + 1)$. This triple violates the triangle inequality, since $\lfloor k/2 \rfloor + \lfloor k/2 \rfloor = k$ and $k < \delta + 1$. Thus the triangle types $(2, 2\lfloor k/2 \rfloor - 2, 2\lfloor k/2 \rfloor - 1)$ must be forbidden in Γ .

To see that some triangle type of perimeter $2\delta + 1$ is indeed realized in Γ , we argue according to the parity of δ . If δ is even, then the triangle type $(1, \delta, \delta)$ is the image of $(\delta, \delta/2, \delta/2)$ under ρ^{-1} , which is of geodesic type. Thus the triangle type $(1, \delta, \delta)$ is realized in Γ . If δ is odd, we consider the triangle type $(3, \delta - 1, \delta - 1)$. This has the image $(\delta - 1, \frac{\delta-1}{2}, \frac{\delta-1}{2})$ under ρ^{-1} , which again is of geodesic type. Thus the triangle type $(3, \delta - 1, \delta - 1)$ is realized in Γ .

This proves the claim. In particular, $C > 2\delta + 1$.

We turn our attention now to C and C' . Here we use some additional structure theory. Since $K_2 = \delta$, we may apply Lemma 2.2.3 to get that $C = 2\delta + \delta' + 1$ and $C' = C + 1$ where δ' is the diameter of Γ_δ . It remains to show that $\delta' = \delta$, or in other words that a triangle of type (δ, δ, δ) occurs in Γ .

The permutation ρ^{-1} takes (δ, δ, δ) to $(\frac{\delta+\epsilon}{2}, \frac{\delta+\epsilon}{2}, \frac{\delta+\epsilon}{2})$, where ϵ is the parity of δ .

We first consider the case when δ is even. Proposition 2.2.1 gives us that $(\delta/2, \delta/2, 2)$ is realized in $\widetilde{\Gamma}$ and therefore the triangle type $(\delta, \delta, 4)$ is realized in Γ . Thus $\delta' = \text{diam}(\Gamma_\delta) \geq 4$. Moreover, Γ_δ is connected by Fact 2.2.12. Thus the distance 2 is realized in Γ_δ and hence the triangle type $(\delta/2, \delta/2, 1)$ is realized in $\widetilde{\Gamma}$. Therefore we may again apply Fact 2.2.12 to get that $\widetilde{\Gamma}_{\delta/2}$ is connected. The distance δ occurs in $\widetilde{\Gamma}_{\delta/2}$, as seen from the geodesic type $(\delta/2, \delta/2, \delta)$, so the distance $\delta/2$ also occurs in $\widetilde{\Gamma}_{\delta/2}$. Thus the triangle type $(\delta/2, \delta/2, \delta/2)$ is realized in $\widetilde{\Gamma}$ and the triangle type (δ, δ, δ) is realized in Γ .

We now consider the case when δ is odd. In this case we need the triangle type $(\frac{\delta+1}{2}, \frac{\delta+1}{2}, \frac{\delta+1}{2})$ to be realized in $\widetilde{\Gamma}$. If we can show that the diameter of $\widetilde{\Gamma}_{\frac{\delta+1}{2}}$ is at least $\frac{\delta+1}{2}$, then we may argue as we did for δ even.

Let ϵ' be the parity of $j = \frac{\delta+1}{2}$. We show that $\widetilde{\Gamma}$ contains the triangle type $(j, j, j + \epsilon')$. The value $j + \epsilon'$ is even and we claim that $j + (j + \epsilon')/2 \leq \delta$, or equivalently $3j + \epsilon' \leq 2\delta$. This is clearly true for $\delta \geq 5$, and for $\delta = 3$, we have that $j = 2$ and $\epsilon' = 0$, and thus this inequality holds for all odd δ . Thus Proposition 2.2.1 tells us that $(j, j, j + \epsilon')$ is realized in $\widetilde{\Gamma}$, and thus are done with the case when $\frac{\delta+1}{2}$ is even.

We are left then with the case δ odd and $\frac{\delta+1}{2}$ is odd. Note that this implies that $\delta \geq 5$. As in the case when δ even, we may deduce that $\widetilde{\Gamma}_{\frac{\delta+1}{2}}$ is connected. The distance $\frac{\delta+1}{2} + 1$ occurs in $\widetilde{\Gamma}_{\frac{\delta+1}{2}}$, since by Proposition 2.2.1 the triangle type $(\frac{\delta+1}{2}, \frac{\delta+1}{2}, \frac{\delta+1}{2} + 1)$ is realized in $\widetilde{\Gamma}$. Thus the connectivity of $\widetilde{\Gamma}_{\frac{\delta+1}{2}}$ yields that the distance $\frac{\delta+1}{2}$ is also realized in $\widetilde{\Gamma}_{\frac{\delta+1}{2}}$ and therefore (δ, δ, δ) is realized in Γ .

Our claim is now complete. □

Our analysis of the parameter values compatible with a twist of the form τ_ϵ is somewhat more involved, and involves some exceptional cases, as seen in Table 4.1. We break up our

analysis into a series of lemmas.

Lemma 4.2.3. *Let Γ be a metrically homogeneous graph of generic type and diameter at least 3 such that Γ^{τ_ϵ} is metrically homogeneous. Suppose moreover that Γ is bipartite. Then $\delta \equiv \epsilon \pmod{2}$ and $C_0 = 2(\delta + \epsilon) + 2$.*

Proof. When $\delta \not\equiv \epsilon \pmod{2}$, then $\tau_\epsilon^{-1} = \tau_\epsilon$ maps the geodesic $(\frac{\delta+\epsilon-1}{2}, \frac{\delta+\epsilon-1}{2}, 2\frac{\delta+\epsilon-1}{2})$ to either $(\frac{\delta+\epsilon-1}{2}, \frac{\delta+\epsilon-1}{2}, \delta + \epsilon - 2\frac{\delta+\epsilon-1}{2})$ or $(\delta + \epsilon - \frac{\delta+\epsilon-1}{2}, \delta + \epsilon - \frac{\delta+\epsilon-1}{2}, \delta + \epsilon - 2\frac{\delta+\epsilon-1}{2})$. In either case, this implies that a triangle type of odd perimeter in Γ is being mapped to a geodesic in Γ^{τ_ϵ} . This is a contradiction, since bipartite graphs have no triangle types of odd perimeter. Thus τ_ϵ is not a viable twist.

When $\delta \equiv \epsilon \pmod{2}$, the parities of the distances between elements of Γ are preserved under τ_ϵ . Thus the image Γ^{τ_ϵ} is bipartite, with the same parts as Γ .

We turn our attention now to C_0 and we show first that $C_0 \geq 2(\delta + \epsilon) + 2$. If $\epsilon = 0$, this holds by definition. If $\epsilon = 1$, then τ_ϵ maps the triangle type $(\delta, \delta, 2)$ to $(1, 1, 2)$ which is a geodesic and therefore is realized in Γ^{τ_ϵ} . Thus $C_0 > 2\delta + 2$, and therefore $C_0 \geq 2(\delta + \epsilon) + 2$.

Thus in order to show that $C_0 = 2(\delta + \epsilon) + 2$, it suffices to prove that any triangle type of perimeter at least $2(\delta + \epsilon) + 2$ is forbidden in Γ . Working towards a contradiction, we assume that such a triangle type is realized in Γ . Fact 2.2.15 then says that Γ has a triangle of type (δ, δ, d) with $d \geq 2\epsilon + 2$.

If $\epsilon = 0$, then we consider a pair of vertices $u, v \in \Gamma_\delta$ at distance d . Using homogeneity, we take u', v' adjacent to u and v respectively so that $d(u', v) \geq \min(d + 1, \delta - 1)$ and

$$d(u', v') \geq \min(d(u', v) + 1, \delta - 1) \geq \min(d + 2, \delta - 1).$$

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

As Γ is bipartite we have $u', v' \in \Gamma_{\delta-1}$ and thus u', v' and the basepoint form a triangle of type $(\delta - 1, \delta - 1, d')$. The permutation τ_0 maps this triangle type to $(1, 1, \tau_0(d'))$. As Γ^{τ_0} is bipartite we find $\tau_0(d') = 2$ and hence $d' = 2$. But as $\delta \equiv \epsilon \pmod{2}$, δ is even and thus

$$d' \geq \min(d + 2, \delta - 1) \geq 3,$$

a contradiction.

If $\epsilon = 1$, then our assumption and Fact 2.2.15 would imply that Γ contains the triangle type (δ, δ, d) for some $d \geq 4$. The permutation τ_1 sends this triangle type to $(1, 1, \tau_1(d))$ and as Γ^{τ_1} is bipartite we have $\tau_1(d) = 2$ and $d = 2$, which is a contradiction.

Therefore every triangle type of perimeter at least $2(\delta + \epsilon) + 2$ is forbidden, and thus $C_0 = 2(\delta + \epsilon) + 2$. \square

Lemma 4.2.4. *Let Γ be a metrically homogeneous graph of generic type such that Γ^{τ_ϵ} is metrically homogeneous. Suppose furthermore that Γ is not bipartite. Then one of the following holds:*

- *The unique distance occurring in $\Gamma_{\delta+\epsilon-1}$ is 2, and if $\epsilon = 0$ then $\delta \geq 4$;*
- *$\epsilon = 0, \delta = 3$; in this case, we have $\Gamma \simeq \Gamma_{1,2,7,8}^3$ (the generic antipodal graph of diameter 3);*
- *Γ is in one of the exceptional cases listed with $\epsilon = 1$ and $\delta \leq 4$, and the distance δ occurs in Γ_δ .*

Proof. We prove this result via a series of claims.

