

UCLA

UCLA Electronic Theses and Dissertations

Title

Descriptive Aspects of Locally Ckecable LAbelling and Constraint Satisfaction Problems

Permalink

<https://escholarship.org/uc/item/3xh1z3qh>

Author

Thornton, Riley

Publication Date

2022

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Descriptive Aspects of Locally Checkable Labelling
and Constraint Satisfaction Problems

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Mathematics

by

Riley Thornton

2022

© Copyright by
Riley Thornton
2022

ABSTRACT OF THE DISSERTATION

Descriptive Aspects of Locally Checkable Labelling and Constraint Satisfaction Problems

by

Riley Thornton

Doctor of Philosophy in Mathematics

University of California, Los Angeles, 2022

Professor Andrew S. Marks, Chair

We explore the descriptive set theory of problems originating in theoretical computer science. We investigate a few special cases of locally checkable labelling problems using a variety of Borel and measurable techniques. Our result clarifies a small part of the connection between descriptive set theory and distributed computing. And, we adapt tools from the algebraic approach to constraint satisfaction problems to the Borel setting. In particular we draw a direct connection between NP-completeness and Σ_2^1 -completeness.

The dissertation of Riley Thornton is approved.

Itay Neeman

Artem Chernikov

Timothy Austin

Andrew S. Marks, Committee Chair

University of California, Los Angeles

2022

To my grandads

TABLE OF CONTENTS

1	Introduction	1
1.1	Locally checkable labelling problems	3
1.1.1	Distributed complexity	4
1.1.2	Lovász local lemma	6
1.1.3	Overview of LCL results	8
1.2	Constraint satisfaction problems	9
1.2.1	Polymorphisms and pp definitions	10
1.2.2	Overview of results	11
1.3	Notation	12
1.4	Basic Tools	18
1.4.1	Locally countable relations and the Luzin–Novikov theorem	19
1.4.2	Invariant measures	20
1.4.3	Augmenting flows	22
1.4.4	Shift graphs and Bernoulli graphings	24
1.4.5	Weak containment and continuous model theory	28
1.4.6	Borel counterexamples: Baire category and determinacy	33
1.4.7	Dichotomies, Gandy–Harrington forcing, and complexity	35
2	Locally Checkable Labelling Problems	40
2.1	Orientations	40
2.1.1	Measurable constructions	42
2.1.2	Examples and applications	48

2.1.3	Complexity	53
2.2	Cayley diagrams	57
2.2.1	Measurable constructions	60
2.2.2	Examples and applications	63
2.2.3	Complexity	76
3	Constraint Satisfaction Problems	79
3.1	Introduction	79
3.1.1	Results and conjectures	81
3.2	Background on CSPs	85
3.3	Simple constructions	95
3.4	Intractable CSPs	101
3.5	Essentially classical CSPs	106
3.6	Effectivizable CSPs	112
A	Codings and edge coloring	118
	References	127

LIST OF FIGURES

1.1	The relation between some distributed and descriptive complexity classes on regular trees	5
2.1	A GRR for $D_\infty = \langle r, s : rsrs = s^2 = e \rangle$ induced by $\{r, s, sr, r^2s\}$	60
2.2	Two copies of T_3 and the Diestel–Leader graph	68
2.3	$B_1(x)$ in G	70
3.1	An intractable directed graph	104
A.1	An inverter component and diagram representation	119
A.2	A variable setting component	120
A.3	An or gate	120
A.4	The reduction applied to $(\neg v_1 \vee v_2 \vee \neg v_3) \wedge (v_1 \vee \neg v_2 \vee v_3)$	121

ACKNOWLEDGMENTS

The contents of this thesis were influenced by more people than I could possibly remember. Here is an attempt. Thank you to my advisor Andrew Marks for a lot of good advice. Thank you to Alekos Kechris, Alex Kastner, Alex Mennen, Alex Tenenbaum, Asaf Shani, Clark Lyons, Forte Shinko, Jan Grebík, Jay Williams, Kevin Matthews, Konrad Wroebel, Omer Ben Neria, Shaun Allison, Tyler Arrant, and Zoltán Vidnyánszky for many illuminating conversations. Thank you to Alex Frederick, Bar Roytman, Ben Johnsrude, Blaine Talbut, Denali Molitor, Don Laackman, Erin George, Gyu Eun Lee, Ian Coley, Juniper Bahr, Madeline Barnicle, Stephanie Wang, Zach Norwood, and the organizers of UAW2865 for making UCLA a much nicer place to work. Thanks to my family for a great deal of support. And thanks to my friends—especially Darya, David, Grant, Hunter, Jonas, Justice, Pearl, Quentin, Tandr , Viora, and Yuval—for, among other things, helping me survive a pandemic.

VITA

- 2013-2016 Reactor Operator, Reed Research Reactor
- 2016 B.A. Mathematics, Reed College
- 2016-2021 Teaching Assistant, Mathematics Department, UCLA
- 2021-2022 Fellow, Mathematics Department, UCLA

PUBLICATIONS

Δ_1^1 effectivization in Borel Combinatorics (submitted May 2021)

Factor maps for automorphism groups via Cayley diagrams (submitted January 2021)

Orienting Borel Graphs, Proceedings of the AMS 150(4) (2022) 1779-1793

Cubes and Their Center, Acta Mathematica Hungarica 152(2) (2017) 291-313

The Homogeneous Spectrum of Milnor-Witt K-Theory, J. Algebra 459 (2016) 376-388

CHAPTER 1

Introduction

Descriptive combinatorics asks questions like: What graphs on \mathbb{R} have a Lebesgue measurable 2-coloring? What posets on \mathcal{N} have Borel maximal antichains? More generally, when does a combinatorial problem on a Polish space have a definable solution? The term definable here is vague, but its scope includes such notions as continuous, measurable, Borel, and ordinal definable. These kinds of questions turn out to have deep and shockingly diverse connections to many other areas of mathematics. For instance:

1. Equidecompositions as in the Banach–Tarski paradox and Tarski circle squaring problem are equivalent to matchings in certain auxiliary graphs on the underlying space
2. Understanding the structure of cardinalities in inner models like $L(\mathbb{R})$ involves delicate analysis of equivalence relations on \mathbb{R} .
3. In classical combinatorics, limiting objects often carry a measurable structure, and definable solutions in the limit give uniform information about the finite objects

This thesis focuses in large part on interactions between computer science and descriptive combinatorics specifically in the Borel and measurable settings. In particular, we look at two classes of problems—Locally Checkable Labelling problems (LCLs) and Constraint Satisfaction Problems (CSPs)—which were first isolated by computer scientists interested in various notions of complexity. We will see that the computational notions of complexity parallel descriptive set theoretic notions. And in fact, in some cases, direct implications can be drawn.

The formalism of Locally Checkable Labelling problems comes from distributed computing. An LCL is specified by some collection of local constraints, an instance of an LCL is a graph, and we want to find a labelling which meets these constraints in every neighbourhood of the graph. Computer scientists are often interested in how much communication is needed between vertices of a graph to solve all instances of a particular LCL, and this gives rise to various distributed complexity classes. One can similarly ask when all instances of a given LCL have a definable solution in the various descriptive set theoretic senses of definability, and this gives rise to descriptive set theoretic complexity classes. It turns out these descriptive set theoretic and distributed complexity classes fit into a common picture of locality, and methods have been fruitfully translated between the two fields. For a survey, see [BCG21b][GR21]. In the first half of this thesis, we investigate a handful of LCLs related to the problem of generating a graph using functions. Our results are purely descriptive set theoretic, but since their original publication many of them have been seen to fit nicely into this picture of locality. We give more context in Section 1.1.

The archetypal CSP is 3SAT, i.e. the satisfaction problem for ternary disjuncts of literals. A CSP is described by a domain D and a collection of relations on D (that is, a CSP is described by a finite relational structure \mathcal{D}). An instance of $\text{CSP}(\mathcal{D})$ is a collection of variables and specification of relations from \mathcal{D} we want to hold among them (that is, an instance is a relational structure in the same signature). And, a solution to an instance \mathcal{X} of $\text{CSP}(\mathcal{D})$ is an assignment of values from \mathcal{D} to the variables in \mathcal{X} which satisfies the constraints imposed by \mathcal{X} (that is, a homomorphism from \mathcal{X} to \mathcal{D}). Details are given in Section 3.1. CSPs are a very rich class of problems, but still a very structured class. Indeed, many algorithmic complexity questions about CSPs, such as which CSPs are NP-complete or solvable by linear relaxation, are decidable using methods from algebra. The second half of this thesis adapts these algebraic tools to the study of Borel CSPs and draws direct comparisons between algorithmic and descriptive set theoretic notions of complexity. For instance, we show that any NP-complete CSP has a Σ_2^1 -complete corresponding Borel CSP.

We sketch the relevant algebra in Section 1.2 and give details in Chapter 3.

1.1 Locally checkable labelling problems

We parameterize LCLs by sets of finite labelled graphs. To this end, let $\mathcal{L}(G, A)$ be the set of labellings of (vertices, edges, etc. of) the graph G by elements of A , and let \mathbb{G}_\bullet be the set of (isomorphism classes of) finite rooted graphs.

Definition 1.1.1. An **LCL** is a set of labellings in $\bigcup_{G \in \mathbb{G}_\bullet} \mathcal{L}(G, A)$ for some fixed A with all elements of some fixed radius n . We call n the **radius of definition** of the LCL.

If φ is an LCL, and x is a vertex of a graph H , we say a labelling $f : H \rightarrow A$ **meets φ at x** if there is a graph isomorphism $g : K \rightarrow B_n(x)$ such that $(f \upharpoonright B_n(x)) \circ g \in \varphi$ (where we root $B_n(x)$ at x).

A labelling of a graph H is a **φ -decoration** or **solution** to the LCL φ on H if it meets φ at every vertex.

Of course we are interested in Borel and p.m.p. graphs which admit Borel labelling which solve an LCL (possibly ignoring a measure 0 or measure ϵ set of neighbourhoods). These kind of definable LCLs have shown up implicitly in dynamics, probability, and geometry. For instance:

Proposition 1.1.2. For any shift of finite type $X \subseteq A^\Gamma$, there is an LCL φ so that $a : \Gamma \curvearrowright Y$ admits a (Borel, measurable, continuous) homomorphism into X if and only if $\text{Sch}(a, E)$ admits a (Borel, measurable, continuous) φ -decoration.

Proof. Recall that a shift of finite type is a Γ -invariant subset $X \subseteq A^\Gamma$ for some finite A defined by a finite list of configurations c_1, \dots, c_n with finite windows $W_1, \dots, W_n \subseteq \Gamma$ so that $x \in X$ if and only if $(\forall i, \gamma) (\gamma \cdot x) \upharpoonright W_i \neq c_i$.

Let k be such that $W_1, \dots, W_n \subseteq B_k(e)$, and let φ be the set of configurations on $B_k(e)$

which avoid the c_i s. Then, if $f : Y \rightarrow X$ is a homomorphism, $g(x) = f(x)(e)$ is a φ -decoration of $\text{Sch}(a, E)$. And similarly, if g is a φ -decoration, then $f(x)(\gamma) = g(\gamma \cdot x)$ is a homomorphism. (Recall here that the Cayley graph is defined by the action of Γ on itself on the left, so Γ acts on labellings f by $(\gamma \cdot f)(\delta) = f(\gamma^{-1} \cdot_r \delta) = f(\delta\gamma)$.) \square

In particular, for any marked group (Γ, E) , $\text{Cay}(\Gamma, E)$ admits an Γ -FIID φ -decoration if and only if $\mathcal{S}(\Gamma, E)$ admits measurable φ -decoration. See also Subsection 1.4.4.

1.1.1 Distributed complexity

LCLs were first isolated in distributed computing. Computer scientists measure how much communication is needed solve an LCL on graphs of a given size in various models. In recent years, surprising formal connections have been drawn between the LOCAL model of computation and descriptive set theory. These connections provide some important context for the results in this thesis. Details of the model are not so important for us, but can be found in [BCG21b]

Definition 1.1.3. $\text{RLOCAL}(f)$ is the set of LCLs solvable in $f(n)$ rounds of communication in the randomized LOCAL model on graphs of size n . $\text{LOCAL}(f)$ is the set of LCLs solvable in $f(n)$ rounds in the deterministic model.

For LCLs on regular trees¹, the connections between LOCAL algorithms and descriptive set theory are summarized in the picture and theorem below.

Theorem 1.1.4. For an LCL φ , the following hold when we restrict to regular acyclic graphs of some fixed degree n :

- If $\varphi \in \text{LOCAL}(O(\log^*(n))) = \text{LOCAL}(o(\log(n))) = \text{RLOCAL}(o(\log \log(n)))$, then φ is solvable on every Borel graph

¹When we speak of finite d -regular trees we mean that only leaves are allowed to be of degree other than d , and we allow constraints to be violated in any neighbourhood containing a leaf.

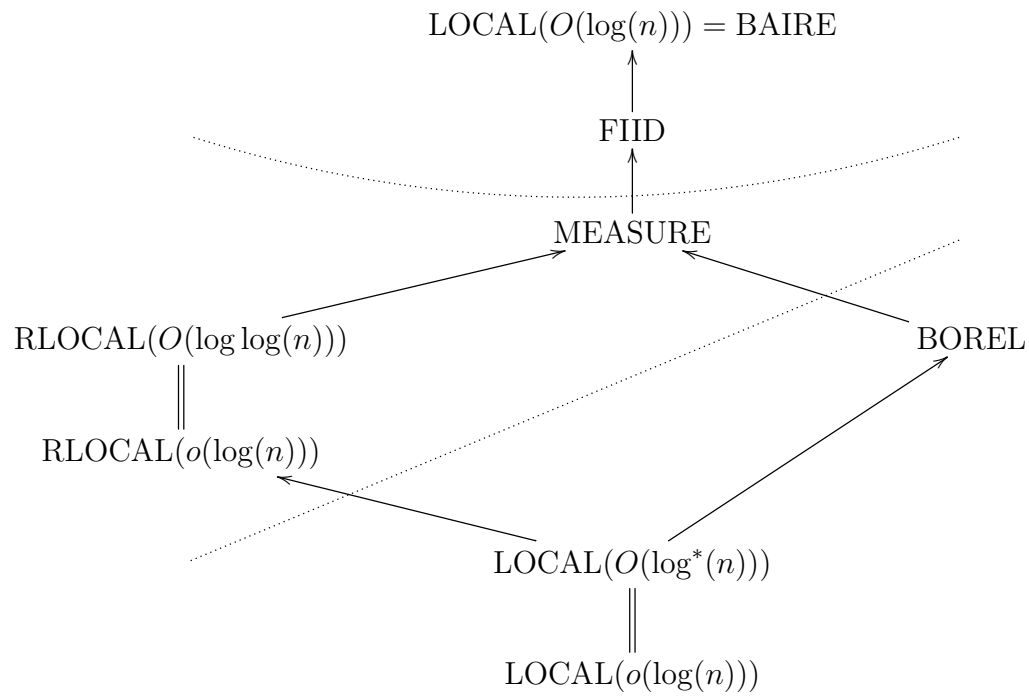


Figure 1.1: The relation between some distributed and descriptive complexity classes on regular trees

- If $\varphi \in \text{RLOCAL}(o(\log(n))) = \text{RLOCAL}(O(\log \log(n)))$, then every p.m.p. graph has a measurable φ -decoration
- $\varphi \in \text{LOCAL}(O(\log(n)))$ if and only if any Borel graph admits a Baire measurable φ -decoration
- If T_n has an $\text{Aut}(T_n)$ -FIID φ -decoration then every Borel graph has a Baire measurable φ -decoration

There are similar pictures for other classes of graphs: d -dimensional grids, bounded degree graphs, etc. We focus on trees here for simplicity.

1.1.2 Lovász local lemma

We can also describe some of these local complexity classes in terms of bounds for the Lovász local lemma. The Lovász local lemma is a powerful tool from probabilistic combinatorics which roughly says that if we have a collection of events which are not too dependent on each other and each occur with high probability, then they can all occur simultaneously. Moser and Tardos gave a simple randomized algorithm for solving instances of the Lovász local lemma, and this and other algorithms have been adapted to distributed computing and descriptive combinatorics [MT10][Ber19]. To state things precisely we need some definitions.

Definition 1.1.5. An **instance of the local lemma** is a pair (G, φ) where G is a graph (say on X) and φ is an LCL.