Claim 4.2.4.1. *If $\delta = 3$ and $\epsilon = 0$ then $\Gamma \simeq \Gamma_{1,2,7,8}^3$.*

Proof of Claim 4.2.4.1. By Fact 2.2.10, any triangle type realized in the canonical metrically homogeneous graph $\Gamma_{K_1, K_2, C, C'}^3$ with the same numerical parameters as Γ will also be realized in Γ . Thus any forbidden triangle types must be directly excluded by one of the parameters. The triangle type $(2, 3, 3)$ corresponds under τ_0 to the triple $(1, 1, 3)$, which violates the triangle inequality. Therefore the type $(2, 2, 3)$ must be excluded from Γ , and is hence excluded either by the value of a parameter K_1 or K_2 , or by the value of C_1 . As Γ is not bipartite, the type $(2, 2, 3)$ is not excluded by K_1 . By Definition 2.2.7, the triangle types excluded by K_2 have odd perimeter p and

$$p > 2K_2 + 2 \min(i, j, k).$$

Thus in our case, $7 > 2K_2 + 2 \min(2, 2, 3) = 2K_2 + 4$ and hence $K_2 = 1$. By the definition of K_1 we have $K_1 = 1$ as well, and as $\delta - 1 = 2$, Fact 2.2.13 gives a contradiction. So this triangle type is not excluded by K_2 . Thus the triangle type $(2, 2, 3)$ can only be excluded by the parameter

C_1 , that is, $C_1 = 7$. Now we may use the classification from [ACM16, Theorem 1] to identify Γ (Fact 2.2.10).

Claim 4.2.4.2. *If $\delta + \epsilon > 3$, then the only possible distances which may occur in $\Gamma_{\delta+\epsilon-1}$ are 2 and $\delta + \epsilon - 1$, and at least the distance 2 does occur.*

Proof of Claim 4.2.4.2. Consider a triangle type of the form $(\delta + \epsilon - 1, \delta + \epsilon - 1, k)$ realized in Γ with $1 \leq k \leq \delta$. The permutation τ_ϵ maps this triangle type to $(1, 1, \tau_\epsilon(k))$. Thus $\tau_\epsilon(k) \leq 2$ and k is either 2 or $\delta + \epsilon - 1$, as claimed. For $k = 2$, as $\delta + \epsilon > 3$ we have $(1, 1, \tau_\epsilon(k)) = (1, 1, 2)$, which is a geodesic type. So the distance 2 does occur.

Claim 4.2.4.3. *If $\epsilon = 0$ and $\delta \geq 4$, then the distance $\delta - 1$ is not realized in $\Gamma_{\delta-1}$.*

Proof of Claim 4.2.4.3. Suppose that $\epsilon = 0$ and that $\Gamma_{\delta-1}$ does realize the distance $\delta - 1$. Then $(\delta - 1, \delta - 1, \delta - 1)^{\tau_0} = (1, 1, 1)$ is realized in Γ^{τ_ϵ} . By Fact 2.2.13, the triangle type $(\delta - 1, \delta - 1, 1)$ is also realized in Γ^{τ_0} . Thus its inverse image $(1, 1, \delta - 1)$ must be in Γ , and therefore $\delta \leq 3$, a contradiction.

Claim 4.2.4.4. *If $\epsilon = 1$ and the distance δ is realized in Γ_δ , then $\delta \leq 4$ and Γ is one of the listed exceptional cases.*

Proof of Claim 4.2.4.4. We show first that there is a triangle of type $(2, 2, \delta)$ in Γ , and in particular $\delta \leq 4$. Note that since δ is realized in Γ_δ , we have $K_1 = 1$ for Γ^{τ_1} since (δ, δ, δ) in Γ corresponds to $(1, 1, 1)$ in Γ^{τ_1} . We may apply Proposition 2.2.1 to get that there is a triangle of type $(2, 2, 1)$ in Γ^{τ_1} , and therefore a triangle of type $(2, 2, \delta)$ in Γ . Therefore $\delta \leq 4$.

Case 1: $\delta = 3$

We are assuming that the triangle type $(3, 3, 3)$ is realized in Γ and we have also shown that the triangle type $(2, 2, 3)$ is realized in Γ , since $\delta = 3$. By Claim 4.2.4.2, the triangle type $(2, 3, 3)$ is also realized in Γ , and the triangle type $(1, 3, 3)$ is not realized in Γ . In particular, there are triangles of perimeters 7, 8, 9 in Γ and thus $C = 10$ and $C' = 11$. Moreover, $K_1 \leq K_2 \leq 2$. Since all the metrically homogeneous graphs of diameter 3 are known, we refer to the catalog in [ACM16] to deduce that $K_2 = 2$ and $K_1 = 1$ or 2. These correspond to the first and third exceptional cases in Table 4.1, as claimed.

Case 2: $\delta = 4$

In this case, by assumption, the triangle type $(4, 4, 4)$ is realized in Γ . Thus, $C_0 = 14$ and Γ is not antipodal. The triangle type $(3, 4, 4)$ would be mapped under τ_1 to $(1, 1, 3)$ and therefore must be forbidden. As this is the only possible triangle type of perimeter 11, we see that $C_1 \leq 11$. The triangle type $(2, 3, 4)$ is mapped to $(1, 2, 3)$ and thus is realized in Γ . Since this triangle type has perimeter 9, we also have that $C_1 = 11$.

The triangle type $(4, 4, 1)$ is forbidden from being realized in Γ since it would be mapped under τ_1 to $(1, 1, 4)$. We see however that the triangle type $(2, 2, 1)$ is realized in Γ , since it is mapped to $(2, 2, 4)$. Thus $K_2 < 4$ and $K_1 \leq 2$.

We may therefore apply Fact 2.2.13 to get that $(1, 3, 3)$ is realized in Γ , yielding $K_2 = 3$.

Once again, $K_1 = 1$ or 2 corresponds to the first and third exceptional cases for $\delta = 4$, as claimed.

Thus the claim holds in all cases.

We argue now that Lemma 4.2.4 follows from Claims 4.2.4.1, 4.2.4.2, 4.2.4.3, and 4.2.4.4. Recall from Proposition 4.1.1 that we may assume $\delta \geq 3$.

If $\delta + \epsilon < 4$, then $\delta = 3$ and $\epsilon = 0$, and we arrive at the second case mentioned the lemma with Claim 4.2.4.1 providing the additional information about Γ .

We suppose then that

$$\delta + \epsilon \geq 4.$$

If the distance $\delta + \epsilon - 1$ does not occur in Γ_δ then by Claim 4.2.4.2 the only distance occurring in Γ_δ is 2, as in the first case of our lemma. Thus we finally suppose that

$$\text{The distance } \delta + \epsilon - 1 \text{ occurs in } \Gamma_\delta.$$

Then Claim 4.2.4.3 shows that $\epsilon = 1$. As we are supposing that the distance δ is realized in Γ_δ , Claim 4.2.4.4 shows that we are in one of the corresponding exceptional cases with $\delta \leq 4$. \square

Lemma 4.2.5. *Let Γ be a metrically homogeneous graph such that Γ^{τ_0} is metrically homogeneous. Suppose moreover that Γ is not bipartite. Then Γ is antipodal and $K_1 = \lfloor \frac{\delta}{2} \rfloor$, $K_2 = \lceil \frac{\delta}{2} \rceil$.*

Proof. Note that for this lemma we are only working with $\epsilon = 0$.

In Lemma 4.2.4 the third case is excluded, and in the second case our lemma holds since the graph is $\Gamma_{1,2,7,8}^3$. Thus we restrict ourselves to the remaining case of Lemma 4.2.4, where the unique distance realized in $\Gamma_{\delta-1}$ is 2 and $\delta \geq 4$.

Claim 4.2.5.1. *For $u \in \Gamma_{\delta-1}, v \in \Gamma_{\delta}$, we have $d(u, v) \leq \delta - 2$.*

Proof of Claim 4.2.5.1. If $d(u, v) = \delta$, then we take $v' \in \Gamma_{\delta-1}$ adjacent to v . That would imply that $d(u, v') \geq \delta - 1$, but since both u and v' are in $\Gamma_{\delta-1}$, we have $d(u, v') = 2$. As $\delta \geq 4$, this is a contradiction.

If $d(u, v) = \delta - 1$, we find a v' adjacent to v such that $d(u, v') = \delta$. By the previous paragraph, $v' \notin \Gamma_{\delta}$. But then $v' \in \Gamma_{\delta-1}$, and since the unique distance realized in $\Gamma_{\delta-1}$ is 2, we would have that $\delta = 2$, a contradiction.

Claim 4.2.5.2. *Γ is antipodal.*

Proof of Claim 4.2.5.2. Working towards a contradiction, we assume there are two distinct points u, v in Γ_{δ} and we assume without loss of generality that $\delta' = d(u, v) = \text{diam}(\Gamma_{\delta})$. We notice first that there must be a $u' \in \Gamma_{\delta-1}$ adjacent to u with $d(u', v) \geq \min(\delta' + 1, \delta - 1)$. Indeed, if $\delta' < \delta$, then by the homogeneity of Γ_{δ} , we may take u' to be adjacent to u with $d(u', v) = \delta' + 1$. Then u' is in $\Gamma_{\delta-1}$. If on the other hand $\delta' = \delta$, we may simply take any u' in $\Gamma_{\delta-1}$ that is adjacent to u .

Take $u_1 \in \Gamma_{\delta-1}, v_1 \in \Gamma_{\delta}$ with $d(u_1, v_1)$ maximal; in particular, $d(u_1, v_1) \geq d(u, v')$. By the previous claim $d(u_1, v_1) < \delta$ and thus there is v' adjacent to v_1 with $d(u_1, v') = d(u_1, v_1) + 1$. By the choice of u_1, v_1 , we have $v' \notin \Gamma_{\delta}$, so $v' \in \Gamma_{\delta-1}$. But

$$d(u_1, v') \geq d(u', v) + 1 \geq \min(\delta' + 2, \delta) > 2$$

a contradiction.

Thus, Γ must be antipodal.