Let n be the radius of definition of φ and suppose φ is an LCL containing k -labellings. We define

$$d(G, \varphi) := \max_{x \in X} |B_n(x)| - 1$$

and

$$p(G, \varphi) := \sup_{x \in X} \mathbb{P}(f \text{ meets } \varphi \text{ at } x)$$

where this probability is taken uniformly over k -labellings of the finite graph $B_n(x)$.

For instance, if φ is the LCL defining edge d -coloring and G is d -regular, then $d(G, \varphi) = d$, and $p(G, \varphi) = 1 - (d!/d^d)$. The local lemma can be defined more generally, but for our discussion it suffices to look at LCLs.

It turns out that problems in $\text{RLOCAL}(o(\log(n)))$ are exactly the problems which can be turned into instance of the local lemma satisfying a certain polynomial bound.

Theorem 1.1.6 ([CP17]). If $\varphi \in \text{RLOCAL}(o(\log(n)))$, then for any constant c , φ is equivalent to an instance of the local lemma with $pd^c = O(1)$.

Theorem 1.1.7. [BCG21b] On trees, there is an $\text{RLOCAL}(O(\log \log(n)))$ algorithm for solving instances of the local lemma with $pd^8 \leq c$ for some small enough c .

Further, with a strong enough exponential bound we can find solutions in $O(\log^*(n))$ deterministic time (and thus we have Borel solutions).

Theorem 1.1.8 ([BGR20]). If φ is an LCL and for any G $p(\varphi, G)2^{d(\varphi, G)} < 1$, then $\varphi \in \text{LOCAL}(O(\log^*(n)))$.

And linear bounds as in the original statement of the Lovász local lemma give FIID solutions (at least on trees). For instance:

Theorem 1.1.9 ([Ber19]). If Γ is nonamenable and $(\varphi, \text{Cay}(\Gamma, E))$ is an instance of the local lemma with $epd < 1$, then $\mathcal{S}(\Gamma, E)$ has a measurable φ -decoration.

So in Figure 1.1.1, the region marked by dotted lines correspond to bounds on the local lemma. In the top region a linear bound is sufficient, in the middle region polynomial bounds suffice, and in the bottom region an exponential bound is needed. There are many open questions about optimality of these bounds and about the picture for other graph classes.

1.1.3 Overview of LCL results

In this thesis we consider a handful of LCLs, many related to orientation problems. For instance:

Definition 1.1.10. BO is the problem of finding a balanced orientation of a given graph (i.e. an orientation with indegree and outdegree equal at all vertices), and SO is the problem of finding a sinkless orientation (or equivalently a sourceless orientation) of a given graph (i.e. an orientation with outdegree at least 1 at all vertices).

Note that if G is d -regular, then $p(\text{SO}, G) = 1/2^d$ and $d(\text{SO}, G) = d$. So, we have $p(\text{SO}, G)2^{d(\text{SO}, G)} = 1$. In particular, $\text{SO} \in \text{RLOCAL}(o(\log(n)))$. We prove the following:

Theorem 1.1.11. The regular tree T_{2n} has an $\text{Aut}(G)$ -FIID balanced orientation, but for all d there is a d -regular p.m.p. acyclic graph with no measurable balanced orientation. So $\text{BO} \in \text{FIID} \setminus \text{MEASURE}$.

Theorem 1.1.12. For any d there is a d -regular acyclic Borel graph with no sinkless orientation.

This means that $p2^d < 1$ is a sharp bound for instances of the local lemma in the bottom region Figure 1.1.1. See Theorem 2.1.15 and Proposition 2.1.20.

In fact, we give optimal bounds on the minimum outdegree needed to orient many p.m.p. graphs. Our bounds mimic a classical bound given by the Edmonds matroid covering theorem. The method of proof suggests connections to another model of distributed computing, the so-called CONGEST model. See Theorem 2.1.7 and the following discussion.

The other class of LCLs we study are Cayley diagrams. A Cayley diagram is an assignment of generators of a marked group (Γ, E) to a Cayley graph of Γ which encodes an action of Γ . For instance, a Cayley diagram from C_2^{*n} with the usual generators is an edge n -coloring of the n -regular tree. We show that if a Cayley graph G admits an (approximate)

$\text{Aut}(G)$ -FIID Cayley diagram, then the distinction between (approximate) $\text{Aut}(G)$ - and Γ -FIID solutions to LCLs vanishes. We also give a number of examples of when these can and cannot be found, resolving a handful of questions from measurable graph combinatorics along the way.

For instance, we lift sharp asymptotic bounds for the density of $\text{Aut}(T_n)$ -FIID independent sets to bounds $\text{Aut}(\Gamma)$ -FIID independent sets when $\text{Cay}(\Gamma, E)$ is a tree. In the language of distributed computing, we find bounds on randomized algorithms for independent sets on trees already equipped with edge colorings or Schreier decorations. This answers a long-standing question of Marks and Keechris. See Thm 2.2.21

We also compute the projective complexity of Borel versions of these and related problems. For instance, we show that the set of Borel equivalence relations which admit a Borel selector is Σ_2^1 -complete in the codes (Theorem 2.1.25). This is in sharp contrast to the complexity of smooth relations.

1.2 Constraint satisfaction problems

For a finite relational structure \mathcal{D} , $\text{CSP}(\mathcal{D})$ is the problem of testing if a given structure in the same signature admits a homomorphism into \mathcal{D} .

Definition 1.2.1. For relational structures $\mathcal{D} = (D, R_1^{\mathcal{D}}, \dots, R_n^{\mathcal{D}})$ and $\mathcal{E} = (E, R_1^{\mathcal{E}}, \dots, R_n^{\mathcal{E}})$, a **homomorphism** $f : \mathcal{D} \rightarrow \mathcal{E}$ is a function $f : D \rightarrow E$ so that $(x_1, \dots, x_n) \in R_i^{\mathcal{D}}$ implies $(f(x_1), \dots, f(x_n)) \in R_i^{\mathcal{E}}$ for all i .

The class of CSPs is broad and contains many well-studied problems. For instance, $\text{CSP}(K_n)$ is the problem of n -coloring graphs. And, it is straightforward to code k SAT and systems of linear equations over \mathbb{F}_n as CSPs for finite relational structures. However, there is a rich algebraic structure to CSPs that allows us to answer many algorithmic questions about these problems.

1.2.1 Polymorphisms and pp definitions

For a finite relational structure \mathcal{D} , the polymorphisms of \mathcal{D} are the operations which preserve solutions to $\text{CSP}(\mathcal{D})$. More formally:

Definition 1.2.2. A **polymorphism** of \mathcal{D} is a homomorphism $f : \mathcal{D}^n \rightarrow \mathcal{D}$ for some n , and $\text{Pol}(\mathcal{D})$ is the set of polymorphisms of \mathcal{D} .

Here, we use the so-called categorical product for structures. For example, one can check that $\text{Pol}(K_n)$ is generated by projections and permutations for $n > 2$. For $n = 2$ we have some extra structure; K_2 is preserved by the following operations:

$$(x, y, z) \mapsto x + y + z \pmod{2}$$

$$\text{maj}(x, y, z) = \text{the repeated value among } x, y, z.$$

Note that $\text{CSP}(K_n)$ is polynomial time solvable if and only if $n = 2$. This is not a coincidence; the presence of nontrivial polymorphisms is, in some sense, responsible for the tractability of $\text{CSP}(K_2)$. To explain the general phenomenon we need some definitions

Definition 1.2.3. Say that a relation R is **pp definable** in a structure \mathcal{D} if

$$R(x_1, \dots, x_n) \Leftrightarrow (\exists y_1, \dots, y_m) \bigwedge_i \alpha_i(x_1, \dots, x_n, y_1, \dots, y_m)$$

where each α_i is either a relation in \mathcal{D} or an equality.

If every relation \mathcal{E} is a pp definable in \mathcal{D} , then any instance of \mathcal{E} can be turned into an instance of $\mathcal{D} \cup \{(\text{=})\}$ by adding in dummy variables. And in the finitary setting, we can eliminate (=) by taking a quotient. So, any algorithm for \mathcal{D} gives an algorithm for \mathcal{E} .

Theorem 1.2.4 ([Jea98]). $\text{Pol}(\mathcal{D}) \subseteq \text{Pol}(\mathcal{E})$ if and only if \mathcal{D} pp defines \mathcal{E}

Corollary 1.2.5. If $\text{Pol}(\mathcal{D}) \subseteq \text{Pol}(\mathcal{E})$, then $\text{CSP}(\mathcal{E})$ is polynomial time reducible to $\text{CSP}(\mathcal{D})$.

Abstractly, this means that any algorithmic question has an answer in terms of $\text{Pol}(\mathcal{D})$. (This can be generalized quite a bit, see Theorem 3.2.6). The most spectacular example of this is the so-called CSP dichotomy theorem:

Theorem 1.2.6 ([Bul17][Zhu17]). If there is some $f \in \text{Pol}(\mathcal{D})$ so that

$$(\forall a, e, r) f(a, r, e, a) = f(r, a, r, e) \quad (\star)$$

then $\text{CSP}(\mathcal{D}) \in \text{P}$. Otherwise, $\text{CSP}(\mathcal{D})$ is NP-complete.

1.2.2 Overview of results

The theory of polymorphisms and pp definitions does not immediately apply to the descriptive set theoretic setting. In particular, the reductions given by a pp definition require taking a quotient, and quotients by Borel equivalence relations can be pathological (consider, for instance, \mathbb{R}/\mathbb{Q}). Nonetheless, with some work this algebraic approach can be adapted.

Definition 1.2.7. Say that \mathcal{D} is **essentially classical** if, whenever a Borel instance of $\text{CSP}(\mathcal{D})$ has a solution, it has a Borel solution.

We can completely classify essentially classical structures. They turn out to be exactly the so-called width 1 structures. In terms of polymorphisms:

Theorem 1.2.8 (see Theorem 3.1.11). \mathcal{D} is essentially classical if and only if \mathcal{D} has a polymorphism f of high enough arity so that $f(x_1, \dots, x_n) = f(y_1, \dots, y_n)$ whenever $\{x_1, \dots, x_n\} = \{y_1, \dots, y_n\}$.

We can also adapt half of the CSP dichotomy theorem to descriptive setting. Projective complexity is the relevant measure of complexity for Borel problem (see Subsection 1.4.7), and a CSP is at most Σ_2^1 -complete.

Theorem 1.2.9 (see Theorem 3.4.3). If \mathcal{D} does not have a polymorphism satisfying (\star) above, then the Borel version, $\text{CSP}_B(\mathcal{D})$, is Σ_2^1 -complete.

Corollary 1.2.10 ($P \neq NP$). If $\text{CSP}(\mathcal{D})$ is NP-complete, then $\text{CSP}_B(\mathcal{D})$ is Σ_2^1 -complete.

And we have partial converses, such as a descriptive set theoretic analogue of the Hell-Nešetřil theorem:

Theorem 1.2.11. For an undirected graph G , $\text{CSP}_B(G)$ is Σ_2^1 -complete if and only if G is not bipartite.

1.3 Notation

We will record some notation in this section. Most of this standard, and we will try to recall the sticking points throughout the text. We will give the formalism for LCLs here as they will be relevant in the next section.

We will be concerned with graphs throughout. For us, a graph on a vertex set X is a symmetric irreflexive subset of X^2 . An edge in a graph $G \subseteq X^2$ is pair $(x, y) \in G$. We write $\{x, y\} \in G$ to mean $(x, y) \in G$ and $(y, x) \in G$ when we want to emphasize that G is symmetric. When we say that a graph G on a Polish space X is Borel, we mean that G is Borel in the product spaces X^2 . Given a graph G , E_G is the connectedness relation of G :

$$xE_Gy \Leftrightarrow (\exists x = z_1, z_2, \dots, z_{n+1} = y)(\forall i = 1, \dots, n) (z_i, z_{i+1}) \in G$$

For a relation $R \subseteq X \times Y$, and $x \in X$, then section over x is $R_x = \{y : (x, y) \in R\}$. Similarly, the section under $y \in Y$ is $R^y = \{x : (x, y) \in R\}$. Given a graph G on X , the degree of G is $\text{deg}(x) := |G_x|$. We will often assume this is uniformly bounded for $x \in X$. We call such graphs bounded degree graphs. We write $B_n(x, G)$ for the ball of radius n in G centered on x . When we refer to $B_n(x, G)$ as a rooted graph, we mean to choose x as the root. Often we suppress G in the notation.

Given a pair $e = (x, y) \in X^2$, we write $-e = (y, x)$, and $-A = \{-e : e \in A\}$. A directed graph is an irreflexive relation, and a directed graph D is an orientation of a graph G if

$G = D \cup -D$ and $D \cap -D = \emptyset$. For x a vertex of a directed graph D , we write

$$B_n^+(x, D) := \{y : (\exists x = x_0, \dots, x_n = y)(\forall i < n) x_i = x_{i+1} \text{ or } (x_i, x_{i+1}) \in D\}.$$

Similarly, $B_n^-(x, D) = B_n^+(x, -D)$. We will often suppress D in the notation. We say (x_0, \dots, x_n) is a directed path of length n in D if $(x_i, x_{i+1}) \in D$ for $i < n$. So, $B_n^+(x)$ is the set of vertices y with a directed path of length at most n from x to y .

Many of the graphs we study will come from group actions. Throughout, Γ will represent a countable discrete group and e will be the identity element of Γ . By a generating set for Γ , we mean a set $E \subseteq \Gamma$ so that $e \notin E$ and $\Gamma = \bigcup_{n \in \mathbb{N}} (E \cup E^{-1})^n$. Note that we disallow e (to avoid loops in our graphs) and we allow ourselves to take inverses. By a marked group, we mean a group Γ equipped with a finite generating set E .

Given an action $a : \Gamma \curvearrowright X$ of a marked group (Γ, E) , we can form the associated Schreier graph on X :

$$\text{Sch}(a, E) = \{(x, \gamma \cdot x) : x \in X, \gamma \in E \cup E^{-1}\}.$$

Γ acts on itself by left multiplication.

$$\lambda : \Gamma \curvearrowright \Gamma$$

$$\gamma \cdot_\lambda x = \gamma x$$

(The choice of left or right multiplication is arbitrary, but will cascade throughout this text.)

The Cayley graph of (Γ, E) is the Schreier graph of this action:

$$\text{Cay}(\Gamma, E) = \text{Sch}(\lambda, E)$$

Note that Γ embeds into $\text{Aut}(\text{Cay}(\Gamma, E))$ by *right* multiplication. That is, writing $\langle \gamma \rangle$ for the image of γ in $\text{Aut}(\text{Cay}(\Gamma, E))$, $\langle \gamma \rangle(x) = x\gamma^{-1}$.

For A a Polish space and G a graph with vertex set X , an A -labelling of H is a partial function

$$\ell : \bigsqcup_n \{(x_1, \dots, x_n) : (\forall i, j \leq n) x_i E_G x_j\} \rightarrow A.$$

We will mostly be concerned with vertex and edge labellings, i.e. labelling $f : X \rightarrow A$ and $g : G \rightarrow A$. We denote by $\mathcal{L}(H, A)$ the space of A -labellings of H . If H is countable, this comes with a natural Polish product topology.

The group $\text{Aut}(H)$ acts on $\mathcal{L}(H, A)$ by shifting indices. For instance, if ℓ is a vertex labelling we have:

$$f \cdot \ell(x) = \ell(f^{-1}(x)),$$

and if ℓ is an edge labelling:

$$f \cdot \ell(x, y) = \ell(f^{-1}(x), f^{-1}(y)).$$

So, Γ acts on labellings as follows:

$$\gamma \cdot \ell(x) = \langle \gamma \rangle \cdot \ell(x) = \ell(\langle \gamma \rangle^{-1}(x)) = \ell(x\gamma)$$

(Again the choice of which side γ appears on is arbitrary, but we need to be sure the correct number of inverses show up.)

Often we want to restrict to the free part of an action: for $a : \Gamma \curvearrowright X$,

$$\text{Free}(X) = \{x : (\forall \gamma \neq e) \gamma \cdot x \neq x\}.$$

In the case that X is a space of labellings equipped with a shift action, there is some ambiguity in whether we mean free with respect to the full automorphism group or some countable subgroup. We typically mean the former when the label space is uncountable. For (Γ, E) a countable marked group, write $\text{Free}(\Gamma, E)$ for the free part of $[0, 1]^\Gamma$ (i.e. the space of total $[0, 1]$ -vertex labellings) under the action of $\text{Aut}(\text{Cay}(\Gamma, E))$. This contains, for instance, the injective labellings, so has measure 1 with respect to the product of Lebesgue measure.