Claim 4.2.5.3. *$K_1 = \lfloor \frac{\delta}{2} \rfloor$ and $K_2 = \lceil \frac{\delta}{2} \rceil$.*

Proof of Claim 4.2.5.3. We begin by considering any triangle type of the form $(i, i, 1)$ which is realized in Γ . If $\min(i, \delta - i)$ is even then τ_0 sends this triangle type to $(i, i, \delta - 1)$. The triangle

inequality and the perimeter bound afforded by antipodality would then yield

$$\delta - 1 \leq 2i \leq \delta + 1.$$

If $\min(i, \delta - i)$ is odd, the corresponding inequalities found would be

$$\delta - 1 \leq 2(\delta - i) \leq \delta + 1$$

which are equivalent to the inequalities in the even case. Thus we have that $K_1 \geq \lfloor \frac{\delta}{2} \rfloor$ and $K_2 \leq \lceil \frac{\delta}{2} \rceil$. Recalling that $K_1 + K_2 = \delta$ for antipodal graphs — see, for example, [ACM16, pg. 13] — our claim is shown.

Claims 4.2.5.2 and 4.2.5.3 prove the lemma. \square

The following will be used in the proof of Lemma 4.2.7.

Lemma 4.2.6. *Let Γ be a metrically homogeneous graph such that $\tilde{\Gamma} = \Gamma^{\tau_1}$ is also metrically homogeneous. Suppose moreover that Γ satisfies the following conditions.*

- $\delta = 3$ or 4
- The distance δ is realized in Γ_δ
- The numerical parameters associated with Γ are those associated with one of the exceptional cases in Table 4.1, namely one of the following.

$K_1 \leq 2$	$K_2 = \delta - 1$	$C = \delta + 7$	$C' = 3\delta + 2$
$K_1 = 1$	$K_2 = \delta - 1$	$C = 2\delta + 3$	$C' = C + 1$

If the unique distance occurring in $\tilde{\Gamma}_\delta$ is 2 then $K_1 = 2$ and $\tilde{\Gamma}$ is also one of the exceptional cases listed, with parameters

$\tilde{K}_1 = 1$	$\tilde{K}_2 = \delta - 1$	$\tilde{C} = 2\delta + 3$	$\tilde{C}' = 2\delta + 4$
-------------------	----------------------------	---------------------------	----------------------------

Proof. If $K_1 = 1$, then the triangle type $(1, 1, 1)$ is realized in Γ , and therefore the triangle type (δ, δ, δ) is realized in $\tilde{\Gamma}$, contradicting the assumption that the unique distance in $\tilde{\Gamma}_\delta$ is 2. Thus, $K_1 = 2$.

In these cases the triangle type (δ, δ, δ) is realized in Γ , so the type $(1, 1, 1)$ is realized in $\widetilde{\Gamma}$, yielding

$$\widetilde{K}_1 = 1$$

The distance δ in $\widetilde{\Gamma}$ corresponds to the distance 1 in Γ , but the relation $d(x, y) = 1$ is not a pairing on Γ . That is, given a vertex $v \in \Gamma$, there is more than one vertex v' such that $d(v, v') = 1$. Thus $\widetilde{\Gamma}$ is not antipodal, and we may apply Fact to get

$$\widetilde{K}_2 \geq \delta - 1$$

The distance 1 does not occur in $\widetilde{\Gamma}_\delta$, since the triple $(1, 1, \delta) = (\delta, \delta, 1)^{\tau_1}$ violates the triangle inequality. Thus

$$\widetilde{K}_2 = \delta - 1$$

Since by assumption the unique distance in $\widetilde{\Gamma}_\delta$ is 2, Lemma 2.2.3 of Section 2.2 tells us that $\widetilde{C}' = 2\delta + 4$, and $\widetilde{C} = 2\delta + 3$ or $\widetilde{C} = 2\delta + 1$.

Thus it remains to show that

$$\widetilde{C} \neq 2\delta + 1$$

We argue that a triangle of type $(2, \delta - 1, \delta)$ is realized in $\widetilde{\Gamma}$.

The image $(2, \delta - 1, \delta)^{\tau_1} = (2, \delta - 1, 1)$ is a geodesic, and therefore is realized in Γ if $\delta = 4$. It is of type $(2, 2, 1)$ if $\delta = 3$. This is realized in Γ since $K_2 = 2$; hence $(2, \delta - 1, \delta)$ is realized in $\widetilde{\Gamma}$, leaving $\widetilde{C} = 2\delta + 3$.

This concludes the proof. □

Lemma 4.2.7. *Let Γ be a metrically homogeneous graph such that Γ^{τ_1} is metrically homogeneous and assume that Γ is not bipartite. Then either the parameters K_1, K_2, C, C' have the values $\lfloor \frac{\delta+1}{2} \rfloor, \lceil \frac{\delta+1}{2} \rceil, 2\delta + 3, 2\delta + 4$ respectively, or Γ is in one of the exceptional cases of Table 4.1 with $\delta \leq 4$.*

Proof. If the distance δ occurs in Γ_δ , then Lemma 4.2.4 provides our desired result. So we suppose that the unique distance occurring in Γ_δ is 2.

If the distance δ occurs in Γ^{τ_1} then by Lemma 4.2.4 the assumptions of Lemma 4.2.6 are applicable to Γ^{τ_1} , yielding that Γ is one of the exceptional cases of Table 4.1. Thus we also suppose that the unique distance realized in $(\Gamma_\delta)^{\tau_1}$ is 2.

Fact 2.2.15 then implies that no triangle realized in Γ or in Γ^{τ_1} can have perimeter greater than $2\delta + 2$.

Consider a triangle type of the form $(i, i, 1)$ realized in Γ . The permutation τ_1 sends this triangle type to either (i, i, δ) or $(\delta - i + 1, \delta - i + 1, \delta)$. The triangle inequality and the perimeter bounds applied to these two triangle types yield the following pairs of inequalities in the two cases, respectively.

$$\begin{aligned}\delta &\leq 2i \leq \delta + 2 \\ \delta &\leq 2(\delta - i + 1) \leq \delta + 2.\end{aligned}$$

Note that these two pairs of inequalities are in fact equivalent and may be written in the form

$$\left\lfloor \frac{\delta + 1}{2} \right\rfloor \leq i \leq \left\lceil \frac{\delta + 1}{2} \right\rceil$$

Given that Γ is not bipartite, the triangle type $(i, i, 1)$ must be realized for some finite i . Thus $\lfloor \frac{\delta+1}{2} \rfloor \leq K_1 \leq K_2 \leq \lceil \frac{\delta+1}{2} \rceil$.

We now consider separately the case when δ is odd and the case when δ is even.

If δ is odd, then the values K_1 and K_2 are squeezed to be

$$K_1 = K_2 = \frac{\delta + 1}{2}.$$

Since $\tau_1^{-1} = \tau_1$, the same applies for Γ^{τ_1} . Thus $(1, \frac{\delta+1}{2}, \frac{\delta+1}{2})$ is realized in Γ^{τ_1} and gets mapped under τ_1 to $(\frac{\delta+1}{2}, \frac{\delta+1}{2}, \delta)$, so Γ realizes a triangle type of perimeter $2\delta + 1$. The graph Γ also realizes the triangle type $(\delta, \delta, 2)$ of perimeter $2\delta + 2$. By Fact 2.2.15, Γ contains no triangle of perimeter larger than $2\delta + 2$. Thus,

$$C = 2\delta + 3 \text{ and } C' = 2\delta + 4.$$

If δ is even, then τ_1 maps the geodesic triangle type $(1, \delta/2, \delta/2 + 1)$ from Γ^{τ_1} to $(\delta/2, \delta/2 + 1, \delta)$ in Γ . This triangle type has perimeter $2\delta + 1$. Using the same reasoning as in the case of δ odd, we have again that $C = 2\delta + 3$ and $C' = 2\delta + 4$.

We address now the parameters K_1 and K_2 . In the even case, the above inequalities are

$$\delta/2 \leq K_1 \leq K_2 \leq \delta/2 + 1.$$

The geodesic type $(\delta/2, \delta/2, \delta)$ is mapped to $(1, \delta/2, \delta/2)$ and hence $K_1 = \delta/2$.

As Γ_δ has diameter 2, Lemma 2.2.2 yields that $\text{diam}(\Gamma_{\delta/2+1}) = \delta$. The same may be said for Γ^{τ_1} . Thus the triangle type $(\delta/2 + 1, \delta/2 + 1, \delta)$ is in Γ^{τ_1} and therefore the triangle type $(1, \delta/2 + 1, \delta/2 + 1)$ is in Γ . This tells us that $K_2 = \delta/2 + 1$. \square

Proof of Proposition 4.2.2. Lemmas 4.2.1 and 4.2.2 deal with the cases of ρ and ρ^{-1} . Lemma 4.2.3 treats τ_1 in the bipartite case. Lemmas 4.2.5 and 4.2.7 treat the cases of τ_0 and τ_1 respectively, in the non-bipartite case, including the exceptional cases. \square

4.2.2 Sufficiency of the restrictions of the parameters

We aim at the following.

Proposition 4.2.3. *Let σ be one of the permutations ρ, ρ^{-1}, τ_0 , or τ_1 , with $\delta \geq 3$. Let Γ be a metrically homogeneous graph whose numerical parameters are given in Table 4.1. Then Γ is twistable by the corresponding permutation σ .*

We begin with the case of ρ .

Lemma 4.2.8. *Let Γ be a metrically homogeneous graph of generic type with numerical parameters*

$$K_1 = 1 \qquad K_2 = \delta \qquad C = 2\delta + 2 \qquad C' = 2\delta + 3.$$

Then Γ^ρ is metrically homogeneous.

Proof. Here we make use of Fact 4.1.1 from Section 2.2, which tells us that in order to show that Γ^ρ is metrically homogeneous, it suffices to check that the triangle type (i, j, k) is not realized in Γ^ρ for $i + j < k$, and that the triangle type $(1, k, k + 1)$ is realized in Γ^ρ for $1 \leq k < \delta$. Thus we need to check that the corresponding triangle type $(i, j, k)^{\rho^{-1}}$ is not realized in Γ and that $(1, k, k + 1)^{\rho^{-1}}$ is realized in Γ .

Claim 4.2.8.1. *For $i + j < k$, the triangle type $(i, j, k)^{\rho^{-1}}$ is not realized in Γ .*

Proof of Claim 4.2.8.1. We argue according to the parities of i, j and k .

If i, j and k are all even, then the triple $(i, j, k)^{\rho^{-1}}$ violates the triangle inequality, since $i/2 + j/2 < k/2$. If i and j are even but k is odd, then the triple (i, j, k) still violates the triangle inequality. Indeed, then $i + j \leq k - 2$ and therefore $i + j + k \leq 2k - 2 \leq 2\delta - 2$. Then we conclude that

$$i/2 + j/2 \leq \delta - k/2 - 1 \leq \delta - \frac{k-1}{2}.$$

As $(i, j, k)^{\rho^{-1}} = (i/2, j/2, \delta - \frac{k-1}{2})$, this triple violates the triangle inequality.

If i and j are both odd, then we show that the perimeter of $(i, j, k)^{\rho^{-1}}$ is greater than $2\delta + 1$, thereby violating the C, C' bounds. Indeed, if k is also odd, then $(i, j, k)^{\rho^{-1}}$ is $(\delta - \frac{i-1}{2}, \delta - \frac{j-1}{2}, \delta - \frac{k-1}{2})$ and therefore has perimeter $3\delta - \frac{i+j+k-3}{2}$. Since $i + j < k \leq \delta$, we have that

$$3\delta - \frac{i+j+k-3}{2} > 3\delta - \frac{2k-3}{2} \geq 2\delta + 3/2.$$

If k is even, then $(i, j, k)^{\rho^{-1}} = (\delta - \frac{i-1}{2}, \delta - \frac{j-1}{2}, k/2)$ and thus has perimeter $2\delta + \frac{k-i-j}{2} + 1 > 2\delta + 1$.

If i and j have opposing parity, then we show that $(i, j, k)^{\rho^{-1}}$ violates the triangle inequality. We assume without loss of generality that i is odd and j is even. If k is even, then $(i, j, k)^{\rho^{-1}} = (\delta - \frac{i-1}{2}, j/2, k/2)$ and $j/2 + k/2 < \delta - \frac{i-1}{2}$ because $i + j + k < 2k - 1 \leq 2\delta - 1$. If k is odd, then $(i, j, k)^{\rho^{-1}} = (\delta - \frac{i-1}{2}, j/2, \delta - \frac{k-1}{2})$ and $(\delta - \frac{k-1}{2}) + j/2 < \delta - \frac{i-1}{2}$ because $i + j < k$.

Thus none of these triples may be realized in Γ and our claim is shown.

Claim 4.2.8.2. For $1 \leq k < \delta$, the triangle type $(1, k, k+1)^{\rho^{-1}}$ is in Γ .

Proof of Claim 4.2.8.2. By definition, $(1, k, k+1)^{\rho^{-1}}$ is either $(\delta, k/2, \delta - k/2)$ or $(\delta, \frac{k-1}{2}, \delta - \frac{k-1}{2})$. In either case, $(1, k, k+1)^{\rho^{-1}}$ is of geodesic type and therefore will indeed be realized in Γ .

The lemma follows. □

Lemma 4.2.9. Let Γ be a metrically homogeneous graph of generic type with finite diameter δ with associated numerical parameters

$$K_1 = \delta \qquad K_2 = \delta \qquad C = 3\delta + 1 \qquad C' = C + 1$$

Then $\Gamma^{\rho^{-1}}$ is metrically homogeneous.

Proof. Our reasoning proceeds as in the proof of the previous lemma. That is, we show that $\Gamma^{\rho^{-1}}$ is metrically homogeneous by verifying that the triangle type (i, j, k) for $i + j < k \leq \delta$ is

not realized in $\Gamma^{\rho^{-1}}$ while the triangle types $(1, k, k+1)$ for $1 \leq k < \delta$ are. We work instead with the graph Γ and the images of these triples under ρ .

Claim 4.2.9.1. *For $i + j < k \leq \delta$, the triangle type $(i, j, k)^\rho$ is not in Γ .*

Proof of Claim 4.2.9.1. If $k \leq \delta/2$, then $(i, j, k)^\rho = (2i, 2j, 2k)$. Since $2i + 2j < 2k$, this triple is not realized in Γ .

If $i \leq \delta/2 < j$, then

$$(i, j, k)^\rho = (2i, 2(\delta - j) + 1, 2(\delta - k) + 1).$$

Since $2i + 2(\delta - k) + 1 < 2(\delta - j) + 1$, this triple cannot be realized in Γ .

The remaining case to consider is $i, j \leq \delta/2 < k$. In this case,

$$(i, j, k)^\rho = (2i, 2j, 2(\delta - k) + 1).$$

The perimeter here is $2(\delta + i + j - k) + 1$ which is odd and less than $2\delta + 1$, and thus by Corollary 2.2.2 is also excluded from being realized in Γ .

Thus no such triangle type (i, j, k) will be realized in Γ .

Claim 4.2.9.2. *For $1 \leq k < \delta$, the triangle type $(1, k, k+1)^\rho$ is in Γ .*

Proof of Claim 4.2.9.2. If $k \neq \lfloor \delta/2 \rfloor$, then $(1, k, k+1)^\rho$ is of geodesic type and is realized in Γ .

If $k = \lfloor \delta/2 \rfloor$, then

$$(1, k, k+1)^\rho = (2, 2k, 2(\delta - (k+1)) + 1)$$

which has perimeter $2\delta + 1$. We apply Corollary 2.2.2 to see that this triangle type must be realized in Γ .

□

Lemma 4.2.10. *Let Γ be a metrically homogeneous graph of generic type. Suppose that $\delta = 3$ or 4 and that the parameters K_1, K_2, C, C' for Γ are among those in the table below. Then Γ^{τ_1} is also metrically homogeneous.*

Proof. By Lemma 4.1.1 the structure Γ^{τ_1} will be a metrically homogeneous graph if it does not realize any triple violating the triangle inequality, and does realize all geodesic types of the form $(1, k, k+1)$ with $1 \leq k < \delta$.

δ	K_1	K_2	C	C'	δ	K_1	K_2	C	C'
3	∞	0	7	10	4	1	3	11	14
3	1	2	10	11	4	1	3	11	12
3	1	2	9	10	4	2	3	11	12
3	2	2	9	10	4	2	3	11	14
3	2	2	10	11					

Thus for $\delta = 3$, the triple $(1, 1, 3)$ must be omitted by Γ^{τ_1} and the triangle types $(1, 1, 2)$ and $(1, 2, 3)$ must be realized by Γ^{τ_1} .

For $\delta = 4$, the triples $(1, 1, 3)$, $(1, 1, 4)$ and $(1, 2, 4)$ must be omitted by Γ^{τ_1} while the triangle types $(1, 1, 2)$, $(1, 2, 3)$, $(1, 3, 4)$ must be realized by Γ^{τ_1} .

This translates into the following restrictions for Γ when $\delta = 3$:

Forbidden	Realized
$(1,3,3)$	$(1,2,3), (2,3,3)$

and the following restrictions for Γ when $\delta = 4$:

Forbidden	Realized
$(3,4,4), (1,4,4), (1,2,4)$	$(2,4,4), (1,3,4), (2,3,4)$

Equivalently, setting aside geodesic types and triples violating the triangle inequality, we must verify that Γ satisfies the following four conditions.

- $K_2 < \delta$;
- Γ realizes the triangle type $(\delta, \delta, 2)$;
- for $\delta = 4$, Γ realizes the triangle type $(2, 3, 4)$;
- for $\delta = 4$, Γ does not realize the triangle type $(3, 4, 4)$.

Indeed, in every line in our table above, $K_2 < \delta$. In addition, whenever $\delta = 4$ we have $C = 11$ and therefore the triangle type $(3, 4, 4)$ will be forbidden. This disposes of the first and last conditions.

In every line in our table, $C_0 \geq 2\delta + 4$, and therefore Γ contains a triangle of perimeter $2\delta + 2$. By Fact 2.2.15, Γ_δ realizes some distance $d \geq 2$. If $K_1 = 1$, then Facts 2.2.12 and 2.2.13 tell us that Γ_δ is connected and therefore contains a pair of vertices at distance 2. If $K_1 > 1$, then Fact 2.2.14 tells us that Γ_δ is connected by the edge relation $d(x, y) = 2$. Thus in either case the distance 2 occurs in Γ_δ and our second condition is satisfied.

Suppose then that $\delta = 4$. To conclude our proof, we must show that the triangle type $(2, 3, 4)$ is realized in Γ . Thus we must find vertices $u \in \Gamma_3$ and $v \in \Gamma_4$ such that $d(u, v) = 2$.

In every line in our table, $K_2 = 3$, and therefore Γ_3 contains an edge. Fact 2.2.12 tells us then that Γ_3 is connected. Moreover we also see that Γ_3 has diameter at least 2, since Γ_2 has diameter 4. Thus Γ_3 realizes the distance 2. Take then $v \in \Gamma_4$ and define I_v to be the set of neighbors of v in Γ_3 . If $I_v = \Gamma_3$ then by homogeneity every vertex of Γ_4 is adjacent to every vertex of Γ_3 and Γ_4 has diameter at most 2, a contradiction. So $\Gamma_3 \neq I_v$.

Since Γ_3 is connected, there is a vertex $u \in \Gamma_3 \setminus I_v$ adjacent to some vertex $v' \in I_v$. It then follows that $d(u, v) = 2$ and u, v and the basepoint form the desired triangle. \square

Lemma 4.2.11. *Let Γ be a bipartite metrically homogeneous graph of generic type with diameter $\delta \equiv \epsilon \pmod{2}$ where $\epsilon = 0$ or 1 and $C_0 = 2(\delta + \epsilon) + 2$. Then Γ^{τ^ϵ} is metrically homogeneous.*

Proof. By Fact 4.1.1, it suffices to check that the triangle types (i, j, k) for

$$i + j < k \leq \delta$$

are not in Γ^{τ^ϵ} and the triangle types $(1, k, k + 1)$ for $1 \leq k < \delta$ are in Γ^{τ^ϵ} . We work in Γ with their images under τ_ϵ .