Given a marked group (Γ, E) , the associated shift graph is the Schreier graph of the shift action on $\text{Free}(\Gamma, E)$. The choice of base space here turns out to be not so important, see Subsection 1.4.4. Given a vertex transitive graph G , we can similarly associate the so-called

Bernoulli graphing. (The difference between the shift graph and the Bernoulli graphing is subtle but important.) Pick a root e for G and let $\text{Aut}_e(G) = \{f \in \text{Aut}(G) : f(e) = e\}$. The vertex set is $\tilde{X} = \text{Free}([0, 1]^G) / \text{Aut}_e(G)$, the space of vertex labellings of G up to rooted isomorphism. This is a standard Borel space as $\text{Aut}_e(G)$ is compact. Note that \tilde{X} inherits the quotient measure from $\text{Free}([0, 1]^G)$. Let M be the set of automorphisms that move e to one of its neighbours. The edge set of the Bernoulli graphing is

$$\tilde{\mathcal{S}}(G) := \{(\text{Aut}_e(G) \cdot x, \text{Aut}_e(G) \cdot y) : (\exists f \in M) f(x) = y\}.$$

If $G = \text{Cay}(\Gamma, E)$, we can equivalently write

$$\tilde{\mathcal{S}}(G) = \{(\text{Aut}_e(G) \cdot x, \text{Aut}_e(G) \cdot \gamma x) : x \in \text{Free}(\Gamma, E), \gamma \in E\}.$$

We will see that the quotient measure is invariant under this $\tilde{\mathcal{S}}(G)$ and that $\tilde{\mathcal{S}}(G)$ captures the $\text{Aut}(G)$ -FIID labellings of G . (See 1.4.4.)

Note that if $\Gamma \leq \text{Aut}_e(G)$ acts transitively on G , then $\text{Aut}(G) = \Gamma \text{Aut}_e(G) = \text{Aut}_e(G)\Gamma$. That is, any automorphism factors as an element of Γ and $\text{Aut}_e(G)$ with the element of γ showing up on either side. For $G = \text{Cay}(\Gamma, E)$ such factorizations are unique and related by the following commutation relations:

$$r\langle\gamma\rangle = \langle r(\gamma^{-1}) \rangle^{-1} (\langle r(\gamma^{-1}) \rangle r\langle\gamma\rangle). \quad (1.1)$$

Let \mathbb{G}_\bullet be the set of (isomorphism types of) finite rooted graphs. A locally checkable labelling problem (or LCL) is described by a set of local constraints, i.e. a finite $\varphi \subseteq \bigcup_{G \in \mathbb{G}_\bullet} \mathcal{L}(G, A)$ for some A finite. We view φ as a set of allowed configurations. The radius of definition is the largest radius of the domain of any configuration in φ .

Definition 1.3.1. An **LCL** is a finite subset $\varphi \subseteq \bigcup_{G \in \mathbb{G}_\bullet} \mathcal{L}(G, A)$ so that the domain of each $f \in \varphi$ has the same radius n . We call n the **radius of definition** of φ .

If φ is an LCL, and x is a vertex of a graph H , we say a labelling $f : H \rightarrow A$ **meets** φ **at** x if there is a graph isomorphism $g : K \rightarrow B_n(x)$ such that $(f \upharpoonright B_n(x)) \circ g \in \varphi$ (where we root $B_n(x)$ at x).

A labelling is a φ -**decoration** or **solution** to the LCL φ if it meets φ at every vertex. If the vertex set of H carries a measure μ , a μ -**measurable** φ -**decoration** is a measurable labelling f which meets φ on a co-null set. A sequence of labellings $\langle f_i : i \in \omega \rangle$ is an **approximate** φ -**decoration** if $\mu(\{x : f \text{ does not meet } \varphi \text{ at } x\}) \rightarrow 0$.

So if G is a countable graph, a random φ -decoration of a graph G is a measure μ on $\mathcal{L}(G, A)$ so that $\mu(\{f : f \text{ meets } \varphi \text{ at } x\}) = 1$ for all $x \in G$. If further μ is invariant under some group $\Gamma \subseteq \text{Aut}(G)$ that acts transitively on G , then it suffices that

$$\mu(\{f : f \text{ meets } \varphi \text{ at } x\}) = 1$$

for any particular x . Approximate random decorations can be defined analogously.

Recall that, for $\Gamma \leq \text{Aut}(G)$, an Γ -invariant measure μ on X is Γ -FIID if there is an equivariant map $f : [0, 1]^G \rightarrow X$ so that the pushforward of the power of Lebesgue measure λ^G is μ . We typically drop the word random from the phrase “ Γ -FIID random (approximate) φ -decoration.”

We will also be concerned with the density with which we can meet an LCL:

Definition 1.3.2. For φ an LCL and H a graph on a probability space (X, μ) ,

$$\rho^\varphi(H) := \sup \{ \mu(\{x : f \text{ meets } \varphi \text{ at } x\}) : f \in \mathcal{L}(X, A) \text{ measurable} \}.$$

For instance, the independence number of H is $\rho^\varphi(H)$, where φ says that ℓ is a vertex $\{0, 1\}$ -labelling with $\ell^{-1}(1)$ independent and $\ell(x) = 1$.

We will also consider relational structures more generally. A relation on a set A of arity k is a subset of A^k . We will write $R(x_1, \dots, x_k)$ to mean $(x_1, \dots, x_k) \in R$. We will switch freely between these two notations. In particular a unary predicate is just a subset of A . Some relations, like equality, are typically written in infix notation. We will write these in parenthesis to emphasize that we mean the associated set. For instance, we have $(=) = \{(a, b) \in A^2 : a = b\}$.

By a relational structure we mean a tuple $\mathcal{D} = (D, R_1, R_2, \dots)$, where D is a set and each R_i is a relation on D . Relational structures will be denoted throughout by script letters: $\mathcal{D}, \mathcal{E}, \mathcal{X}$, etc. Unless otherwise stated, the domain of a structure will be denoted by the unscripted letter: D, E, X , etc.

We say that a structure is finite to mean that its domain is finite and that it comes equipped with only finitely many non-empty relations. When \mathcal{D} is a finite relational structure, we say that \mathcal{X} is an instance of \mathcal{D} when it has the same signature. That is, if $\mathcal{D} = (D, R_1, \dots, R_n)$, then $\mathcal{X} = (X, S_1, \dots, S_n)$ where S_i has the same arity as R_i . We refer to S_i as the interpretation of R_i in \mathcal{X} and write $S_i = R_i^{\mathcal{X}}$. The superscript may be dropped if the context is clear. If \mathcal{X} is an instance of \mathcal{D} , we say \mathcal{D} is a template of \mathcal{X} .

We will want to perform a few operations on structures. For $\mathcal{D} = (D, R_1, \dots, R_n)$ a structure and R a relation on D , we abuse notation and write $\mathcal{D} \cup \{R\}$ for (D, R_1, \dots, R_n, R) . For $A \subseteq D$ we write $\mathcal{D} \upharpoonright A$ for the structure with domain A equipped with relations $R_i \cap A^k$ for each i (where R_i has arity k .) And, if R and S are relations both of arity k on A and B respectively, then $R \sqcup S$ is the relation on $A \sqcup B$ defined by

$$R \sqcup S = \{e \in A^k \sqcup B^k : R(e) \text{ or } S(e)\}$$

A homomorphism between relational structures in the same signature is a map which preserves relations. That is, $f : \mathcal{D} \rightarrow \mathcal{E}$ is a homomorphism if, for all R

$$(x_1, \dots, x_n) \in R^{\mathcal{D}} \Rightarrow (f(x_1), \dots, f(x_n)) \in R^{\mathcal{E}}.$$

For products, we take the categorical product (i.e. the product in the category of structures with homomorphisms). In detail, given \mathcal{D}, \mathcal{E} and a relation R

$$R^{\mathcal{D} \times \mathcal{E}}(x, y) :\Leftrightarrow R^{\mathcal{D}}(x) \text{ and } R^{\mathcal{E}}(y).$$

For a set A , $\pi_i^k : A^k \rightarrow A$ is the i^{th} projection map

$$\pi_i^k(a_1, \dots, a_k) = a_i.$$

And for $J = \{j_1, \dots, j_n\} \subseteq \{1, \dots, k\}$, π_J^k is the projection onto the coordinates in J :

$$\pi_J^k(a_1, \dots, a_k) = (a_{j_1}, \dots, a_{j_n}).$$

We will drop the superscript when the context is clear.

Our infinite structures will all live on Polish spaces. When we say a structure on X is Borel, we mean each relation is Borel in the product space. By standard universality arguments, we can typically restrict our attention to the space $\mathcal{N} = \mathbb{N}^{\mathbb{N}}$. Most reasonable codings of Borel sets will work for our arguments. See Appendix A for some details.

We will work throughout with Borel and projective complexity classes. See, e.g. [Mar19] for the basic definitions. Most relevantly, Σ_1^1 is the class of sets in Polish spaces defined with an existential quantifier over \mathbb{R} with Borel predicates, Π_1^1 is the class of sets defined with a universal quantifier, and Σ_2^1 is the class of $(\exists\forall)$ -definable sets. The lightface symbols Σ_1^1 , Π_1^1 , etc. denote effective analogues of these.

Many results in this paper refer to computational complexity classes. Almost anywhere P (or NP) appears it can be replaced with the algebraic notion of tractability (or intractability). See Definition 3.2.8 and the following comments. In particular most theorems in this paper are true even if $P=NP$. We will point out any exceptions by putting “ $(P \neq NP)$ ” before any theorems which require this assumption.

1.4 Basic Tools

Most of the tools needed in this thesis are explained in the survey on Borel graph theory by Pikhurko [Pik21], Marks’s notes on effective descriptive set theory [Mar19], and the survey of continuous model theory by Ben Yaacov, Berenstein, Ward Henson, and Usvyatsov [BHU06]. We will summarize some of the main points here and explain a give more details on a few of the topics not found in these surveys (namely effectivization, the Bernoulli graphing associated to a countable graph, and weak containment for graphs). The most notable

omission from this section is the machinery of toasts and asymptotic Borel dimension, which have resolved many questions about the descriptive combinatorics of nilpotent (and other elementary amenable) group actions. See for instance [GR21]. This section can safely be skipped and referred back to as needed.

1.4.1 Locally countable relations and the Luzin–Novikov theorem

In the next chapter of this thesis we will primarily be concerned with Borel relations whose sections are all countable (or even of bounded cardinality). It turns out that all such relations are generated by countably many Borel functions. In the case of equivalence relations, these functions can even be corralled into a group action. We will return to general of question of how many and how nice these functions can be later.

Definition 1.4.1. A relation $R \subseteq X \times Y$ is locally countable if $|R_x| \leq \aleph_0$ for all $x \in X$.

Theorem 1.4.2 ([Pik21, Theorem 3.8]). If $R \subset X \times Y$ is a locally countable Borel relation, then $R = \bigcup_i f_i$ where each $f_i : X \rightarrow Y$ is a Borel partial injection.

Theorem 1.4.3 ([Pik21, Theorem 7.1]). If $R \subseteq X^2$ is a Borel, locally countable, symmetric and total relation, then $R = \bigcup_i f_i$, where each f_i is a total Borel involution. In particular, if R is an equivalence relation, then it is the orbit equivalence relation of an action of $\bigoplus_{n \in \omega} C_2^{*n}$.

Together with the following proposition, these results make many complexity calculations much simpler for locally countable relations.

Proposition 1.4.4. For any Borel A and injective Borel function f , $f(A)$ is also Borel. And for any Borel partial function f , $\text{dom}(f)$ is Borel.

This means any set obtained by a countable length procedure operating on locally countable Borel inputs will be Borel. For instance:

Corollary 1.4.5. For any Borel set A and locally countable Borel graph G , the connected component of A , $[A]_G = \{y : (\exists x \in A) yE_G x\}$ is Borel.

Proof. Indeed, if $G = \bigcup f_i$ with each f_i a partial Borel injection, then $N(A) = \bigcup_i f_i(A)$ is Borel by the above proposition. So, $[A]_G = \bigcup_i N^i(A)$ is Borel. \square

1.4.2 Invariant measures

In the measurable setting we will often focus on probability measure preserving (or p.m.p.) graphs. In other sources these are called graphings. These are locally countable graphs on a measure space satisfying measure theoretic version of the handshake lemma which allows us to measure sets of edges:

Definition 1.4.6. A Borel graph G on a standard probability space (X, μ) is **probability measure preserving** or **p.m.p.** if, whenever $f \subseteq G$ is injective and $A \subseteq X$ is Borel

$$\mu(A) = \mu(f(A)).$$

Equivalently if, for $A \subseteq G$ Borel

$$\tilde{\mu}(A) := \int_{x \in X} |A_x| d\mu = \int_{x \in X} |A^x| d\mu$$

where $A_x = \{y : (x, y) \in A\}$ and $A^x = \{y : (y, x) \in A\}$.

We say that μ is an invariant measure for G

We will often abuse notation and write μ for $\tilde{\mu}$. The measure $\tilde{\mu}$ allows us to use edge counts to track modifications to labellings and ensure almost everywhere convergence of algorithms. The following argument Conley and Tamuz is a good illustration.

Theorem 1.4.7. If G is a bounded degree p.m.p. graph then there is a vertex labelling of G by $\{0, 1\}$ so that the label at each vertex differs from at least half of its neighbours.

Proof. A theorem of Kechris–Solecki–Todorcevic [Pik21, Theorem 5.12] says that every bounded degree Borel graph has a finite Borel coloring. So, we can find a sequence of sets of vertices, $\langle X_i : i \in \omega \rangle$, where each X_i is independent and each vertex appears infinitely often.

Let f_0 be any $\{0, 1\}$ labelling of the vertices and define f_i inductively by flipping every vertex in X_i with too many same-colored neighbours. That is,

$$f_{i+1}(x) = \begin{cases} 1 - f_i(x) & x \in X_i \text{ and for most neighbours } y \text{ of } x, f_i(x) = f_i(y) \\ f_i(x) & \text{otherwise} \end{cases}$$

Intuitively, flipping colors in this manner always decreases the fraction of edges between same-colored vertices, so the process must converge in the limit. To make this precise, consider $A_i = \{(x, y) \in G : f_i(x) = f_i(y)\}$ and $E_i = \{x : f_i(x) \neq f_{i+1}(x)\}$. Note that $\mu(A_{i+1}) \leq \mu(A_i) \leq \mu(G)$. Since

$$A_i \Delta A_{i+1} \subseteq \{(x, y) : x \text{ or } y \in E_i\}$$

and $|(A_i)_x| > |(A_{i+1})_x|$ for $x \in E_i$ we have $\mu(A_i) - \mu(A_{i+1}) \geq \mu(E_i)$. This gives us a telescoping sum:

$$\sum_i \mu(E_i) \leq \sum_i \mu(A_i) - \mu(A_{i+1}) = \mu(A_0) < \infty.$$

By the Borel–Cantelli lemma, $f_i(x)$ converges for almost all x , say to $f(x)$. If $f(x) = f(y)$ for most neighbours y of x , then there is some i where, for any $j > i$, $f(x) = f_j(x)$ and $f(y) = f_j(y)$ for all neighbours y of x . But then x would have flipped at some stage after i . □

We will sometimes also work with measures that are not quite invariant, but where measure does not vary too much along edges. These are so-called quasi-invariant measures.

Definition 1.4.8. A measure μ is G **quasi-invariant** if there is a measurable function $\rho : G \rightarrow \mathbb{R}^+$ so that, for $f \subseteq G$ injective and A Borel

$$\mu(f(A)) = \int_{x \in A} \rho(x, f(x)) d\mu.$$

Such a ρ is called a **Rado-Nikodym cocycle** for μ . We typically write $\rho = \text{ess sup}_G \rho(x, y)$.

One can check that, for instance, the argument for Theorem 1.4.7 works whenever ρ is sufficiently small. A discussion of quasi-invariance and equivalent conditions can be found in [KM04].

1.4.3 Augmenting flows

There are two important features of the above argument. First we have some greedy algorithm we want to run on a Borel graph. We cannot operate on vertices one by one, so we use some general coloring results to solve a scheduling problem and operate on entire independent sets of vertices at once. The second feature is that we show that the sets of points we need to modify at each step (E_i in the above argument) have summable measure, so our construction will work in the end by the Borel–Cantelli argument. Showing this summability usually relies on some invariance of measure. The next two lemmas package these ideas usefully for us. We will prove the lemma for vertex labellings, though the argument adapts easily to more general labellings.