Claim 4.2.11.1. *For $i + j < k \leq \delta$, the triangle type $(i, j, k)^{\tau_\epsilon}$ is not in Γ .*

Proof of Claim 4.2.11.1. We observe that since $\delta \equiv \epsilon \pmod{2}$, the permutation τ_ϵ preserves parity. Therefore, if $i + j + k$ is odd then $(i, j, k)^{\tau_\epsilon}$ is forbidden, as there are no triangle types of odd perimeter in bipartite graphs.

If i, j , and k are all even, then they will all be fixed by τ_ϵ . So the triangle type (i, j, k) is excluded from Γ .

Thus it remains to consider the case in which one of i, j, k is even and the other two are odd.

If i and j are odd and k is even, then

$$j^{\tau_\epsilon} + k^{\tau_\epsilon} = \delta + \epsilon - j + k < \delta + \epsilon - i = i^{\tau_\epsilon}$$

and once again by the triangle inequality the triple $(i, j, k)^{\tau_\epsilon}$ will not be realized in Γ .

Finally, suppose that i is even and j and k are odd. In this case,

$$i^{\tau_\epsilon} + k^{\tau_\epsilon} = \delta + \epsilon + i - k < \delta + \epsilon - j = j^{\tau_\epsilon}$$

so again the triple $(i, j, k)^{\tau_\epsilon}$ violates the triangle inequality.

Claim 4.2.11.2. *For $1 \leq k < \delta$, the triangle type $(1, k, k + 1)^{\tau_\epsilon}$ is in Γ .*

Proof of Claim 4.2.11.2. If k is even, then

$$(1, k, k + 1)^{\tau_\epsilon} = (\delta + \epsilon - 1, k, \delta + \epsilon - (k + 1))$$

which is of geodesic type and therefore is realized in Γ .

If k is odd, then

$$(1, k, k + 1)^{\tau_\epsilon} = (\delta + \epsilon - 1, \delta + \epsilon - k, k + 1)$$

We consider the two values of ϵ separately.

Suppose first that

$$\epsilon = 0.$$

Then by hypothesis the graph Γ is antipodal and the triangle type $(1, k, k + 1)^{\tau_\epsilon}$ is $(\delta - 1, \delta - k, k + 1)$. Replacing one of the vertices v of this triangle type by its antipodal vertex v' yields the triangle type $(1, k, k + 1)$, by the antipodal law (Fact 2.2.8). Therefore the original triangle type must in be Γ since this triangle is of geodesic type.

Now suppose

$$\epsilon = 1.$$

In this case, $C_0 = 2\delta + 4$ and $(i, j, k)^{\tau_\epsilon} = (\delta, \delta - k + 1, k + 1)$. Moreover, we may apply Fact 2.2.15 to get that $\text{diam}(\Gamma_\delta) = 2$. Since $(\delta, \delta - k + 1, k + 1)$ is invariant under the substitution of $\delta - k$ for k , we may assume that $k \leq \delta/2$. Applying Lemma 2.2.2 then, we get that $\text{diam}(\Gamma_{\delta - k + 1}) = 2k$. We may take $u, v \in \Gamma_{\delta - k + 1}$ at distance $2k$ and $u', v' \in \Gamma_\delta$ at distance $k - 1$ from u and v respectively. It follows that $d(u', v') \geq 2$ and hence $d(u', v') = 2$. This implies that $d(u, v') \leq k + 1$. By the triangle inequality, $d(u, v') \geq k + 1$, and therefore $d(u, v') = k + 1$. The triangle formed by u, v' and the basepoint has type $(\delta - k + 1, \delta, k + 1)$, and hence this triangle type is indeed in Γ .

□

Lemma 4.2.12. *Let Γ be a metrically homogeneous graph of generic type with diameter δ , with the numerical parameters*

$$K_1 = \left\lfloor \frac{\delta + \epsilon}{2} \right\rfloor \quad K_2 = \left\lceil \frac{\delta + \epsilon}{2} \right\rceil \quad C = 2(\delta + \epsilon) + 1 \quad C' = C + 1.$$

Then Γ^{τ_ϵ} is metrically homogeneous.

Proof. By Fact 4.1.1, it suffices to show that violations of the triangle equality are not in Γ^{τ_ϵ} and geodesics of the form $(1, k, k + 1)$ for $1 \leq k < \delta$ are in Γ^{τ_ϵ} . We work in Γ with the images of these triples under τ_ϵ . We begin with the first point.

Claim 4.2.12.1. *For $i + j < k \leq \delta$, the triple $(i, j, k)^{\tau_\epsilon}$ is not in Γ ;*

Proof of Claim 4.2.12.1. For brevity, we refer to the parity of $\min(h, \delta + \epsilon - h)$ as the $(\delta + \epsilon)$ -parity of h . Note that by definition h and h^{τ_ϵ} have the same $(\delta + \epsilon)$ -parity. We break our argument into cases, based on the relative $(\delta + \epsilon)$ -parities of i , j , and k .

Case 1. The $(\delta + \epsilon)$ -parity of i and j are both even.

If the $(\delta + \epsilon)$ -parity of k is also even, then $(i, j, k)^{\tau_\epsilon} = (i, j, k)$, which violates the triangle inequality.

If the $(\delta + \epsilon)$ -parity of k is odd and $\delta + \epsilon - k \geq k$, then $(i, j, k)^{\tau_\epsilon}$ still violates the triangle inequality.

If the $(\delta + \epsilon)$ -parity of k is odd and $\delta + \epsilon - k < k$, then $k > (\delta + \epsilon)/2$.

In that case, $\min(k, \delta + \epsilon - k) = \delta + \epsilon - k$ is odd. Then the perimeter of $(i, j, k)^{\tau_\epsilon}$ is $i + j + \delta + \epsilon - k < \delta + \epsilon$. Since $K_1 = \lfloor \frac{\delta + \epsilon}{2} \rfloor$, the perimeter of $(i, j, k)^{\tau_\epsilon}$ must be even. Thus i^{τ_ϵ} and j^{τ_ϵ} have opposite parity, and we assume without loss of generality that i^{τ_ϵ} is even and j^{τ_ϵ} is odd.

From this we infer that $(\delta + \epsilon)$ is odd, that i is even and $i^{\tau_\epsilon} = i < \frac{\delta + \epsilon}{2}$, and that $j^{\tau_\epsilon} = j > \frac{\delta + \epsilon}{2}$.

We therefore conclude that

$$i^{\tau_\epsilon} + k^{\tau_\epsilon} = i + \delta + \epsilon - k < i + \delta + \epsilon - (i + j) = \delta + \epsilon - j < j = j^{\tau_\epsilon}$$

and hence the triple $(i, j, k)^{\tau_\epsilon}$ violates the triangle inequality.

Case 2. The $(\delta + \epsilon)$ -parity of i and j are both odd.

If the $(\delta + \epsilon)$ -parity of k is even, then the perimeter of $(i, j, k)^{\tau_\epsilon}$ is $2(\delta + \epsilon) + k - i - j > 2$ which is forbidden by the C, C' bounds.

Suppose the $(\delta + \epsilon)$ -parity of k is odd. In this case, the perimeter of $(i, j, k)^{\tau_\epsilon}$ is

$$3(\delta + \epsilon) - (i + j + k) > 3(\delta + \epsilon) - 2k \geq 2K_2 + 2(\delta_\epsilon - k) = 2K_2 + 2k^{\tau_\epsilon}$$

and thus violates the condition associated to K_2 .

Without loss of generality, the remaining case is the following.

Case 3. The $(\delta + \epsilon)$ -parities of i and j are even and odd respectively.

If the $(\delta + \epsilon)$ -parity of k is odd, then we have that

$$i^{\tau_\epsilon} + k^{\tau_\epsilon} = i + (\delta + \epsilon) - k < (\delta + \epsilon) - j = j^{\tau_\epsilon}$$

and therefore this triple violates the triangle inequality.

If the $(\delta + \epsilon)$ -parity of k is even then the perimeter of $(i, j, k)^{\tau_\epsilon}$ is

$$i + (\delta + \epsilon - j) + k \geq (\delta + \epsilon) + 2i \geq 2K_2 + 2i$$

and the bound associated with K_2 is violated.

We may finally assume that

$$i > \frac{\delta + \epsilon}{2} \quad j \leq \frac{\delta + \epsilon}{2} \quad k > \frac{\delta + \epsilon}{2}.$$

Then the perimeter of $(i, j, k)^{\tau_\epsilon}$ is

$$i + (\delta + \epsilon) - j + k \geq (\delta + \epsilon) + 2i > 2(\delta + \epsilon)$$

which violates the C bounds.

We conclude then that every such triple $(i, j, k)^{\tau_\epsilon}$ is forbidden in Γ .

Claim 4.2.12.2. For $1 \leq k < \delta$ the triangle type $(1, k, k + 1)^{\tau_\epsilon}$ is in Γ .

Proof of Claim 4.2.12.2. Note that $\min(k, \delta + \epsilon - k) \leq \lfloor \frac{\delta + \epsilon}{2} \rfloor$.

If $\min(k, \delta + \epsilon - k) = \lfloor \frac{\delta + \epsilon}{2} \rfloor$ then $(1, k, k + 1)^{\tau_\epsilon} = (1, k, k + 1)$, with k and $k + 1$ either fixed for swapped. So we may assume

$$\min(k, \delta + \epsilon - k) < (\delta + \epsilon)/2.$$

If $\min(k, \delta + \epsilon - k)$ is even, then $(1, k, k + 1)^{\tau^\epsilon} = (\delta + \epsilon - 1, k, \delta + \epsilon - k - 1)$. This is itself of geodesic type, and therefore must be realized in Γ .

Now suppose that $\min(k, \delta + \epsilon - k)$ is odd. If $\epsilon = 0$, then Γ is antipodal and the triangle type under consideration is $(k + 1, \delta - k, \delta - 1)$.