Definition 1.4.9. For an LCL φ and labelling ℓ of G , $E(\varphi, \ell) = \{x : \ell \text{ does not meet } \varphi \text{ at } x\}$.

An r -**augmentation** of a labelling ℓ for an LCL φ is a connected subgraph G' of size at most r with a labelling ℓ' of G' so that changing ℓ to ℓ' on G' satisfies φ at more points. That is, for $f := \ell' \cup (\ell \setminus \ell \upharpoonright G')$, $E(\varphi, \ell) \supseteq E(\varphi, f)$.

Lemma 1.4.10. Let G be a locally finite Borel graph on a space X , ℓ be a Borel labelling. Then, for any $r > 1$, there is a Borel labelling ℓ' of G with no r -augmentations for φ and r Borel maps, $f_1, \dots, f_r : E(\varphi, \ell) \setminus E(\varphi, \ell') \rightarrow X$ such that

- $d(x, f_i(x)) \leq r$
- If $\ell(x) \neq \ell'(x)$, then $x \in \text{im}(f_i)$ for some i .

In the next lemma we use the Borel–Cantelli lemma to get measurable decorations. However, we only use measures to analyze the convergence of our decorations. The labellings we construct are Borel and independent of the measure.

Lemma 1.4.11. Suppose μ is a G quasi-invariant measure and that there are ϵ_r so that $\sum_r r \rho^r \epsilon_r < \infty$ and whenever a labelling ℓ of G has no r -augmentation then $\mu(E(\varphi, \ell)) < \epsilon_r$. Then, G has a μ -measurable φ decoration.

In fact, there is a single Borel labelling which meets φ off of a μ -null set for any measure μ which meets these hypothesis. This yields purely Borel results for graphs which support rich enough measures.

Proof. Define ℓ_i inductively as follows: let ℓ_0 be any labelling, and let ℓ_{i+1} be a labelling with no $(i+1)$ -augmentation as in the previous lemma. If μ is a quasi-invariant measure as in the statement, then

$$\mu(\{x : \ell_i(x) \neq \ell_{i+1}(x)\}) \leq r\rho^r E(\varphi, \ell_i) \leq r\rho^r \epsilon_r$$

is summable, so $\ell(x) = \lim_i \ell_i(x)$ is defined almost everywhere. And, if ℓ had an r -augmentation, say with domain G' , then at some state in this construction ℓ_i stabilizes at every point in G' . But then ℓ_i has an r -augmentation for a cofinal set of indices i , which is a contradiction. So, $E(\varphi, \ell) \leq \epsilon_r$ for all r and must be 0. \square

To get rapid decay in $E(\varphi, \ell)$ we often use expansion in the underlying graph.

Definition 1.4.12. A p.m.p. graph is **expansive** with **expansion constant** λ if

$$\lambda := \inf\{\mu(\partial A)/\mu(A) : 0 < \mu(A) < 1/2\} > 1.$$

For example, the shift graph and Bernoulli graphing for any non-amenable Cayley graph are expansive (indeed this is equivalent to non-amenableity). The Lyons–Nazarov matching theorem was one of the first instances of a measurable augmentation argument.

Theorem 1.4.13. If G is an expansive d -regular p.m.p. graph with no odd cycles, then G has a measurable perfect matching

Proof. Suppose that M is matching with no r -augmentation. Consider the set E of unmatched vertices in M . It suffices to show that $\mu(M) \leq O(\lambda^{-r/4})$. One way to get an augmentation of M is to find an alternating path (as in most proofs of Hall’s theorem). So, define E_i inductively by

- $E_0 = E$,
- $E_{2i+1} = N(E_{2i})$,
- $E_{2i+2} = \{x : \{x\} \times E_{2i+1} \cap M \neq \emptyset\}$.

For $i < r/4$ each E_{2i} must be independent, and each E_{2i+1} must be in the domain of M (otherwise we would have an r -augmentation). So, $\mu(E_i) < 1/2$, $\mu(E_{2i+1}) = \mu(E_{2i+2})$, and $\mu(E_{2i}) \leq \lambda^{-1}\mu(E_{2i+1})$. Chaining this all together, we have $\mu(E) \leq \lambda^{-r/4}$. \square

1.4.4 Shift graphs and Bernoulli graphings

We will want a somewhat more detailed analysis of the shift graph associated to a group and the Bernoulli graphing associated to a countable graph.² Recall that the vertex sets of these graphs are both spaces of labellings (in latter case considered only up to rooted isomorphism). The first fact we will want is that in many cases the choice of label spaces do not matter. First, in the Borel setting we have the following:

Theorem 1.4.14. [ST16] For any Polish space A and group Γ , there is a Borel homomorphism from the shift action $\Gamma \curvearrowright \text{Free}(A^\Gamma)$ to the shift $\Gamma \curvearrowright \text{Free}(2^\Gamma)$.

In particular, since the action of Γ is free in both cases, the homomorphism restricts to a graph embedding on each orbit of $\Gamma \curvearrowright \text{Free}(A^\Gamma)$. It follows that if we have some Borel decoration of $\text{Free}(2^\Gamma)$, then we can pull it back to a decoration of any $\text{Free}(A^\Gamma)$, in particular to $\mathcal{S}(\text{Gamma}, E)$.

In the approximate setting, the Abért–Weiss theorem says that we can transport approximate decorations between the shift graphs on $\text{Free}(A^\Gamma)$ for various Γ . In fact it says a great deal more, see Subsection 1.4.5 and Theorem 1.4.24.

²These names are somewhat unfortunate as they are easily confused with the phrase Bernoulli shift, but the three objects are distinct. These names are also somewhat standard.

For measurable decorations, matters are a little more complicated, but we have the following:

Definition 1.4.15. Given p.m.p. actions of a group Γ on space probability spaces $a : \Gamma \curvearrowright (X, \mu)$ and $b : \Gamma \curvearrowright (Y, \nu)$, say that a **factors onto** b if there is an equivariant measure preserving map $f : X \rightarrow Y$.

Theorem 1.4.16 ([Bow19]). If Γ is nonamenable, then then any two shifts $\Gamma \curvearrowright (A^\Gamma, \mu^\Gamma)$, $\Gamma \curvearrowright (B^\Gamma, \nu^\Gamma)$ factor onto each other.

For amenable groups, $([0, 1]^\Gamma, \lambda^\Gamma)$ (where λ is a Lebesgue measure) factors onto any other shift. So, measurable decorations can always be lifted from any Bernoulli shift to $\mathcal{S}(\Gamma, E)$, and if Γ is nonamenable then decorations can be transported from $\mathcal{S}(\Gamma, E)$ to any other shift.

Probabilists often want to know when a measure is a factor of independent identically distributed variables or FIID.

Definition 1.4.17. A p.m.p. action $\Gamma \curvearrowright (X, \mu)$ is Γ -FIID if there is a factor map

$$f : ([0, 1]^\Gamma, \lambda^\Gamma) \rightarrow (X, \mu).$$

The general problem of testing when a measure is FIID is usually referred to as Ornstein theory (after foundational results of Ornstein for the case $\Gamma = \mathbb{Z}$.) When we speak of a Γ -FIID solution to some labelling problem on a countable graph G (where $\Gamma \leq \text{Aut}(G)$), we mean a Γ -FIID measure on the space of solutions. It turns out Γ -FIID labellings of $\text{Cay}(\Gamma, E)$ are equivalent to measurable labellings of $\mathcal{S}(\Gamma, E)$:

Proposition 1.4.18. Let φ be any LCL, Γ be a group with a finite generating set E

1. $\text{Cay}(\Gamma, E)$ admits a Γ -FIID φ -decoration if and only if the shift graph $\mathcal{S}(\Gamma, E)$ admits a measurable φ -decoration
2. $\text{Cay}(\Gamma, E)$ admits a Γ -FIID approximate φ -decoration if and only if $\mathcal{S}(\Gamma, E)$ admits an approximate φ -decoration

Proof. We prove the proposition for vertex labellings. The proof for general labellings is much the same.

For (1), suppose Φ is a measurable d -labelling of $\mathcal{S}(\Gamma, E)$ and $x \in \text{Free}(\Gamma, E)$. Define

$$F(x)(\gamma) = \Phi(\gamma^{-1} \cdot x). \quad (1.2)$$

Then,

$$\begin{aligned} F(\delta \cdot x)(\gamma) &= \Phi(\gamma^{-1} \delta \cdot x) \\ &= F(x)(\delta^{-1} \gamma) \\ &= \delta \cdot F(x)(\gamma), \end{aligned}$$

so F is a factor map. And if a factor map F is given, the same algebra shows we can define a measurable labelling Φ via (1.2).

Since $\gamma \mapsto \gamma^{-1} \cdot x$ is an isomorphism from $B_n(e)$ to $B_n(x)$, Φ meets φ at x if and only if $F(x)$ meets φ at e . So the pushforward measure of F is a Γ -FIID φ -decoration if and only if Φ is a measurable φ -decoration. Applying this argument to each decoration in a sequence separately yields the approximate result, (2). \square

Something similar holds $\tilde{\mathcal{S}}(G)$: $\text{Aut}(G)$ -FIID labellings of G are equivalent to measurable labellings of $\tilde{\mathcal{S}}(G)$. First we need a more careful analysis of $\tilde{\mathcal{S}}(G)$. We restrict to the case $G = \text{Cay}(\Gamma, E)$ to ease notation (and because this is the main case we will be interested in.)

Proposition 1.4.19 ([CKT12, Lemma 7.9]). Let $G = \text{Cay}(\Gamma, E)$. For any $x \in \text{Free}(\Gamma, E)$, the map $\gamma \mapsto (R\langle\gamma\rangle) \cdot x$ is an isomorphism between G and the component of $R \cdot x$ in $\tilde{\mathcal{S}}(\Gamma, E)$.

Sketch of proof. For any $x \in \text{Free}$, $\gamma \mapsto \langle\gamma\rangle \cdot x$ embeds G into $\mathcal{S}(\Gamma, E)$, so we need to show the quotient map

$$\begin{aligned} \pi : \mathcal{S}(\Gamma, E) &\longrightarrow \tilde{\mathcal{S}}(\Gamma, E) \\ x &\longmapsto R \cdot x \end{aligned}$$

restricts to a graph isomorphism on components. Injectivity is clear as we are working in $\text{Free}(\Gamma, E)$.

Since any automorphism in R permutes the elements of E and their inverses, the relation (1.1) above tells us that any element of $r\langle\gamma\rangle \cdot x \in R\langle\gamma\rangle \cdot x$ is matched up with an element $R \cdot x$ by an edge in $\mathcal{S}(\Gamma, E)$ corresponding to the group element $r(\gamma^{-1})^{-1}$. This means the quotient map preserves edges and non-edges between elements of the same component, and applied inductively shows that the quotient map is surjective between components. \square

Note that it follows from the proof that $\tilde{\mathcal{S}}(G)$ preserves the quotient measure λ^G/R .

Proposition 1.4.20 ((c.f. the proof of [CKT12, theorem 7.7])). Let φ be any LCL, G be a G be a vertex transitive graph. Then G admits an $\text{Aut}(G)$ -FIID (approximate) φ -decoration if and only if the graph $\tilde{\mathcal{S}}(G)$ admits a measurable (approximate) φ -decoration.

Proof. Again, we prove the proposition for vertex labellings.

Let Γ be a countable group that acts transitively on G , and abbreviate $\text{Aut}_e(G)$ as R . If $\tilde{\Phi}$ is a measurable (approximate) φ -decoration of $\tilde{\mathcal{S}}(\Gamma, E)$, then Proposition 1.4.19 implies that $\Phi(x) = \tilde{\Phi}(R \cdot x)$ is a measurable (approximate) labelling of $\mathcal{S}(\Gamma, E)$ which is invariant under the action of R . As above, this gives rise to a Γ -factor map F . Since $\text{Aut}(G) = \Gamma R$, it suffices to check that F is also an R -factor map. By the relation 1.1, for any $\gamma \in \Gamma$, $\langle\gamma\rangle^{-1}r = r'\langle r^{-1}(\gamma)\rangle^{-1}$ for some r' . Then

$$\begin{aligned}
F(r \cdot x)(\gamma) &= \Phi(\langle\gamma\rangle^{-1}r \cdot x) \\
&= \Phi(r'\langle r^{-1}(\gamma)\rangle^{-1} \cdot x) \\
&= \Phi(\langle r^{-1}(\gamma)\rangle^{-1} \cdot x) \\
&= F(x)(r^{-1}(\gamma)) \\
&= (r \cdot F(x))(\gamma).
\end{aligned}$$

If F is an $\text{Aut}(G)$ -FIID φ -decoration, then as above F gives rise to an R -invariant measurable decoration of $\mathcal{S}(\Gamma, E)$, $\tilde{\Phi}$. Since $\tilde{\Phi}$ is R -invariant $\Phi(R \cdot x) = \tilde{\Phi}(x)$ defines a measurable φ -decoration of $\tilde{\mathcal{S}}(\Gamma, E)$. Again the approximate version follows from the same argument. \square

The last fact we need about these graphs is that they are expansive when Γ or G is nonamenable.

Theorem 1.4.21. Let Γ be a nonamenable group and G be a nonamenable graph.

1. The p.m.p. graph $\mathcal{S}(\Gamma, E)$ (equipped with the measure λ^Γ) is expansive [LN11].
2. The p.m.p. graph $\tilde{\mathcal{S}}(G)$ (equipped with the quotient measure) is expansive [CKT12].

1.4.5 Weak containment and continuous model theory

We will often be interested in when a problem can be solved off of an arbitrarily small (though not necessarily null) set. Weak containment and the model theory of measure algebras are useful language for these approximate problems. Weak containment of actions is well-known among descriptive set theorists and extends to p.m.p. graphs with only some minor technical subtleties. We record some of the basic ideas here.

The main sticking points are that (1) without an action to give canonical names to edges, it is not clear that any labelling of a graph can be coded by vertex labellings, and (2) it is not clear how to define the ultraproduct of a general relation on a measure space. It turns out that general labellings can always be coded by vertex labellings following the approach of Hatami, Lovász, and Szegedy, so (1) is easily sidestepped [HLS14]. To fix (2) we define ultraproducts only for *marked* graphs, i.e. graphs equipped with a sequence of generating involutions, and we note that graph theoretic questions do not depend on the choice of marking.

We will work in the simplified case of p.m.p. graphs where all components are vertex transitive and vertex transitive, say to a graph G with root e . In particular, all of our LCLs

will all be given by configurations supported on $B_n(e)$ for some n .

Definition 1.4.22. For a probability space (X, μ) , graph Borel graph G on X , and $r, k \in \mathbb{N}$, the set of (r, k) -**local statistics** of G is the set of measures on k -labellings of the rooted graph $B_r(e)$ induced by μ and measurable labellings of G . In symbols, for $f \in \mathcal{L}(G, k)$, let $P_r(f)$ be the measure on $\mathcal{L}(B_r(e), k)$ defined by

$$P_r(f)(A) = \mu(\{x : f \upharpoonright B_r(x) \in A\}).$$

Then the set of local statistics is

$$Q_{r,k}(G) = \{P_r(f) : f \in \mathcal{L}(G, e) \text{ is measurable}\}.$$

We similarly define the (r, k) -**local vertex statistics** of G as

$$Q_{r,k}^v(G) = \{P_r(f) : f \text{ is a measurable vertex labelling}\}.$$

An action $a : \Gamma \curvearrowright X$ weakly contains an action $b : \Gamma \curvearrowright Y$ if for all r, k , $Q_{r,k}^v(\text{Sch}(b, E)) \subseteq \overline{Q_{r,k}^v(\text{Sch}(a, E))}$ where the closure is taken in the weak topology on the space of measures. Similarly, a p.m.p. graph G weakly contains a p.m.p. graph H if for all r, k , $Q_{r,k}(H) \subseteq \overline{Q_{r,k}(G)}$.

Weak containment was first defined for actions, and it is straightforward to check that weak containment of actions implies weak containment of their Schreier graphs using the fact that edges in an action are indexed by generators of the group. But it turns out that this extra group structure is unnecessary to code labellings as vertex labellings.