We may therefore replace one of the vertices of this triangle types with its opposite to obtain the triangle type $(1, k, k + 1)$. Since the latter triangle type is realized in Γ , the former one is as well.

If $\epsilon = 1$, then we require the triangle type $(\delta, \delta - k + 1, k + 1)$ to be realized in Γ . We follow the proof of Claim 4.2.8.2 in the proof of Lemma 4.2.11. We argue that Γ_δ has diameter 2, that we may suppose without loss of generality that $k \leq \delta/2$, and that there are vertices $u, v \in \Gamma_{\delta-k+1}$ at distance $2k$. Then taking $u', v' \in \Gamma_\delta$ at distance $k - 1$ from u, v respectively, we find $d(u', v') = 2$ and argue as previously that the triangle type formed by u, v' and the basepoint has the desired type. This proves the claim.

These two claims show that Γ^{τ^ϵ} is indeed a metrically homogeneous graph. □

Proof of Proposition 4.2.3. This follows immediately from Lemmas 4.2.8, 4.2.9, 4.2.10, 4.2.11, and 4.2.12. □

We may therefore conclude the following.

Proof of Proposition 4.2.1. This follows immediately from Propositions 4.2.2 and 4.2.3. □

We now combine Propositions 4.1.2 and 4.2.1 to get the following.

Theorem 3. *Let $\delta \geq 2$ be fixed (potentially infinite), and let σ be a non-trivial permutation of the language of metrically homogeneous graphs of diameter δ .*

If there is a metrically homogeneous graph Γ of generic type such that Γ^σ is again a metrically homogeneous graph, then σ is one of the permutations ρ, ρ^{-1}, τ_0 , or τ_1 from Proposition 4.1.2.

Conversely, if σ is one of the permutations ρ, ρ^{-1}, τ_0 , or τ_1 , with $\delta \geq 3$, then there is a metrically homogeneous graph Γ for which Γ^σ is again a metrically homogeneous graph. Furthermore, the metrically homogeneous graphs Γ whose images Γ^σ are also metrically homogeneous are precisely those with the numerical parameters K_1, K_2, C, C' as in Table 4.1.

4.3 Algebra of an Age

Here we will consider the graded \mathbb{Q} -algebra \mathcal{A}^G introduced by Peter Cameron in the study of growth rates of profile functions, measuring the number of orbits of G on unordered sets. Here G will be the group $\text{Aut}(\Gamma)$ with Γ a metrically homogeneous graph, so the number of orbits of G on sets of order n is the number of structures on n vertices embedding into Γ (as a metric space), taken up to isomorphism.

We aim to show (Theorem 4) that this algebra is a polynomial algebra when Γ is a known metrically homogeneous graph of generic type, that is, one of the graphs

$$\Gamma_{K_1, K_2, C, C', S}^\delta$$

with admissible parameters, where K_1 is finite and either $C > 2\delta + 1$, or else $C = 2\delta + 1$ with δ even (i.e., setting aside the bipartite case, and the antipodal case for δ odd).

The algebra in question may be defined as follows.

Definition 4.3.1. [Cam97] Let Ω be a set, and G a group acting on Ω . The reduced incidence algebra \mathcal{A} associated with the partial order of finite subsets of Ω [Rot64] is the graded \mathbb{Q} -algebra defined as follows. For each n let V_n be the vector space of \mathbb{Q} -valued functions on the set of n -element subsets of Ω . Then $\mathcal{A} = \bigoplus V_n$ with multiplication determined by

$$(fg)(X) = \sum_{X=X_1 \sqcup X_2} f(X_1)g(X_2).$$

Then G acts naturally on \mathcal{A} and \mathcal{A}^G denotes the subalgebra of G -invariant functions.

Remark 4.3.1. Equivalently, $\mathcal{A} = \bigoplus V_n^G$ where V_n^G is the space of functions constant on G -orbits, which may be identified with the space of functions on the G -orbits. If G is the automorphism group of a homogeneous structure Γ , then the orbits on sets of order n are the isomorphism types of substructures of Γ of order n , which make up the so-called ‘‘age’’ of Γ .

Our main result is the following.

Theorem 4. Let $(\delta, K_1, K_2, C, C', S)$ be an admissible parameter sequence with K_1 and δ finite, and let Γ be the corresponding metrically homogeneous graph, with automorphism group G . If $C = 2\delta + 1$, suppose that δ is even. Then the associated algebra \mathcal{A}^G is a polynomial algebra in infinitely many variables.

Cameron gave a criterion in purely structural terms sufficient to establish that the algebra is polynomial, with many applications, among them the case in which the structure Γ is the random graph. We will show that his method applies also in our general case.

Cameron's Criterion

We present a modified version of Cameron's framework, with slightly narrower assumptions than his.

Definition 4.3.2. *Let \mathcal{C} be a class of finite structures, closed under isomorphism. We write $A \subseteq B$ for the partial substructure relation (e.g., subgraph is partial substructure, whereas induced subgraph is a substructure).*

1. *A decomposition operator for \mathcal{C} consists of a binary operation $+$ on \mathcal{C} satisfying the following conditions.*

- *Functoriality: On isomorphisms of structures in \mathcal{C} , we have that for any pair of isomorphisms $i: A \rightarrow A'$ and $j: B \rightarrow B'$, the operator $+$ satisfies*

$$A + B \simeq A' + B';$$

- *Additivity: $|A + B| = |A| + |B|$;*
- *Unique decomposition: The commutative semigroup $(\mathcal{C}, +)$ is freely generated by its indecomposable elements.*

2. *A decomposition operator for \mathcal{C} is free if there is a partial order on \mathcal{C} satisfying the following.*

For A in \mathcal{C} , if A is partitioned into induced substructures B_1, \dots, B_k , then

$$B_1 + \dots + B_k \leq A.$$

Remark 4.3.2. *If a decomposition operator on \mathcal{C} is free then there is a canonical partial order \leq associated with the theory. Namely, one considers the transitive closure of the relation*

$$B \leq^+ A$$

defined on \mathcal{C} by

$B = B_1 + \cdots + B_k$ for some partition of A into induced substructures.

In particular, this partial order is also invariant under isomorphism. In practice however, one proves freeness by specifying a suitable partial order.

Example 4.3.1. *The decomposition of graphs as disjoint sums of connected graphs is a decomposition theory. It is free with respect to the subgraph relation.*

More subtle examples are found in [Cam97].

The point of this is the following.

Theorem 4.3.1. [Cam97, Theorem 2.1] *If $G = \text{Aut}(\Gamma)$ is the automorphism group of a homogeneous structure for a finite relational language, and the age \mathcal{C} of Γ has a free decomposition operator, then the algebra \mathcal{A}^G is the polynomial algebra with generators corresponding to the isomorphism types of indecomposable elements of \mathcal{C} .*

The statement given in [Cam97] is phrased in more general terms.

Example 4.3.2. [Cam97, Example 1] *The algebra associated with the random graph is a polynomial algebra.*

Remark 4.3.3. *In the case of graphs, a dual decomposition theory is obtained by switching edges and non-edges. The dual relation to “subgraph” is “supergraph.”*

The dual theory is free with respect to the supergraph relation.

The following is immediate.

Lemma 4.3.1. *Let \mathcal{C} have the free decomposition operator $+$ and let \mathcal{C}' be a $+$ -closed hereditary subset of \mathcal{C} . Then $+$ is a free decomposition operator for \mathcal{C}' .*

Example 4.3.3. [Cam97, Example 1 (cont.)] *The algebra associated with the generic K_n -free graph is a polynomial algebra.*

Direct sum operations

A straightforward generalization of the decomposition theories considered above for graphs is the following.

Definition 4.3.3. Let \mathcal{L} be a relational language and E a distinguished binary relation symbol in \mathcal{L} . For \mathcal{L} -structures A, B , define

$$A +_E B$$

to be the structure consisting of the disjoint union of A and B together with all relations $E(a, b)$ and $E(b, a)$, for (a, b) in $A \times B$.

Lemma 4.3.2. With the notation of Definition 4.3.3, the operation $+_E$ is a free decomposition operator on the class of finite \mathcal{L} -structures.

Proof. The operation is clearly functorial and additive.

If we associate to each structure $A \in \mathcal{C}$ the graph A^c which is the graph complement of A with edge relation E , then a decomposition of A corresponds to a decomposition of A^c as a disjoint sum. So unique decomposition follows.

Similarly, if A is partitioned into induced subgraphs B_1, \dots, B_k , then the disjoint sum B_1^c, \dots, B_k^c is contained in A^c . So we define $B \leq^+ A$ on \mathcal{C} by $B^c \subseteq A^c$, and freeness follows. \square

Corollary 4.3.2.1. Let \mathcal{C} be a hereditary class of binary relational structures. Let E be a distinguished binary symmetric relation in the language. If \mathcal{C} is closed under the operation $+_E$, then this operation provides a free decomposition theory for \mathcal{C} .

In practice, the case one has in mind in the above is the following: \mathcal{C} is the class of finite substructures of a homogeneous structure in a binary relational language L . The language L consists of names for the orbits on pairs of distinct elements. Furthermore, \mathcal{C} should have a transitive automorphism group (otherwise, the single relation E would be replaced by a finite set of relations, complicating the notation).

In the case of graphs, E is either the edge or the non-edge relation.

4.3.1 The case of metrically homogeneous graphs

The question now becomes, what sorts of generalized disjoint sum operations are available in the language of metrically homogeneous graphs. Here we replace the notation $+_E$ by the notation $+_i$, where i is the distance corresponding to the binary relation E .

The following is implicit in [ABH⁺17].

Lemma 4.3.3. *Let Γ be a 3-constrained metrically homogeneous graph with parameters $(\delta, K_1, K_2, C, C')$, where $\delta \geq 3$. Let $M \in [\delta]$. Then the following are equivalent.*

- *The class \mathcal{C} of finite substructures of Γ is closed under the operation $+_M$.*
- $\max(K_1, \delta/2) \leq M \leq \min(K_2, (C - \delta - 1)/2)$.

Proof. As any triangles occurring in a composition $A+_MB$ have type (M, M, i) for some distance $i \leq \delta$, and all such triangles occur in some composition, this is equivalent to the requirement that all triangles of type (M, M, i) embed into Γ , and thus is covered by [ABH⁺17, Observation 4.1]. We give some additional details for the proof.

The condition

$$M \geq \delta/2$$

is necessary and sufficient to ensure that structures in $\mathcal{C}+_M\mathcal{C}$ satisfy the triangle inequality.

Similarly, the condition $2M+\delta < C$ is necessary and sufficient to ensure that $\mathcal{C}+_M\mathcal{C}$ respects the perimeter bound.

It remains to consider constraints on triangles of odd perimeter. By considering triangles of type $(M, M, 1)$, we find that the conditions

$$K_1 \leq M \leq K_2$$

are necessary.

It remains to check their sufficiency. There are three conditions on triangles of type (M, M, i) corresponding to the parameters K_1, K_2 .

$$2M + i \geq K_1 \qquad M + i \leq 2K_2 + 2M \qquad 2M \leq 2K_2 + 2i$$

If $K_1 \leq M \leq K_2$, the first and third inequalities are clearly satisfied, and the second inequality becomes

$$i \leq 2K_2 + M.$$

We know that $M \geq K_1$, and in particular this implies that K_1 is finite. The conditions for admissibility imply in all cases that $2K_2 + K_1 \geq \delta$. Since $M \geq K_1$ and $i \leq \delta$, it follows that $i \leq 2K_2 + M$, as required.

The lemma follows. □

We require a similar lemma for the general case, in which Henson constraints occur.

Lemma 4.3.4. *Let Γ be a metrically homogeneous graph of generic type and of known type, with associated parameters $(\delta, K_1, K_2, C, C', \mathcal{S})$, where $\delta \geq 3$. Let $M \in [\delta]$. Then the following are equivalent.*

- *The class \mathcal{C} of finite substructures of Γ is closed under the operation $+_M$.*
- *$\max(K_1, \delta/2) \leq M \leq \min(K_2, (C - \delta - 1)/2)$ and in addition*
 - *If there is a constraint $H \in \mathcal{S}$ in which the distance δ occurs, then $M < \delta$.*

Proof. In general there are two notions of Henson constraint which apply: one in the case $C = 2\delta + 1$, and one in the remaining cases.

When $C = 2\delta + 1$, our conditions imply $M = \delta/2$ and in particular δ is even. In this setting, the Henson constraints involve distances 1 and $\delta - 1$, and as $M \neq 1, \delta - 1$ in this case, the additional condition is both vacuous and unnecessary, and the previous lemma suffices.

So we come to the main case in which $C > 2\delta + 1$ and \mathcal{S} is a family of $(1, \delta)$ -spaces. In this case, we have $M \geq \delta/2 > 1$, so if the distance δ does not occur in a Henson constraint, then once again the additional constraint is both vacuous and unnecessary.

So we come down to the case in which the distance δ does occur in some (minimal forbidden) Henson constraint H . Then H is itself $+_\delta$ -decomposable, and as the factors are not forbidden, we require $M < \delta$. Conversely, if $M < \delta$, then no conflicts can arise.

This completes consideration of all cases. □

The following is closely related to [ABH⁺17, Lemma 5.1]: the conditions for the completion process used there are slightly more restrictive than those required here, where we take only disjoint sums.

Lemma 4.3.5. *The following conditions on an admissible sequence of parameters $(\delta, K_1, K_2, C, C', \mathcal{S})$ are equivalent.*

- *There is a parameter M for which $\mathcal{A}_{K_1, K_2, C, C', \mathcal{S}}^\delta$ is $+_M$ -closed.*
- *δ and K_1 are finite. If $C = 2\delta + 1$, then δ is even.*

Proof. We refer to the conditions on M given in Lemma 4.3.4.

Suppose first a suitable parameter M exists. The conditions $M \geq K_1, \delta/2$ imply that K_1 and δ are both finite. If $C = 2\delta + 1$, the conditions $\delta/2 \leq M \leq (C - \delta - 1)/2$ imply that $M = \delta/2$ and δ is even.

Conversely, with δ and K_1 finite, we use the minimum value

$$M = \max(K_1, \lceil \delta/2 \rceil).$$

So we first require the numerical conditions

$$\max(K_1, \lceil \delta/2 \rceil) \leq \min(K_2, (C - \delta - 1)/2).$$

We know that $K_1 \leq K_2$ by definition, and $2\lceil \delta/2 \rceil \leq \delta + 1 \leq C - 2\delta - 1$ unless $C = 2\delta + 1$, and in this case as δ is even, the required inequality still holds. The other two inequalities required are

$$2K_1 + \delta + 1 \leq C$$

$$\lceil \delta/2 \rceil \leq K_2.$$

Here one must examine the conditions on admissible parameters in detail. There are three cases, the first of which has already been excluded. We give the conditions which separate these three cases along with some of the relevant side conditions which apply in each case.

$$\text{I . } K_1 = \infty;$$

$$\text{II . } K_1 < \infty, C = 2K_1 + 2K_2 + 1 \leq 2\delta + K_1; \text{ and } K_1 + K_2 \geq \delta;$$

$$\text{III . } K_1 \leq \infty, C > 2\delta + K_1; \text{ and } \delta \leq (3/2)K_2.$$

In case (II), we have

$$C = 2K_1 + 2K_2 + 1 \geq 2K_1 + (K_1 + K_2) + 1 \geq 2K_1 + \delta + 1$$

$$\delta \leq K_1 + K_2 \leq 2K_2$$

so the relevant inequalities hold in this case.

In case (III), we have

$$C > 2\delta + K_1 \geq \delta + 2K_1 + 1$$

$$\delta/2 \leq (3/4)K_2 \leq K_2$$

and again the relevant inequalities hold.

This disposes of the purely numerical constraints. The final point to check is the following: if $M = \delta$, then \mathcal{S} does not contain a Henson constraint in which the distance δ occurs. By our choice of M , this would mean

$$K_1 = \delta < \infty.$$

The characterization of admissibility in such cases implies that no Henson constraint in \mathcal{S} involves the distance δ . More precisely, in case (II), 1-cliques are allowed as Henson constraints, and in case (III), no Henson constraints are allowed.

Thus all conditions are verified and the lemma follows. □

We may now apply the general theory to prove Theorem 4.

Proof of Theorem 4. Our hypotheses on the parameters are those necessary for the application of Lemma 4.3.5. So we have a value M for which the associated class of finite structures is closed under $+_M$. Therefore by Lemma 4.3.2, the operator $+_M$ provides a free decomposition operator for the class.

By Cameron's criterion Theorem 4.3.1, the associated algebra is polynomial.

The indecomposable elements are those which are connected after deleting edges with weight M . For any n , any configuration consisting of a point a and n neighboring points is connected with respect to weight 1 edges. So there are infinitely many indecomposable isomorphism types and therefore the polynomial algebra has infinitely many generators. □

4.3.2 The bipartite antipodal case

We now examine one of the cases not covered by Theorem 4.

Lemma 4.3.6. *Let Γ be the generic bipartite antipodal graph of diameter 3, with $G = \text{Aut}(\Gamma)$. Then the associated algebra \mathcal{A}^G is a polynomial algebra in three variables.*

Furthermore, the associated class \mathcal{C} has a free decomposition operator.

Proof. If A is a finite bipartite antipodal graph with parts A_1, A_2 , let

$$\alpha(A) = (k, m, n)$$

where $m = \min(|A_1|, |A_2|)$, $n = \max(|A_1|, |A_2|)$, and k is the number of antipodal pairs in A .

Claim 4.3.6.1. *The function α induces a bijection between the set of isomorphism types of bipartite finite antipodal graphs and the set S of triples (k, m, n) satisfying $k \leq m \leq n$.*

We define a map β from S to bipartite finite antipodal graphs by setting $\beta(k, m, n) = (V_1, V_2)$ with $|V_1| = m$, $|V_2| = n$, $d(a_i, b_i) = 3$ for a_i, b_i which are k elements of V_1, V_2 , respectively, and remaining distances 1 between V_1 and V_2 and 2 within V_1 or V_2 .

Then $\alpha \circ \beta$ is the identity on S . We claim that $\beta \circ \alpha$ is also the identity.

If $\alpha(A) = (k, m, n)$, then we may suppose that $|A_1| = m$, $|A_2| = n$. Since there are exactly k pairs (a_i, b_i) at distance 3, with $a_i \in A_1$ and $b_i \in B_2$, clearly $A \simeq B$.

The claim follows.

Now S is a semigroup under pointwise addition, and the elements $x = (0, 0, 1)$, $y = (0, 1, 1)$, $z = (1, 1, 1)$ are indecomposable. Any element (k, m, n) may be written uniquely as

$$k(1, 1, 1) + (m - k)(0, 1, 1) + (n - m)(0, 0, 1)$$

so the semigroup is freely generated by x, y, z .

We may transfer this semigroup structure to the age of Γ .

There is also a natural partial order \leq on S given by $(k_1, m_1, n_1) \leq (k_2, m_2, n_2)$ if $k_1 \leq k_2$, $m_1 \leq m_2$, and $m_1 + n_1 = m_2 + n_2$ (this last condition is unimportant but will hold in all cases of interest).

We transfer this partial order to the age of Γ as well. Then the final assumption in Cameron's criterion (Definition 4.3.2) is that if A is partitioned into induced substructures B_1, \dots, B_ℓ , with $\alpha(A) = (k, m, n)$ and $\alpha(B_i) = (k_i, m_i, n_i)$, then

$$\left(\sum k_i, \sum m_i, \sum n_i \right) \leq (k, m, n),$$

i.e.,

$$\sum k_i \leq k \qquad \sum m_i \leq m \qquad \sum m_i + \sum n_i = m + n.$$

Clearly $\sum k_i \leq k$ by counting, and m is the sum of terms m_i or n_i with $m_i \leq n_i$, so the second inequality also holds. The final equality holds since the B_i collectively partition A .