Lemma 1.4.23 ([HLS14, Section 2]). For all $r, k \in \mathbb{N}$ there are R, K so that, if $Q_{R,K}^v(H) \subseteq \overline{Q_{R,K}^v(F)}$ then $Q_{r,k}(H) \subseteq \overline{Q_{r,k}(F)}$

Proof. Say H is d -regular. We want to show that the local statistics of vertex colorings captures all of the local statistics. We will sketch the main idea for undirected edge colorings, the general case is almost identical. Given an edge labelling in $B_{r,k}$ of H , say f , using a

coloring result of Kechris–Solecki–Todorćević we can find N large enough (depending only on H) and an edge N -labelling of H g so that $g(e) \neq g(e')$ when e, e' are within distance 3 in the line graph. Then $\tilde{f}(x) = \{(f(e), g(e)) : e \in B_r(x)\}$ has the property that if $e = \{x, y\}$ then $\tilde{f}(x) \cap \tilde{f}(y) = \{(f(e), g(e))\}$. So, we can recover f and g from \tilde{f} and vice versa. \square

Note that if φ is an LCL with radius of definition r , then H has a φ -decoration if and only if there is a measure in $Q_{r,k}(H)$ supported on the configurations allowed by φ . And φ has an approximate φ -decoration if and only if there is such a measure on $\overline{Q_{r,k}(H)}$. Indeed, $\rho^\varphi(H)$ can be computed from $Q_{r,k}(H)$ for the appropriate r, k .

The Abért–Weiss theorem says that $\mathcal{S}(\Gamma, E)$ is minimal with respect to weak containment among all Γ -Schreier graphs [AW11]. A theorem of Hatami, Lovász, and Szegedy says $\tilde{\mathcal{S}}(G)$ is similarly minimal among all graphs which have all components isomorphic to G . (In fact they prove a more general result for unimodular random graphs, but we only need this simple case).

Theorem 1.4.24. [HLS14, Theorem 1.6] If H is a p.m.p. graph and every component of H is isomorphic to G , then H weakly contains $\tilde{\mathcal{S}}(G)$.

Weak containment is intimately tied up with ultraproducts of actions (see, for instance [CKT12]). We will use the language of model theory of metric structures to extend this to p.m.p. graphs (see [BHU06] for proofs and background). In particular, we will define ultraproducts model theoretically via the measure algebra and note that they can be represented by graphs on nonstandard measure spaces using the Löb measure construction.

Definition 1.4.25. For a p.m.p. graph H , a **marking** is a sequence $\langle f_i : i \in \mathbb{N} \rangle$ such that $G = \{(x, f_i(x)) : x \in X, x \neq f_i(x), i = 1, \dots, n\}$. A **marked graph** is a p.m.p. graph equipped with a marking.

If H is a bounded degree p.m.p. graph, then by standard coloring results, H has a marking with finitely many functions. For instance, the Kechris–Solecki–Todorćević [Pik21, Theorem 5.12] coloring theorem says that $2\Delta + 1$ is sufficient (where Δ is a degree bound for G).

Definition 1.4.26. For a complete measure space (X, μ) , the **measure algebra** of X , $M(X)$ is the following structure:

$$M(X) := (MALG(X), d_\mu, \mu, \cup, \cap, \neg)$$

where $MALG(X)$ is the Boolean algebra of measurable (equivalently Borel, equivalently G_δ) sets modulo the null sets, $d_\mu(A, B) = \mu(A \Delta B)$, and $\neg A = \{x \in X : x \notin A\}$.

We conflate a marked p.m.p. graph (H, g_1, \dots, g_n) on vertex set X with the associated metric structure is $(M(X), g_1, \dots, g_n)$, where $g_i : MALG(X) \rightarrow MALG(X)$ is the set-wise action of G_i :

$$g_i(A) = \{g_i(x) : x \in A\}.$$

For function symbols over X f_1, \dots, f_n with codomain X and predicate symbols over X r_1, \dots, r_m with codomain $[0, 1]$, a **formula** is a formal composition of $\inf_X, \sup_x, f_1, \dots, f_n, r_1, \dots, r_m$, and operations on the interval so that the domain is X^n and the codomain is the interval. For instance $\inf_{A_1} \mu(A_1 \cup A_2)/2$ is a formula for measure algebras.

A formula is a sentence if it has no free variables and it is quantifier free if \sup and \inf do not appear in it. For a formula φ and $x_1, \dots, x_n \in X$, we write

$$\mathbb{X} \models \varphi(x_1, \dots, x_n) \leq \epsilon$$

to mean that the function obtained by interpreting the function and predicate symbols as indicated by \mathbb{X} takes value at most ϵ at x_1, \dots, x_n .

We say structures equipped with the same function and relation symbols are elementary equivalent and write $\mathbb{X} \equiv \mathbb{Y}$ if $\mathbb{X} \models \varphi \leq \epsilon \Leftrightarrow \mathbb{Y} \models \varphi \leq \epsilon$ for every sentence φ . And we say that \mathbb{X} is an elementary substructure of \mathbb{Y} and write $\mathbb{X} \prec \mathbb{Y}$ if $\mathbb{X} \subseteq \mathbb{Y}$ and for any $x_1, \dots, x_n \in \mathbb{X}$ and formula φ

$$\mathbb{X} \models \varphi(x_1, \dots, x_n) \leq \epsilon \Leftrightarrow \mathbb{Y} \models \varphi(y_1, \dots, y_n).$$

Given a sequence of metric structures \mathbb{X}_i (all with the same function symbols) and an nonprincipal ultrafilter \mathcal{U} , the **ultraproduct** is the structure $\prod_{\mathcal{U}} \mathbb{X}_i$ defined as follows. For

$x, y \in \prod_i X_i$ say $f \sim_{\mathcal{U}} g$ if $\{i : d(x(i), y(i)) \leq \epsilon\} \in \mathcal{U}$ for all $\epsilon > 0$. The ultraproduct has domain $\prod_i X_i / \sim_{\mathcal{U}}$, and we interpret a function symbol f by $f([x]) = [\langle f_i(x(i)) : i \in \omega \rangle]$ and predicate r by $r(x) = y :\Leftrightarrow \{i : |r_i(x(i)) - y| < \epsilon\} \in \mathcal{U}$ for all $\epsilon > 0$.

In the case that the \mathbb{X}_i are all the same structure \mathbb{X} , we write $\mathbb{X}^{\mathcal{U}} := \prod_{\mathcal{U}} \mathbb{X}_i$ and this structure an **ultrapower**.

The most important facts about the ultraproduct is Łoś's theorem.

Theorem 1.4.27. For any sentence φ , $\prod_i \mathbb{X}_i \models \varphi \leq \epsilon$ if and only $\{i : \mathbb{X}_i \models \varphi \leq \epsilon\} \in \mathcal{U}$. In particular, any ultrapower of \mathbb{X} is always elementary equivalent to \mathbb{X}

It will also be convenient in our context to know that the ultraproduct is realized by p.m.p. functions on a (nonstandard) measure space, and that separable elementary substructures of the ultrapower are represented by p.m.p. graphs on standard measure spaces.

Lemma 1.4.28. If (H, g_1, \dots, g_n) is a marked p.m.p. graph and \mathcal{U} is a non-principle ultrafilter then there is a (nonstandard) measure space X with p.m.p functions f_1, \dots, f_n so that $(H, g_1, \dots, g_n)^{\mathcal{U}} \cong (M(X), f_1, \dots, f_n)$. Further, if \mathbb{X} is a separable structure and $\mathbb{X} \equiv (H, g_1, \dots, g_n)$, then \mathbb{X} is isomorphic to the structure associated with some marked p.m.p. graph.

Proof. Proposition Number in [CKT12] tells us that $(M(Y), g)^{\mathcal{U}}$ is realized by a p.m.p. function. Applying this separately to each g_i gives the first statement. The second statement similarly follows from the results in [BHU06, Chapter 18]. \square

Given a nonstandard space X and f_1, \dots, f_n as in the above theorem, we can consider the graph on X generated by f_1, \dots, f_n . Abusing notation slightly, we refer to this graph as $H^{\mathcal{U}}$. Now Łoś's theorem can be reformulated for p.m.p. graphs as follows:

Lemma 1.4.29. For any p.m.p. graph H and nonprincipal ultrafilter \mathcal{U} , $\overline{Q_{r,k}(H)} = Q_{r,k}(H^{\mathcal{U}})$.

Proof. By Proposition 1.4.23, it is enough to work with vertex labellings. We can code vertex labellings by tuples in the measure algebra: for $A = \langle A_d : d \in D \rangle \in \text{MALG}(X)^D$, $g_A(x) = d \Leftrightarrow x \in A_d \setminus \bigcup_{d' \neq d} A_{d'}$.

By Łoś's theorem, it suffices to show that, for any probability measure p on $\mathcal{L}(B_r(e), k)$ there is some quantifier free φ so that $H \models \varphi(A) = |p - \tilde{g}_A|_{TV}$:

$$|p - \tilde{g}_A|_{TV} = \max\{|p(f) - \mu(\{x : (g_A \upharpoonright B_n(x)) \cong f\})| : f \in \mathcal{L}(B_n(e), d)\}$$

and $\{x : (g_A \upharpoonright B_n(x)) = f\}$ can be constructed from A using Boolean combinations and the marking of H . □

The last fact we will need is a version of the Lowenheim–Skolem theorem.

Theorem 1.4.30. For any structure \mathbb{X} and $x_1, \dots, x_n \in \mathbb{X}$ there is a separable elementary substructure $\mathbb{Y} \prec \mathbb{X}$ so that $x_1, \dots, x_n \in \mathbb{Y}$.

As one a consequence, if there is a p.m.p. graph H with an approximate φ -decoration for some LCL φ , then there is an elementary equivalent p.m.p. graph with a measurable φ -decoration. This is because the ultrapower of any marking of H has a measurable φ -decoration, and by Lowenheim–Skolem we can take some separable elementary equivalent substructure of the ultrapower containing the partition corresponding to this φ -decoration.

1.4.6 Borel counterexamples: Baire category and determinacy

In the purely Borel setting, Baire category is by far the most common method for obtaining counterexamples. Recall that Baire's theorem says any countable intersection of dense G_δ sets in a complete metric space is again dense and G_δ . This gives us notion of largeness (or dually smallness) we call being comeager (or meager).

Definition 1.4.31. A set A in a complete metric space is comeager if it contains a dense G_δ set. And A is **Baire measurable** if $A \Delta U$ is meager for some open set U .

A function $f : X \rightarrow Y$ is **Baire measurable** if $f^{-1}(U)$ is Baire measurable for every open set $U \subseteq Y$.

Every Borel set is Baire measurable and a theorem of Solovay (using an inaccessible cardinal) and Shelah (in ZFC) says that, absent choice, it is consistent for every set to be Baire measurable [She84].

Let us illustrate how Baire category is used with a simple example:

Proposition 1.4.32. There is an acyclic Borel graph with no Baire measurable 2-coloring.

Proof. Let R be an irrational rotation of S^1 , and consider the graph G on S^1 given by

$$(x, y) \in G :\Leftrightarrow x = R(y) \text{ or } y = R(x).$$

Since R is irrational, G is acyclic. If $f : S^1 \rightarrow \{0, 1\}$ is Baire measurable 2-coloring, then $f^{-1}(i)$ is nonmeager for some i . That is, for some meager M and open U , f is constant on $U \setminus M$.

Since R is irrational, every orbit under R^2 is dense. This means $R^{2n+1}(U) \cap U \neq \emptyset$ for some n . Also $R^{2n+1}(M)$ is still meager. So $R^{2n}(U) \cap U \setminus (M \cup R^{2n}(M)) \neq \emptyset$. But this means $f(x)$ and $f(R^{2n+1}(x))$ agree, so f is not a 2-coloring. \square

Corollary 1.4.33. It does not follow from ZF + DC that every acyclic graph is 2-colorable

We will see that for CSPs, every problem that has a Borel instance with no solution comes from a Baire category argument.

Tools for producing Borel counterexamples which do not rely on measure or Baire category are few and far between. The most widely applicable is Marks's determinacy method. Mark's main lemma is below. The proof (which we omit) relies on Martin's celebrated Borel determinacy theorem, a deep result with provable set theoretic left.

Theorem 1.4.34. For any countably infinite groups Γ, Δ and Borel set $A \subseteq \text{Free}(\mathbb{N}^{\Gamma * \Delta})$, either there is an equivariant embedding from $\Gamma \curvearrowright \text{Free}(\mathbb{N}^\Gamma)$ into $\Gamma \curvearrowright A$ or there is an equivariant embedding from $\Delta \curvearrowright \text{Free}(\mathbb{N}^\Delta)$ into $\Delta \curvearrowright \text{Free}(\mathbb{N}^{\Gamma * \Delta}) \setminus A$.

As an application, we can compute the Borel chromatic number of free groups.

Theorem 1.4.35. There is no Borel $2n$ coloring of $\mathcal{S}(F_n, E)$ where E is the usual generating set.

Proof. We proceed by induction. For $n = 1$, Proposition 1.4.32 says that there is some \mathbb{Z} action with no Borel 2-coloring and Theorem 1.4.14 then says that $\mathcal{S}(\mathbb{Z})$ has no 2-coloring. Suppose $f : 2^{F_{n+1}} \rightarrow 2(n+1)$ is a Borel proper coloring. Let $A = \{x : f(x) \in 2n\}$. By the lemma, either 2^{F_n} embeds into A or $2^\mathbb{Z}$ embeds into $2^{F_n} \setminus A$. In the former case, we get an n -coloring of 2^{F_n} and in the latter case we get a 2-coloring of $2^\mathbb{Z}$, these are both impossible by the induction hypothesis. \square

Note that this is optimal by Kechris–Solecki–Todorćevic [Pik21, Theorem 5.12].

1.4.7 Dichotomies, Gandy–Harrington forcing, and complexity

The search for dichotomy theorems is something of an industry in Borel combinatorics. The G_0 dichotomy of Kechris–Solecki–Todorćevic is a prototypical example.

Definition 1.4.36. Let $\langle \sigma_n : n \in \omega \rangle$ be a sequence of binary strings so that $|\sigma_n| = n$ and so that any string has some σ_n as an extension. Define a graph G_0 on 2^ω by

$$(x, y) \in G_0 \Leftrightarrow (\exists i \in 2, z \in 2^\omega, n \in \omega) x = \sigma_n \hat{\ } i \hat{\ } z \wedge y = \sigma_n \hat{\ } (1 - i) \hat{\ } z$$

Theorem 1.4.37 ([KST99]). For any Borel graph G , either G has a countable Borel coloring or there is a Borel homomorphism from G_0 to G .

For a proof see [Mar19, Theorem 3.20]. Miller has shown that a great deal of classical descriptive set theory can be deduced from this theorem. For instance, one can recover

the Luzin–Novikov theorem and Silver’s theorem on Borel equivalence relations (i.e. the continuum hypothesis for Borel quotients) from the G_0 dichotomy.

Other dichotomies include the Harrington–Kechris–Louveau dichotomy characterizing smooth equivalence relations [HKL90], Hjorth and Miller’s dichotomies on end selection [HM09], and Solecki’s dichotomy characterizing Baire class 1 functions [Sol98]. These dichotomies typically have effective proofs via Gandy–Harrington forcing which give slightly more information. For instance, the effective version of the G_0 dichotomy is:

Theorem 1.4.38. For any parameter p , if G is a $\Delta_1^1(p)$ graph, then either G has a $\Delta_1^1(p)$ countable coloring or there is a continuous homomorphism from G_0 into G .

Corollary 1.4.39. If G is Δ_1^1 , then G has a Borel countable coloring if and only if G has a Δ_1^1 countable coloring

Corollary 1.4.40. The set of Borel countably coloring graphs is Π_1^1 in the codes.

Proof. Let C be the set of codes for graphs of Borel countably colorable graphs. The set of codes for Borel graphs is Π_1^1 , and

$$c \in C :\Leftrightarrow c \text{ codes a graph, and } (\exists f \in \Delta_1^1(c))(\forall x \neq y) f(x) \in \omega \wedge f(x) \neq f(y).$$

There are countably many $\Delta_1^1(c)$ codes and there is a Π_1^1 coding of Δ_1^1 sets, so this is a Π_1^1 definition. \square

In general, this kind of effectivization or complexity result can be viewed as a weak dichotomy theorem. Most of the effectivization and complexity consequences of the Gandy–Harrington forcing method are packaged in the lemma below. Readers unfamiliar with the technique can treat this lemma as a black box. We include a brief sketch of the argument for readers who are familiar.

Definition 1.4.41. We say that a property $\Phi(A)$ of a set is an **independence property** if there is some Δ_1^1 property ϕ such that

$$\Phi(A) \Leftrightarrow \neg(\exists \mathbf{x}, y) \mathbf{x} \in A^k \text{ and } \phi(x_1, \dots, x_n, y).$$

A property $\Psi(A)$ of sets is a **closure property** if there is a Δ_1^1 property ψ so that $\Psi(A)$ if and only if

$$(\forall \mathbf{x} \in A^k)(\forall y, z) \psi(\mathbf{x}, y, z) \rightarrow z \in A.$$

For a closure property Ψ and set A , the Ψ -closure of A is $A^\Psi := \bigcup_m f^m(A)$, where

$$f(A) := A \cup \{z : (\exists \mathbf{x} \in A^k)(\exists y)\psi(x, y, z)\}.$$

Theorem 1.4.42. Let Φ be an independence property and Ψ be a closure property. Suppose $X \in \Delta_1^1$ and $X \subseteq \bigcup_{i \in \omega} B_i$, where each B_i is a Borel set so that $(\Phi \wedge \Psi)(B_i)$. Then there is a Δ_1^1 sequence of Δ_1^1 sets $\langle A_i : i \in \omega \rangle$ so that $\Phi(A_i)$ for all i and $X \subseteq \bigcup_{i \in \omega} A_i$.

Proof sketch. Let \mathbb{P}_n be Gandy–Harrington forcing on \mathcal{N}^n , i.e. the set of nonempty Σ_1^1 subsets of \mathcal{N}^n . Suppose $A \in \mathbb{P}_1 \Vdash \dot{x} \in \dot{B}$, where \dot{x} is a name for the generic real and B is a Borel set so that $\Phi(B)$. Then, we must have that $\Phi(A)$ as well, otherwise $p := \{(x, y) \in A^2 : \phi(x, y)\} \in \mathbb{P}_2$ and

$$p \Vdash \dot{x}, \dot{y} \in \dot{B} \wedge \phi(\dot{x}, \dot{y})$$

contradicting the absoluteness of $\Phi(B)$. By a similar argument, we have that if $A \Vdash \dot{x} \in \dot{B}$ and $\Psi(B)$, then $A^\Psi \Vdash \dot{x} \in \dot{B}$.

Now suppose $A \Vdash \dot{x} \in \dot{B}$ where $(\Phi \wedge \Psi)(B)$. Combining the above observations, $(\Phi \wedge \Psi)(A^\Psi)$. By iteratively taking closures and using the reflection theorem [Mar19, Theorem 2.27] we can find $A \subseteq A' \in \Delta_1^1$ where $(\Phi \wedge \Psi)(A')$.

Suppose toward contradiction that X is not covered by Δ_1^1 sets satisfying $(\Phi \wedge \Psi)$, but is covered by Borel sets with this property. Define

$$X' := X \setminus \bigcup \{A \in \Delta_1^1 : (\Phi \wedge \Psi)(A)\}.$$

Note that X' is nonempty and Σ_1^1 , so $X' \Vdash \dot{x} \in \bigcup_i \dot{B}_i$, where each \dot{B}_i is Borel and satisfies $(\Phi \wedge \Psi)$. But then we can refine X' to find a condition A and index i with $A \Vdash \dot{x} \in \dot{B}_i$. But then A is contained a Δ_1^1 set satisfying $(\Phi \wedge \Psi)$ and is disjoint from X' .

Lastly, we can convert from a covering to a sequence as follows. Being a covering by Δ_1^1 sets satisfying $(\Phi \wedge \Psi)$ is a Π_1^1 on Π_1^1 property. The set of all such sets is a Π_1^1 covering, so by reflection, there is a Δ_1^1 covering, which admits a Δ_1^1 enumeration. \square

The last step in the proof where we use reflection to enumerate a cover is typical. It lets us, for instance, construct a countable coloring of a graph given that the graph is covered by independent Δ_1^1 sets.

The proof above relies on the fact that if Φ is a closure or independence property, $A \Vdash \dot{x} \in \dot{B}$, and $\Phi(B)$, then there is a Δ_1^1 superset of A which also satisfies Φ . The paper [Tho21] gives more details, generalizations, and applications. As a quick application, we can effectivize the Luzin–Novikov and Feldman–Moore theorem.

Theorem 1.4.43. If R is a Δ_1^1 locally countable relation, then $R = \bigcup_i f_i$, where each f_i is a Δ_1^1 partial function. Further if R is an equivalence relation, then each f_i can be taken to be an involution.

Proof. Being a partial function contained in R is an independence property:

$$f \subseteq X^2 \text{ is a partial function in } R :\Leftrightarrow f \subseteq R \wedge (\forall x, y, z)(x, y) \in f \wedge (x, z) \in f \rightarrow y = z.$$

Being a partial involution is the conjunction of a closure and independence property:

$$f \text{ is a partial involution } :\Leftrightarrow f \text{ is a partial function, and } (\forall x, y)(x, y) \in f \rightarrow (y, x) \in f.$$

By the classical Luzin–Novikov and Feldman–Moore theorems, Theorems 1.4.2 and 1.4.3, if R is a locally countable Δ_1^1 (equivalence) relation, then R is a union of countably many partial injections (or involutions). Then by Theorem 1.4.42, R is a union of Δ_1^1 partial injections (or involutions). And if R is a union of partial involutions, then we can extend each partial involution to a total involution by setting $f(x) = x$ for $x \notin \text{dom}(f)$. \square

Note that, under any reasonable interpretation of the phrase, a dichotomy theorem also gives a nontrivial complexity upper bound of Δ_2^1 for a Borel combinatorial problem (a solution exists if and only if some canonical obstruction does not exist). So to rule out dichotomies (and effectivization etc.), it suffices to prove a complexity lower bound. The most spectacular example of this is Todorćević and Vidnyansky's result on Borel 3-coloring.

Theorem 1.4.44 ([TV21]). The set of Borel 3-colorable locally finite graphs is Σ_2^1 -complete in the codes.

In the chapter on CSPs, we will use some algebraic machinery to propagate this lower bound to all NP-complete CSPs.

CHAPTER 2

Locally Checkable Labelling Problems

In this chapter we try to solve (or present obstructions to solving) various locally checkable labelling problems (or LCLs). And, we try to fit our results into the picture of local complexity outlined in the introduction. As discussed in the introductions, our results sharpen the bounds for definable versions of the Lovász local lemma and answer some open questions in measurable combinatorics.

2.1 Orientations

Our first family of LCLs amounts to the following: What is the minimum outdegree we need to orient a given graph?

Definition 2.1.1. For G a graph on a set X , we say that an orientation o of G is a **k -orientation** if $\sup_{x \in X} \text{out}_o(x) \leq k$.

Equivalently, a k -orientation is an φ_k -decoration where φ_k is the set of oriented radius 1 graphs where the root has outdegree at most k .

The **orientation number** of G , $\mathfrak{o}(G)$, is the least cardinal k so that G is k -orientable.

$$\mathfrak{o}(G) := \min \{k : \text{there is a } k\text{-orientation of } G\}.$$

We get the standard descriptive variants. For X a standard Borel space, μ a Borel measure on X , and G a graph of X :

$$\mathfrak{o}_B(G) := \min \{k : \text{there is a Borel } k\text{-orientation of } G\}.$$

$$\mathfrak{o}_\mu(G) := \min \{k : \text{there is a } \mu\text{-measurable } k\text{-orientation of } G\}.$$

For example, $\mathfrak{o}(G) = \lceil d/2 \rceil$ whenever G is a d -regular finite graph, $\mathfrak{o}(G) = 1$ whenever G is acyclic, an undirectable forest of lines (as in [KM04, Example 6.8]) is equivalent to a 2-regular Borel graph with $\mathfrak{o}_B(G) = 2$, and Marks finds 3-regular graphs with $\mathfrak{o}_B(G) = 3$ in [Mar13].

This kind of orientation problem shows up when we want to recover functions which generate a given graph. The Luzin-Novikov uniformization theorem says that if $\mathfrak{o}_B(G) \leq \aleph_0$, then it is equal to the number of functions f_1, \dots, f_k so that

$$(x, y) \in G \Leftrightarrow (\exists i) y = f_i(x) \text{ or } x = f_i(y).$$

This is closely related to the parameter k_B studied by Csóka, Lippner, and Pikhurko [CLP16]. And, later, we will use orientations as a stepping stone to recovering group actions generating regular trees.

Orientation numbers provide an interesting test case for adapting methods from finite combinatorics. In the classical setting, $\mathfrak{o}(G)$ is well understood in terms of partitions of G into sidewalks.¹

Proposition 2.1.2. For any graph G on any set X ,

$$\mathfrak{o}(G) = \min \left\{ |S| : S \subseteq \mathcal{P}(G), \bigcup S = G, \text{ and } (\forall s \in S) s \text{ is a sidewalk} \right\}$$

and if $\mathfrak{o}(G)$ is finite,

$$\mathfrak{o}(G) = \max_{S \subseteq X \text{ finite}} \left\lceil \frac{\sum_{x \in S} |G_x \cap S|}{2\rho(S)} \right\rceil$$

where $\rho(S) = |S| - \#\text{acyclic components of } G \upharpoonright S$.

The first statement is elementary. For the second, observe that the sidewalks in a graph are the independent sets of a matroid (usually called the bicircular matroid). The result

¹A sidewalk is graph where each component has at most 1 cycle, i.e. where bicycles are not allowed. These are sometimes called psuedoforests.

then follows from compactness and the Edmonds covering theorem [SU13, Theorem 5.3.2]. Curiously, this characterization in terms of sidewalk covering fails in definable contexts (see Theorem 2.1.19), but we can still recover a measurable version of the Edmonds formula:

Theorem 2.1.3 (See Corollary 2.1.7). If G is a p.m.p. graph with bounded degree, and k is an integer with $k > \text{cost}(G \upharpoonright A)$ for $\mu(A) > 0$, then $\mathfrak{o}_\mu(G) \leq k$.

In many cases, this upper bound turns out to be optimal, and we also have a sharper analysis for expansive graphs (see Theorem 2.1.10). Our proof of Theorem 2.1.3 is robust in the sense that it gives a bound on \mathfrak{o}_μ for μ quasi-invariant in terms of the essential supremum of the Radon–Nikodym cocycle (see Theorem 2.1.7). As in Conley–Tamuz [CT20], this yields a Borel result for graphs of subexponential growth.

These results are enough to compute the measurable orientation numbers of many naturally arising graphs. In particular, we can find optimal FIID orientations for nonamenable groups. We also compute Borel orientation number of some graphs, such as the Hadwiger–Nelson unit distance graph.

2.1.1 Measurable constructions

For this section, G will be a locally countable Borel graph, the letter μ will always stand for a G -quasi-invariant Borel probability measure on X , $\rho(x, y)$ stands for the associated Radon–Nikodym cocycle, and $\rho := \text{ess sup}_{x,y} \rho(x, y)$. In particular μ is G -invariant if and only if $\rho = 1$. We can express the handshake lemma in terms of orientations as follows:

Proposition 2.1.4. For any μ , orientation o ,

$$\int_{x \in X} \text{in}_o(x) d\mu \leq \rho \int_{x \in X} \text{out}_o(x) d\mu.$$

We use similar propositions throughout this section without comment. Readers unfamiliar with quasi-invariant measures can consult [KM04], or consider the p.m.p. case where $\rho = 1$ to get a gist of the arguments.

Definition 2.1.5. Let $\alpha_\mu(G) = \sup\{\text{cost}(G \upharpoonright A) : A \subseteq X, \mu(A) > 0\}$, where the cost of a restriction is computed with respect to the normalized measure $\mu/\mu(A)$. That is,

$$\text{cost}(G \upharpoonright A) = \frac{1}{2\mu(A)} \int_{x \in A} |G_x \cap A| d\mu.$$

For example, if G is d -regular, $\alpha_\mu(G) = \frac{d}{2}$ regardless of the measure μ . And, if G is a finite graph equipped with counting measure, this is essentially the Edmonds formula for \mathfrak{o} .

Proposition 2.1.6. For any μ ,

$$\mathfrak{o}_\mu(G) \geq \left\lceil \frac{2\alpha_\mu(G)}{1 + \rho} \right\rceil.$$

In particular, if μ is G -invariant, $\mathfrak{o}_\mu(G) \geq \lceil \alpha_\mu(G) \rceil$.

Proof. Since $\mathfrak{o}_\mu(G \upharpoonright A) \leq \mathfrak{o}_\mu(G)$ when $\mu(A) > 0$, it suffices to show

$$\mathfrak{o}_\mu(G) \geq \frac{2}{1 + \rho} \text{cost}(G).$$

Suppose G has an orientation o with outdegree bounded by n . Then,

$$\begin{aligned} \int_X |G_x| d\mu &\leq \int_X \text{out}_o(x) d\mu + \int_X \text{in}_o(x) d\mu \\ &\leq (1 + \rho) \int_X \text{out}_o(x) d\mu \end{aligned}$$

so

$$\text{cost}(G) = \frac{1}{2} \int_X |G_x| d\mu \leq \frac{1 + \rho}{2} \int_X \text{out}_o(x) d\mu \leq \frac{1 + \rho}{2} n.$$

□

The next theorem says that, for bounded degree graphs with ρ small this lower bound is close to sharp.

Theorem 2.1.7. If G has bounded degree, then for any $k \in \mathbb{N}$, there is a Borel orientation o such that, for any measure μ , if $k > \rho^2 \alpha_\mu(G)$, then

$$\mu(\{x : \text{out}_o(x) > k\}) = 0.$$

Proof. We will use an augmentation argument. We conflate an orientation with its characteristic function. Let Δ be a degree bound for G . Given an orientation o of G , Define

$$O_o := \{x : \text{out}_o(x) > k\}, \quad \text{and } I_o := \{x : \text{out}_o(x) < k\}.$$

So, if φ_k is the set of diameter-1 graphs oriented with outdegree bounded by k , $O_o = E(\varphi_k, o)$. So, by Lemma 1.4.11, it suffices to show that if o admits no r -augmentation and μ is as in the theorem statement, then $\mu(O_o) \leq c^r$ for some $0 \leq c < 1$. One way to get an augmentation is to find a path $p = (x_0, x_1, \dots, x_n)$ which is an oriented path from O_o to I_o , i.e. such that

1. $(x_i, x_{i+1}) \in o$ for $1 \leq i \leq n$, and
2. $\text{out}_o(x_1) > k$ and $\text{out}_o(x_n) < k$.

We can augment p by flipping every edge along the p : $o'(e) = 1 - o(e)$ for e an edge in p . Then, the degree of vertices in p only changes for endpoints of the path, and these move towards k .²

Fix μ and abbreviate $\alpha := \alpha_\mu(G)$. We verify that if o has no augmenting chains of length at most n , then $\mu(O_o) \leq \left(\frac{\rho\alpha}{k}\right)^n$. Recall

$$B_n^+(A) := \{y : \text{there is an oriented path of length at most } n \text{ from some } x \in A \text{ to } y\}$$

and $B^+(A) := B_1^+(A)$. Notice that $B_{n+m}^+(A) = B_n^+(B_m^+(A))$.

Claim: If every point in $A \subseteq G$ has outdegree at least k , then

$$\mu(B^+(A)) \geq \frac{k}{\rho\alpha} \mu(A).$$

Using the facts that every edge coming out of A ends in $B^+(A)$ and that $\text{in}(x) + \text{out}(x) =$

²Technically, this is not exactly an r -augmentation as we may not satisfy φ_k on the nose anywhere new, but Lemma 1.4.11 easily adapts.

$\deg(x)$ for the graph $G \upharpoonright B^+(A)$, we have

$$\begin{aligned}
k \mu(A) &\leq \int_{x \in A} \text{out}_o(x) d\mu \\
&\leq \frac{1}{2} \left(\int_{x \in A} \text{out}_{o \upharpoonright B^+(A)}(x) d\mu + \rho \int_{x \in B^+(A)} \text{in}_{o \upharpoonright B^+(A)}(x) d\mu \right) \\
&\leq \frac{\rho}{2} \int_{x \in B^+(A)} (\text{out}_{o \upharpoonright B^+(A)}(x) + \text{in}_{o \upharpoonright B^+(A)}(x)) d\mu \\
&\leq \frac{\rho}{2} \int_{x \in B^+(A)} |G_x \cap B^+(A)| d\mu \\
&= \rho \text{cost}(G \upharpoonright B^+(A)) \mu(B^+(A)) \\
&\leq \rho \alpha \mu(B^+(A)).
\end{aligned}$$

Now suppose that o admits no augmenting chains of length at most n . In this case, any oriented path starting in O_o of length at most n must fail to reach I_o . That is, for $i \leq n$, $B_i^+(O_o)$ satisfies the hypotheses of the claim. So,

$$1 \geq \mu(B_n^+(O_o)) \geq \left(\frac{k}{\rho\alpha}\right)^n \mu(O_o).$$

□

Since initial publication, it has been pointed out that the above argument follows a very similar outline to the algorithm given by Ghaffari and Su in [GS17]. Their algorithm runs in polylog rounds in the CONGEST model. Note that polylog time in the LOCAL model is insufficient to draw any descriptive set theoretic conclusions. The CONGEST model tracks not only the number messages sent, but the size of these messages. It is open if this model is as intimately connected to descriptive combinatorics as the LOCAL model, but the above theorem suggests that such connections are possible.