Thus Cameron's criterion applies, and the generators for \mathcal{A}^G as a polynomial algebra correspond to the indecomposable elements x, y, z of S . \square

Open problems

Problem 4.3.1. *Is the associated algebra polynomial also in the remaining cases of the known metrically homogeneous graphs of generic type, namely the general case of bipartite graphs and of antipodal graphs of odd diameter?*

In such cases there is no operator of the form $+_M$ under which the class is closed, but as we have seen in Lemma 4.3.6, there may be a suitable free decomposition operator of another kind.

Problem 4.3.2. *Let \mathcal{F} be a finite set of finite connected graphs and suppose that there is an \aleph_0 -categorical countable universal \mathcal{F} -free graph Γ . Is the associated algebra a polynomial algebra?*

To clarify, in this setting there is a *canonical* \aleph_0 -categorical countable universal \mathcal{F} -free graph (namely, the existentially complete one), and the question applies to this particular graph. Again Cameron's criterion suggests a natural approach to the problem.

References

- [AH06] M. Alfuraidan and J. Hall, *Smith's theorem and a characterization of the 6-cube as distance-transitive graph*, *Journal of Algebraic Combinatorics* **24** (2006), 195–207.
- [ACM16] D. Amato, G. Cherlin, and H. D. Macpherson, *Metrically homogeneous graphs of diameter 3*, Preprint, 2016.
- [ABH⁺17] A. Aranda, D. Bradley-Williams, J. Hubička, M. Karamanlis, M. Kompatscher, M. Konečný, and M. Pawliuk, *Ramsey expansions of metrically homogeneous graphs*, arXiv e-print, arXiv:1707.02612 [math.CO], 2017.
- [BB80] Eiichi Bannai and Etsuko Bannai, *How many P -polynomial structures can an association scheme have?*, *European Journal of Combinatorics* **1** (1980), 289–298.
- [Bod15] M. Bodirsky, *Ramsey classes: examples and constructions*, *Surveys in Combinatorics 2015*, Cambridge Univ. Press, 2015, pp. 1–48.
- [Cam76] P. J. Cameron, *Transitivity of permutation groups on unordered sets*, *Math. Z.* **148** (1976), 127–139.
- [Cam81] ———, *Orbits of permutations groups on unordered sets, II*, *J. London Math. Soc.* **23** (1981), 249–264.
- [Cam90] ———, *Oligomorphic Permutation Groups*, London Mathematical Society Lecture Notes **152** (1990), viii+160.
- [Cam97] ———, *The algebra of an age*, *Model Theory of Groups and Automorphism Groups* (Blaubeuren, 1995), London Math. Soc. Lecture Note Ser., vol. 244, Cambridge Univ. Press, Cambridge, 1997, pp. 126–133. MR 1689855
- [Cam98] ———, *A census of infinite distance transitive graphs*, *Discrete Mathematics* **192** (1998), 11–26.
- [CT17] P. J. Cameron and S. Tarzi, *On the automorphism group of the m -coloured random graph*, Preprint, 2007. Posted to arXiv as arXiv:1709.07831 [math.CO], 2017.
- [Che11] G. Cherlin, *Two problems on homogeneous structures, revisited*, *Model Theoretic Methods in Finite Combinatorics*, number 558 (M. Grohe and J. A. Makowsky, eds.), Contemporary Mathematics, American Mathematical Society, Providence, RI, 2011, pp. 319–415.
- [Che17] ———, *Homogeneous ordered graphs and metrically homogeneous graphs*, http://sites.math.rutgers.edu/~cherlin/Paper/_H2017Full.pdf, 2017, [Online; accessed 10-December-2018].
- [Ell60] R. Ellis, *Universal minimal sets*, *Proceedings of the American Mathematical Society* **11** (1960), no. 4, 540–543.
- [Eng61] E. Engeler, *Unendliche Formeln in der Modelltheorie*, *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik* **7** (1961), 154–160.

- [FT18] J. Falque and N. M. Thiéry, *The orbit algebra of a permutation group with polynomial profile is Cohen-Macaulay*, arXiv e-print, arXiv:1804.03489 [math.CO], 2018.
- [Fra54] R. Fraïssé, *Sur l'extension aux relations de quelques propriétés des ordres*, Ann. Ecole Normale Sup. **7** (1954), 361–388.
- [Gar80] A. Gardiner, *Redrawing distance-regular graphs*, Unpublished manuscript, 1980.
- [Hau14] F. Hausdorff, *Grundzüge der Mengenlehre*, Leipzig: Veit., 1914.
- [Hen71] C. Ward Henson, *A family of countable homogeneous graphs*, Pacific Journal of Mathematics **38** (1971), 69–83.
- [Her98] Bernhard Herwig, *Extending partial isomorphisms for the small index property of many ω -categorical structures*, Israel J. Math **107** (1998), 93–123.
- [HL99] Bernhard Herwig and Daniel Lascar, *Extending partial isomorphisms and the profinite topology on free groups*, Transactions of the American Mathematical Society **352**, (5) (1999), 1985–2021.
- [Hod97] W. Hodges, *A Shorter Model Theory*, Cambridge University Press, 1997.
- [HHLS93] 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), no. 2, 204–218.
- [HKK17] J. Hubička, M. Kompatscher, and M. Konečný, *Forbidden cycles in metrically homogeneous graphs*, arXiv e-print, arXiv:1808.05177 [math.CO], 2017.
- [HN16] J. Hubička and J. Nešetřil, *All those Ramsey classes (Ramsey classes with closures and forbidden homomorphisms)*, arXiv e-print, arXiv:1606.07979 [math.CO], 2016.
- [Kec13] A. S. Kechris, *Dynamics of non-archimedean Polish groups*, European Congress of Mathematics, Eur. Math. Soc., Zürich, 2013, pp. 375–397. MR 3469133
- [KR07] A. S. Kechris and C. Rosendal, *Turbulence, amalgamation and generic automorphisms of homogeneous structures*, Proc. London Math. Soc. **94** (2007), no. 2, 302–350.
- [KPT05] A.S. Kechris, V.G. Pestov, and S. Todorčević, *Fraïssé limits, Ramsey theory, and topological dynamics of automorphism groups*, Geometric and Functional Analysis **15** (2005), no. 1, 106–189.
- [LW80] A. Lachlan and R. Woodrow, *Countable ultrahomogeneous undirected graphs*, Transactions of the American Mathematical Society **262** (1980), 51–94.
- [Lup17] M. Lupini, *Review: Two papers on metrizable universal minimal flows*, The Bulletin of Symbolic Logic **2016** (2017), no. 5, 201–203.
- [Mac82] H.D. Macpherson, *Infinite distance transitive graphs of finite valency*, Combinatorica **2** (1982), 63–69.
- [Mac85] ———, *Growth rates in infinite graphs and permutation groups*, Proceedings of the London Mathematical Society **3(2)** (1985), 285–294.
- [MTT16] J. Melleray, L. Nguyen Van Thé, and T. Tsankov, *Polish groups with metrizable universal minimal flows*, International Mathematics Research Notices **2016** (2016), no. 5, 1285–1307.
- [Mos92] L. Moss, *Distanced graphs*, Discrete Mathematics **102** (1992), 287–305.

- [Neš07] J. Nešetřil, *Metric spaces are Ramsey*, European Journal of Combinatorics **28**(1) (2007), 457–468.
- [NR77] J. Nešetřil and V. Rödl, *A structural generalization of the Ramsey theorem*, Bulletin of the Amer. Math. Soc. **83** (1977), no. 1, 127–128.
- [Pou08] M. Pouzet, *When is the orbit algebra of a group an integral domain? Proof of a conjecture of P. J. Cameron*, Informatique Théorique et Applications (Theor. Inform. Appl.) **42** (2008), 83–103.
- [Rot64] G.-C. Rota, *On the foundations of combinatorial theory I. Theory of Möbius functions*, Z. Wahrscheinlichkeitstheorie und Verw. Gebiete **2** (1964), 340–368.
- [RN59] C. Ryll-Nardzewski, *On the categoricity in power $\leq \aleph_0$* , Bulletin of the Polish Academy of Sciences. Mathematics. Série Sciences Mathématiques, Astronomiques, et Physiques **7** (1959), 545–548.
- [Ser74] J.-P. Serre, *Amalgames et points fixes*, Proceedings of the Second International Conference on the Theory of Groups. Lecture Notes in Mathematics **372** (1974), 633–640.
- [Ser77] ———, *Arbres, amalgames, $S L_2$* , Société Mathématique de France, Astérisque **46** (1977), 1–189.
- [She74] J. Sheehan, *Smoothly embeddable subgraphs*, J. London Math. Soc. **9** (1974), 212–218.
- [Sin17] D. Siniora, *Automorphism groups of homogeneous structures*, Ph.D. thesis, University of Leeds, March 2017.
- [Smi71] D.H. Smith, *Primitive and imprimitive graphs*, Q. J. Math. Oxford **2** (1971), no. 22, 551–557.
- [Sve59] L. Svenonius, *\aleph_0 -categoricity in first-order predicate calculus*, Theoria (Lund) **25** (1959), 82–94.
- [TZ12] K. Tent and M. Ziegler, *On the isometry group of the Urysohn space*, J. London Math. Soc. **87** (2012), no. 1, 289–303.
- [Ury27] P. Urysohn, *Sur un espace métrique universel*, Bull. Sci. Math. **51** (1927), 43–64,74–90.
- [Vee77] W. Veech, *Topological dynamics*, Bull. Amer. Math. Soc. **83** (1977), no. 5, 775–830.
- [YMT17] I. Ben Yaakov, J. Melleray, and T. Tsankov, *Metrizable universal minimal flows of Polish groups have a comeagre orbit*, Geometric and Functional Analysis **27** (2017), no. 1, 67–77.