For p.m.p. graphs, we get the following.

Corollary 2.1.8. If G is p.m.p. with measure μ , then

$$\mathfrak{o}_\mu(G) \in [\alpha_\mu(G), \alpha_\mu(G) + 1].$$

In particular, if $\alpha_\mu(G)$ is not an integer, $\mathfrak{o}_\mu(G) = \lceil \alpha_\mu(G) \rceil$.

It follows that any d -regular p.m.p. graph can be generated by $\lceil (d+1)/2 \rceil$ functions. Similar results have been obtained by Grebík and Pikhurko [GP20, Theorem 1.7]. Using an idea from Conley and Tamuz [CT20], we can get a Borel result for slow-growing graphs.

Corollary 2.1.9. If G has subexponential growth and is d -regular, then $\mathfrak{o}_B(G) \in [\frac{d}{2}, \frac{d}{2} + 1]$.

Proof. In this case, for any vertex x of G , there is an atomic measure with $\rho < \sqrt{1+d}$ whose support contains x . The Borel orientation given by Theorem 2.1.7 witnesses $\mathfrak{o}_\mu(G) \leq \frac{d}{2} + 1$ for all of these atomic measures, so witnesses $\mathfrak{o}_B(G) \leq \frac{d}{2} + 1$. Similarly, Proposition 2.1.6 gives a lower bound. \square

We can also sharpen the analysis for expansive regular graphs. If G is expansive, then edge boundaries in G are large. So $\text{cost}(G \upharpoonright A)$ is bounded away from $\alpha_\mu(G)$ when A has measure less than $1/2$. If G is regular, the problem is symmetric enough that we only need to consider these small sets.

Theorem 2.1.10. If μ is G -invariant, and G is d -regular and expansive then

$$\mathfrak{o}_\mu(G) = \left\lceil \frac{d}{2} \right\rceil.$$

Proof. It is enough to consider the case with d even. We modify the proof of Theorem 2.1.7, making use of the following symmetry. For any orientation o , let $o^{-1} = \{(x, y) : (y, x) \in o\}$. That is, o^{-1} is o with every edge flipped. We have $\text{out}_o(x) = \text{in}_{o^{-1}}(x)$, $B_o^+(A) = B_{o^{-1}}^-(A)$, and

$$I_o = \{x : \text{out}_o(x) < d/2\} = \{x : \text{in}_o(x) > d/2\} = O_{o^{-1}}.$$

Claim 1: There is $c > 1$ such that, for any orientation o , if $A \subseteq X$ satisfies $0 < \mu(B^+(A)) \leq 1/2$

and $\text{out}_o(x) \geq d/2$ for all $x \in A$, then

$$c \mu(A) \leq \mu(B^+(A))$$

Let λ be the expansion constant of G . Then

$$\begin{aligned} \frac{d}{2} \mu(A) &\leq \int_{x \in A} \text{out}_o(x) d\mu \\ &\leq \frac{1}{2} \int_{x \in B^+(A)} (\deg(x) - |G_x \setminus B^+(A)|) d\mu \\ &\leq \frac{1}{2} (d - \lambda) \mu(B^+(A)). \end{aligned}$$

So, $c = \frac{d}{d-\lambda}$ works.

Symmetrically, if $\mu(B^+(A)) \leq \frac{1}{2}$ and $\text{out}_o(A) \leq d/2$ for $x \in A$, then applying Claim 1 to o^{-1} gives

$$c \mu(A) \leq \mu(B^-(A)).$$

As in the proof of 2.1.7, we derive an exponential bound on $\min(\mu(O_o), \mu(I_o))$.

Claim 2: There is $c > 1$ such that, if o admits no augmenting chains shorter than length $2n$, then $\min(\mu(I_o), \mu(O_o)) < c^{-n}$.

In such an orientation, $B_n^+(O_o) \cap B_n^-(I_o) = \emptyset$. So one of them must have measure bounded by $1/2$. Possibly replacing o with o^{-1} , we can assume $\mu(B_n^+(O_o)) \leq 1/2$. By claim 1,

$$\frac{1}{2} \geq \mu(B_n^+(O_o)) \geq c^n \mu(O_o)$$

for some $c > 1$. Thus, $\mu(O_o) \leq \frac{1}{2c^n}$

Now the same augmentation argument as in Theorem 2.1.7 gives an orientation, o , of G with either O_o or I_o empty. In the former case, o is a $(d/2)$ -orientation, in the latter case o^{-1} is a $(d/2)$ -orientation. \square

Note that the orientations produced by the above theorem are balanced.

Corollary 2.1.11. If G is d -regular with d even and μ is G -invariant, then any measurable $(d/2)$ -orientation of d has indegree $(d/2)$ μ -a.e.

Proof. Otherwise, $a = \mu(\{x : \text{out}_o(x) < d/2\}) > 0$, and

$$d/2 = \frac{1}{2} \int_x \text{deg}(x) d\mu = \int_X \text{out}(x) d\mu \leq (d/2)(1 - a) + (d/2 - 1)a < d/2$$

□

Grebík has show that every even-degree Borel graph admits an approximate balanced orientation [Gre22].

2.1.2 Examples and applications

The most widely studied class of Borel graphs are the Schreier graphs associated to actions of finitely generated groups. Recall that we are allowed to take inverses from a generating set to generate a group, and we do not allow the identity to be an element of our generating sets. So, for any free action a , $\text{Sch}(a, E)$ is regular with degree $|E \cup E^{-1}|$. In particular, if E contains no involutions, $\text{Sch}(a, E)$ is $2|E|$ -regular. So, if a is a free p.m.p. action, then $\mathfrak{o}_\mu(\text{Sch}(a, E)) \geq \frac{1}{2}|E \cup E^{-1}|$ by Proposition 2.1.6. Clearly, $\mathfrak{o}_B(\text{Sch}(a, E)) \leq |E|$. We then have

Proposition 2.1.12. If E does not contain γ and γ^{-1} for any γ , and a is a p.m.p action of G , then $\mathfrak{o}_\mu(\text{Sch}(a, E)) = \mathfrak{o}_B(\text{Sch}(a, E)) = |E|$.

The situation is much more interesting for groups with 2-torsion. We equip \mathbb{N}^Γ with any non-atomic product measure μ .

Proposition 2.1.13. Let $\Gamma = C_2^{*2} := \langle a, b : a^2 = b^2 = 1 \rangle$ and $E = \{a, b\}$. Then,

$$\mathfrak{o}_\mu(\mathcal{S}(\Gamma, E)) = 2.$$

In particular, there are p.m.p. graphs with $\mathfrak{o}_\mu(G) \neq \lceil \alpha_\mu(G) \rceil$.

Proof. Each component of $\mathcal{S}(\Gamma, E)$ is an infinite path. Suppose toward contradiction that o is a measurable 1-orientation, so o assigns a direction to each component of $\mathcal{S}(\Gamma, E)$. Consider

$A = \{x : (x, ax) \in o\}$. The group element ab acts ergodically and A is invariant under this action. Thus almost every point is in A or almost no point is in A . In either case, some edge must be missed by o , which is a contradiction. \square

Generalizing this construction, Bencs, Hrušková, and Tóth have found $2d$ -regular graphs with no measurable d -orientation, showing that for general graphs Theorem 2.1.7 is optimal [BsT21]. And more recently, Kun has given acyclic examples.

Theorem 2.1.14 (ess [Kun22]). For any d , there is a $2d$ -regular acyclic p.m.p. graph with no balanced orientation.

Proof. Kun constructs, for every d , a d -regular graph with no nonzero bounded circulation, and of course a balanced orientation gives a circulation by sending 1 unit along every edge in the direction of the orientation. \square

Shift graphs of nonamenable groups are expansive (see e.g. [LN11, Section 3]). So by Theorem 2.1.10, $\mathfrak{o}_\mu(\mathcal{S}(\Gamma, E)) = \lceil \frac{n}{2} \rceil$ when $\Gamma = C_2^{*n}$ with standard generating set E . Using a determinacy argument, we can show that these graphs have strictly larger Borel orientation numbers.

Theorem 2.1.15. For $\Gamma = C_2^{*n} = \langle a_1, \dots, a_n : a_i^2 = 1 \rangle$ and $E = \{a_1, \dots, a_n\}$,

$$\mathfrak{o}_B(\mathcal{S}(\Gamma, E)) = n$$

Proof. It suffices to consider the case when n is even. We proceed by induction. The base case is Proposition 2.1.13.

For the induction step, we use Marks's determinacy argument, Theorem 1.4.34. Suppose the shift graph for $C_2^{*(n-2)}$ does not admit an $(n-3)$ -orientation, and suppose o is an orientation of $\mathcal{S}(\Gamma, E)$ with $\text{out}_o(x) < n$ for all x . Set $H = \langle a_1, a_2 \rangle$ and $K = \langle a_3, \dots, a_n \rangle$, and let

$$A = \{x : (x, a_1x) \notin o \text{ or } (x, a_2x) \notin o\}.$$

By the lemma, there is an equivariant embedding either of $sS(H, \{a_1, a_2\})$ into A or of $\mathcal{S}(K, E \setminus \{a_1, a_2\})$ into $X \setminus A$.

In the first case, if f is the embedding, let \tilde{o} be the pullback orientation on $\mathcal{S}(H, E)$, i.e. for $i = 1, 2$

$$(x, a_i x) \in \tilde{o} \Leftrightarrow (f(x), a_i f(x)) \in o.$$

By the definition of A , \tilde{o} is a 1-orientation of $\mathcal{S}(H, E)$, contradicting the base case.

In the second case, again suppose f is the embedding and let \tilde{o} be the pullback orientation of $\mathcal{S}(\Gamma, E \setminus \{a_1, a_2\})$. Then for all x ,

$$(f(x), a_1 f(x)), (f(x), a_2 f(x)) \in o$$

so

$$|\tilde{o}_x| = |\{a_i : (f(x), a_i f(x)) \in o, i \neq 1, 2\}| = |o_{f(x)}| - 2 < n - 2.$$

But then we have an $n - 3$ orientation of $\mathcal{S}(K, E \setminus \{a_1, a_2\})$, which contradicts the induction hypothesis. □

As pointed out by Bernshteyn, this implies that the locality bounds on the Borel Lovász local lemma, Theorem 1.1.8 are sharp. Note that for the instance of the local lemma defining sinkless orientations of the d -regular tree, $\text{vdeg}(G, \varphi) = 2$, $\text{ord}(G, \varphi) = d$, and $p(G, \varphi) = 1/2^d$ (see Definition 1.1.5). So, Theorem 1.1.8 cannot be improved to $p \cdot \text{vdeg}^{\text{ord}} \leq 1$. Indeed this shows many known definable versions of the local lemma are sharp. Interestingly, lower bounds on distributed algorithms for the Lovász local lemma were first established using sinkless orientations. See [Ber21] for further discussion.

Another nice class of graphs generalizes the Hadwiger–Nelson graph.

Definition 2.1.16. For a Polish group G and $E \subseteq G$ Borel, the associated generalized distance graph is

$$D(G, E) := \{\{x, y\} : xy^{-1} \in E\}.$$

When E is countable, this is just the Schreier graph of $\langle E \rangle$ acting on G by translation. It turns out, when E is uncountable, $\mathfrak{o}_B(D(G, E)) = |\mathbb{R}|$.

Theorem 2.1.17. For any Polish group G and $E \subseteq G$ Borel, $\mathfrak{o}_B(D(G, E)) \leq \aleph_0$ if and only if E is countable.

Proof. If E is countable, then $D(G, E)$ is locally countable, and $\mathfrak{o}_B(D(G, E)) \leq \aleph_0$. So, suppose E is uncountable and fix a sequence of Borel functions $\langle f_i : \mathbb{N} \rangle$ with $f_i \subseteq G$ for each i . We will show that $\bigcup_i f_i$ is not an orientation of G . We may assume E is symmetric, and by the perfect set theorem we may replace E with one of its uncountable closed perfect subsets.

Equipped with the subspace topology, $D(G, E) \subseteq G^2$ is homeomorphic to $G \times E$. There are two natural identifications; a point $(x, e) \in G \times E$ can map to either (x, ex) or (ex, x) . We can translate between the two via a self-homeomorphism of $G \times E$, $(x, e) \mapsto (ex, e^{-1})$.

Define $\tilde{f}_i : G \rightarrow E$ by $\tilde{f}_i(g) = f_i(g)g^{-1}$. Then \tilde{f}_i is Borel and so has a meager graph in $G \times E$. That is,

$$\{(x, e) \in G \times E : (x, e) \neq (x, \tilde{f}_i(x))\} = \{(x, e) : (x, ex) \neq (x, f_i(x))\}$$

is comeager. Symmetrically, $\{(x, e) : (x, ex) \neq (f_i(ex), ex)\} = \{(x, e) : (ex, e^{-1}) \notin \tilde{f}_i\}$ is comeager. So, for a comeager set of (x, e) , then $(x, ex), (ex, x) \notin \bigcup_i f_i$. \square

For $G = \mathbb{R}^n$ and $E = S^{n-1}$, we get the following.

Corollary 2.1.18. The unit distance graph in \mathbb{R}^n ($n \geq 2$) does not have a countable Borel orientation.

This differs from the classical case. In ZFC, the unit distance graph on \mathbb{R}^n always admits a countable orientation, independent of the size of the continuum. See, for example, [AZ18, Theorem 6.2].

Note that the orientation and sidewalk covering numbers can be arbitrarily far apart in the Borel setting.

Theorem 2.1.19. For every $n \in \mathbb{N} \cup \{\aleph_0, |\mathbb{R}|\}$, there is an acyclic Borel graph (in particular a sidewalk) with $\mathfrak{o}_B(G) = n$. Further, if $n \leq \aleph_0$, G can be taken to be locally countable.

Proof. For $n \leq \aleph_0$, F_n , the free group with n generators, is torsion free. So if E is the usual set of generators, $\mathfrak{o}_B(\mathcal{S}(F_n, E)) = n$. And the Schreier graphs of F_n actions are locally countable.

For $n = |\mathbb{R}|$, label the standard generators for F_{2^n} as $E_n = \{a_\sigma : \sigma \in 2^n\}$ and let $g_n : F_{2^n} \rightarrow F_{2^{n-1}}$ be the homomorphism determined by $g_n(a_\sigma) = a_{\sigma \upharpoonright (n-1)}$. Define

$$\Gamma = \lim_{\leftarrow} F_{2^n} = \{f \in \prod_n F_{2^n} : (\forall n) g_n(f(n)) = f(n-1)\}$$

$$E = \{f \in \Gamma : (\forall i) f(i) \in E_i\}.$$

Then E is closed in Γ and uncountable, so by Theorem 2.1.17, $\mathfrak{o}_B(G(\Gamma, E)) > \aleph_0$. Also, since every finite subset of E freely generates a free group, $D(\Gamma, E)$ is acyclic. \square

We end this section by placing some orientation problems on the diagram in Figure 1.1.1.

Proposition 2.1.20. The LCL BO is in the class FIID, but not MEASURE. The LCL SO is in the class MEASURE, but not BOREL.

Proof. Since $\mathcal{S}(T_n)$ is expansive when $n > 2$, trees have FIID balanced orientations by Theorem 2.1.10. The fact that BO \notin MEASURE is Theorem 2.1.14.

For sinkless orientations, Theorem 2.1.7 implies SO can be solved measurably, and Theorem 2.1.15 says SO is not in the class Borel. \square

Note this proposition leaves it open if φ_{so} is in the class RLOCAL($O(\log \log(n))$). This is indeed the case, see [BCG21b].

2.1.3 Complexity

We end this section with some metamathematical considerations. First we prove effectivization for 1- and \aleph_0 -orientations.

Theorem 2.1.21. If G is Δ_1^1 , and $\mathfrak{o}_B(G) \leq \aleph_0$, then $\mathfrak{o}_{\Delta_1^1}(G) \leq \aleph_0$.

Proof. Note that $G \subseteq \mathcal{N}^2$ admits a countable orientation if and only if $G \subseteq \bigcup_i (f_i \cup f_i^{-1})$, where each f_i is a Borel partial function. A set $f \subseteq \mathcal{N}^2$ is a partial function when

$$(\forall(x, y), (a, b) \in f) \neg(x = a \text{ and } y \neq b).$$

This is an independence property, so by Theorem 1.4.42 if G is Δ_1^1 and admits a cover Borel by partial functions, G admits a cover by Δ_1^1 partial functions. \square

Corollary 2.1.22. The set of \aleph_0 -orientable Borel graphs is Π_1^1 -complete in the codes.

See Appendix A for some details on coding. This argument does not yield a dichotomy theorem, but one is given in [Tho21]. We can also prove effectivization for 1-orientations, refining Hjorth and Miller's analysis of end selection [HM09].

Theorem 2.1.23. If G is Δ_1^1 and admits a Borel 1-orientation, then G admits a Δ_1^1 1-orientation.

Proof. First note that G must be a sidewalk. Since cycles are unique when they exist in the components of G , the acyclic part of G is Δ_1^1 . And, G admits a Δ_1^1 selector on the components with a cycle, so G admits a 1-orientation if and only if the acyclic part of G admits a 1-orientation. So, we may restrict our attention to acyclic graphs.

For edges $e = (x, y)$ and $e' = (a, b)$ in G , say that e points at e' , or ePe' , if x, a are in the same component and y is on the unique simple path from x to a . Say e, e' are inconsistent, or eIe' , if x, a are in the same component, e does not point at e' , and e' does not point at

e . And, say that e is below e' , or $e \prec e'$ if ePe' but not $e'Pe$. Note that these are all Δ_1^1 relations, .

We show that G admits a (Borel or Δ_1^1) 1-orientation if and only if $G \subseteq \bigcup_i (A_i \cup A_i^{-1})$ where each A_i is I -independent and closed downwards under \prec , and does not contain (x, y) and (y, x) for any x, y . Then Theorem 1.4.42 gives the result.

Suppose o is a 1-orientation and eIe' . Say the simple path from x to a contains $n + 1$ vertices (including x and a). If both e and e' are in o , then there are at most $n - 1$ edges on this path oriented out of vertices on the path, but there are n edges total. so o must be I -independent. If $e \prec e'$ and $e' \in o$, then the reverse of e , (y, x) is inconsistent with e' , so $e \in o$. Thus, o is closed downwards. And, since o is an orientation, (x, y) and (y, x) are not both in o for any x, y .

Conversely, suppose that $G \subseteq \bigcup_i A_i$ with each A_i closed downwards and independent and does not contain any edge and its reverse. Define B_i inductively as follows:

- $B_0 = A_0$
- $B_{i+1} = B_i \cup (A_{i+1} \setminus B_i^{-1})$.

Then, the B_i are increasing and independent, and their union contains exactly one of (x, y) or (y, x) for every edge (x, y) . By independence, we cannot have (x, y) and $(x, y') \in B_i$ for any i unless $y = y'$, so $\bigcup_i B_i$ is a 1-orientation of G . \square

Hyperfinite Borel equivalence relation are equivalent to countable Borel relations admitting a 1-orientable graphing. (By a graphing of an equivalence relation E we mean a graph G with $E_G = E$.) Determining the complexity of such relations is a longstanding open problem. That problem is still out of reach, but we can settle the complexity of general relations admitting a 1-orientable graphing. The following proposition says it is enough to give a lower bound on the complexity of equivalence relations admitting a Borel selector.

Proposition 2.1.24. For E smooth (but not necessarily countable), the following are equivalent:

1. E admits a 1-orientable graphing
2. E is treeable
3. E admits a Borel selector

Proof. (2) and (3) are equivalent by a result of Hjorth [Hjo08]. To see that (3) implies (1), note that if f is a Borel selector for E , then f generates a graphing of E .

Now we show (1) implies (2). Suppose that $\mathfrak{o}_B(E) \leq 1$ and fix an graphing G witnessing this. Since G is a sidewalk, each component is either a tree or contains a unique cycle. Tossing out the least edge in each cycle (relative to some Borel linear ordering of the underlying space) yields a treeing of E . □

For X Polish, let $F(X)$ be the Effros Borel space of closed subsets of X .

Theorem 2.1.25. The set

$$\mathbf{Sel} := \{E \in F(\mathcal{N}^2) : E \text{ is an equivalence relation with a Borel selector}\}$$

is Σ_2^1 complete.

Proof. We prove this in two steps. Define

$$\mathbf{Uni} := \{R \in F(\mathcal{N}^2) : R \text{ admits a Borel uniformization}\}.$$

We will show

$$\mathbf{FBU} \leq_B \mathbf{Uni} \leq_B \mathbf{Sel}$$

where \mathbf{FBU} is the set of relations with full domain admitting Borel uniformization. Then, by a theorem of Adams and Kechris [AK00], \mathbf{Sel} is Σ_2^1 complete.

FBU \leq_B **Uni**: We want to take a relation R and extend it to a relation R' with full domain in such a way that the R' cannot be uniformized over all the points added to the domain. If R has cofinite domain this cannot be done, so we will replace R by $\mathcal{N} \times R$, and then add noise to extend it to a full domain relation.

Let $N \subseteq \mathcal{N}^3$ be such that $N_{(x,y)} \subseteq \mathcal{N} \setminus \Delta_1^1(x,y)$. Note that if f is a $\Delta_1^1(p)$ function whose graph is contained in N , then $\text{dom}(f) \cap \{p\} \times \mathcal{N} = \emptyset$.

Given $R \subseteq \mathcal{N}^2$ closed, let $R' = (\mathcal{N} \times R) \sqcup N$. If R admits full Borel uniformization, say via f , then so does R' , via $f'(x,y) = f(y)$. If R' admits Borel uniformization, say via $f \in \Delta_1^1(p)$, then for any x ,

$$f(p,x) \in R'_{(p,x)} \cap \Delta_1^1(p,x) \subseteq R'_{(p,x)} \setminus N_{(p,x)} = R_x.$$

So R admits full Borel uniformization via $f'(x) = f(p,x)$.

Identifying \mathcal{N}^2 with \mathcal{N} via a Borel isomorphism as usual, the map $R \mapsto R'$ is a reduction from **FBU** to **Uni**.

Uni \leq_B **Sel**: If $R \subseteq \mathcal{N}^2$ is closed, define

$$(x,y)E_R(x',y') :\Leftrightarrow x = x' \wedge [(x,y), (x,y') \in R \vee y = y'].$$

Then, E_R is a closed equivalence relation.

If R admits Borel uniformization, say via f , then E_R has a Borel selector g defined by

$$g(x,y) = \begin{cases} (x, f(x)) & (x,y) \in R \\ (x,y) & \text{else} \end{cases}$$

If E_R admits a selector, g , then R admits uniformization f via $f(x) = y :\Leftrightarrow g(x,y) = (x,y) \wedge (x,y) \in R$.

Again, \mathcal{N}^4 and \mathcal{N}^2 can be identified, and the map $R \mapsto E_R$ is Borel. So we have a reduction as claimed. \square

Note that, by Harrington–Kechris–Louveau [HKL90], the set of smooth equivalence relations is $\mathbf{\Pi}_1^1$ in the codes. So the preceding theorem gives a strong reason why admitting a selector is not the same as being smooth.

2.2 Cayley diagrams

The second class of LCLs we’ll study are Cayley diagrams, a kind of labelling problem that encodes a free action of a given finitely generated group. In the measurable setting, Cayley diagrams turn out to be intimately tied up with the problem of lifting Γ -factor maps to $\text{Aut}(G)$ -factor maps when $G = \text{Cay}(\Gamma, E)$. And special cases of Cayley diagrams include edge colorings and Schreier decorations for trees as in [Tot21].

To start with, let’s define Cayley diagrams:

Definition 2.2.1. For (Γ, E) a marked group with E finite, let φ_n^Γ be the set of edge labellings of $B_n(e)$ in $\text{Cay}(\Gamma, E)$ by elements of E satisfying the following:

- (i) for every vertex x and every $\gamma \in E$ there is at most one neighbour with $d(x, y) = \gamma$
- (ii) for every vertex x of degree $|E|$ and every $\gamma \in E$, there is some neighbour with $d(x, y) = \gamma$.
- (iii) for any path (x_0, \dots, x_n) in H , $d(x_{n-1}, x_n) \dots d(x_0, x_1) = e$ if and only if $x_0 = x_n$.

(Note that φ_n^Γ depends on E , so this is a slight abuse of notation).

A **Cayley diagram** for a is an edge labelling which is a φ_n^Γ -decoration for all n . A measurable graph H admits an **approximate Cayley diagram** if $\rho_n^\Gamma(H) = 0$ for all n .

For example, a Cayley diagram for C_2^{*n} with the usual generators is an n -coloring of an n -regular acyclic graph. And, a Cayley diagram for \mathbb{Z} with the usual generator is a balanced orientation of the 2-regular acyclic graph.

Note that a graph H admits a (not necessarily measurable) Cayley diagram if and only if every component of H is isomorphic to $\text{Cay}(\Gamma, E)$. And, any Schreier graph of a free Γ action admits a Cayley diagram given by $d(x, \gamma \cdot x) = \gamma$. (In fact, every Cayley diagram comes from such an action, see Proposition 2.2.4). In particular, $\text{Cay}(\Gamma, E)$ admits a Γ -FIID Cayley diagram for any marked group (Γ, E) . The question of which Cayley graphs admit an $\text{Aut}(G)$ -FIID Cayley diagram is much more complicated.

Our main theorems about Cayley diagrams are in the measurable and FIID settings. We show that lifting a Γ -FIID Cayley diagram to an $\text{Aut}(G)$ -FIID Cayley diagram is essentially equivalent to lifting any FIID solution to a locally checkable labelling problem³. And, we show that lifting approximate Cayley diagrams is equivalent to lifting approximate solutions to LCLs (see Theorems 2.2.7 and 2.2.8).

We also establish some general results on when this is possible. Amenable groups and trees always admit $\text{Aut}(G)$ -FIID approximate Cayley diagrams, and ergodicity considerations can rule out FIID Cayley diagrams for many groups. These general results let us transport some deep work on $\text{Aut}(T_n)$ -FIID combinatorics to the study of free groups. In particular, we get an asymptotic calculation of the measurable independence number of the shift graph of F_n , giving a partial answer to [KM04, Problem 5.57] (See Prop 2.2.21).

And we have a number of examples hinting at a rich structure for these Cayley diagram problems. In particular we can build an example of a nilpotent group with an $\text{Aut}(G)$ -FIID Cayley diagram, and we show that any such group must have torsion. This construction relies on Morris, Morris, and Verret's classification of automorphisms for nilpotent Cayley graphs. We show that (for a natural choice of generators) the lamplighter group does not admit an $\text{Aut}(G)$ -FIID Cayley diagram, but that a finite extension of it does. And lastly, we build a nonamenable group without even an approximate $\text{Aut}(G)$ -FIID Cayley diagram.

³More exactly, G admits an $\text{Aut}(G)$ -FIID Cayley diagram if and only if every Γ -FIID solution to a countable sequence of LCLs lifts to an $\text{Aut}(G)$ -FIID solution. This is equivalent to lifting solutions to single problems if, for instance, Γ is finitely presented

Our construction is sharp enough to answer a question of Weilacher about coloring graphs with isomorphic Cayley graphs (See Theorem 2.2.24).

Returning the Borel setting, we compute the complexity of the set of Schreier graphs of free Borel action of \mathbb{Z}^n . These graphs in fact admit a finite basis theorem generalizing Miller's work on directable forests of lines, though the size of the basis we obtain is super-exponential in n .

Before we proceed, we will want a few conventions and pieces of notation in place. Throughout, (Γ, E) will be a marked group with E finite, and G will be $\text{Cay}(\Gamma, E)$. Recall that the group Γ acts naturally by automorphisms on G by

$$\gamma \cdot x = x\gamma^{-1}.$$

To belabour the point, our Cayley graphs are defined as the Schreier graph of left multiplication, so to get graph automorphisms we have our group act by right multiplication.

Definition 2.2.2. We write $\langle \gamma \rangle$ for the image of γ under the natural embedding of Γ into $\text{Aut}(G)$ described above. We write $\text{Aut}_e(G)$ for the group of “rotations” in $\text{Aut}(G)$:

$$\text{Aut}_e(G) = \{r \in \text{Aut}(G) : r(e) = e\}.$$

We will only write $\langle \gamma \rangle$ when we want to make the distinction between group elements and automorphisms clear. And we often abbreviate $\text{Aut}_e(G)$ as R .

Note that every element of $\text{Aut}(G)$ factorizes uniquely as $\langle \gamma \rangle r$ or $r' \langle \gamma' \rangle$ where $\gamma, \gamma' \in \Gamma$ and $r, r' \in \text{Aut}_e(G)$. These two factorizations are related as follows:

$$r \langle \gamma \rangle = \langle r(\gamma^{-1}) \rangle^{-1} (\langle r(\gamma^{-1}) \rangle r \langle \gamma \rangle) \tag{1.1}$$

where $(\langle r^{-1}(\gamma^{-1}) \rangle r \langle \gamma \rangle) \in \text{Aut}_e(G)$. We refer to the identity (1.1) as the commutation relation.

Many of our questions trivialize when $\text{Aut}_e(G)$ is just the identity. Indeed, in this case, there is a unique Cayley diagram on G , so the action of $\text{Aut}(G)$ on the space of Cayley




Figure 2.1: A GRR for $D_\infty = \langle r, s : rsrs = s^2 = e \rangle$ induced by $\{r, s, sr, r^2s\}$

diagrams is already (trivially) FIID (see Proposition 2.2.4). In this case G is a so-called **graphic regular representation**, or GRR, of Γ . A GRR for the infinite dihedral group is shown in Figure 2.1. We will often state theorems only for the nontrivial case where (Γ, E) does not induce a GRR.

2.2.1 Measurable constructions

Throughout this section, we rely on a correspondence between actions, rotations, and Cayley diagrams.

Definition 2.2.3. Let CD be the space of Cayley diagrams on $\text{Cay}(\Gamma, E)$

Proposition 2.2.4. For any $d \in \text{CD}$ there is a unique rotation $r_d \in \text{Aut}_e(G)$ and action $\cdot_d : \Gamma \curvearrowright G$ so that for all $x \in \Gamma, \gamma \in E$ the following diagram commutes:

$$\begin{array}{ccc} x & \xrightarrow{r_d} & r_d(x) \\ \downarrow \gamma & & \downarrow d(x, \gamma x) \\ \gamma x & \xrightarrow{r_d} & \gamma \cdot_d r_d(x) \end{array}$$

meaning

$$r_d(\gamma x) = d(x, \gamma x)r_d(x) = \gamma \cdot_d r_d(x).$$

Similarly, any rotation or action gives a unique way to fill out this diagram.

Proof. Uniqueness is clear since $r_d(e) = e$ and the definition tells us how to compute $r_d(\gamma)$ for $\gamma = \gamma_n \dots \gamma_2 \gamma_1$ by filling out the diagram

$$e \rightarrow \gamma_1 \rightarrow \gamma_2 \gamma_1 \rightarrow \dots \rightarrow \gamma.$$

For existence, it suffices to note that item (iii) in the definition of Cayley diagram tells us that any two expressions for γ in terms of generators gives the same value of $r_d(\gamma)$ in this computation. \square

It will be useful to know how this translation interacts with automorphisms.

Proposition 2.2.5. Let $\text{Aut}_e(G)$ act on CD by shifting indices. For $r \in \text{Aut}_e(G)$, $\gamma \in \Gamma$ and $d \in \text{CD}$, we have

$$r_{r \cdot d} = r_d \circ r^{-1}, \text{ and } r_{\langle \gamma \rangle \cdot d} = \langle r_d(\gamma) \rangle r_d \langle \gamma \rangle^{-1}$$

Proof. Let $h = \langle r_d(\gamma) \rangle r_d \langle \gamma \rangle^{-1}$. We have that, for any $x \in \Gamma$, $\delta \in E$, $(r \cdot d)(x, \delta x) = d(r^{-1}(x), r^{-1}(\delta x))$ and $(\langle \gamma \rangle \cdot d)(x, \delta \gamma) = d(\langle \gamma \rangle^{-1}(x), \langle \gamma \rangle^{-1}(\delta x)) = d(x\gamma, \delta x\gamma)$. And, $h(x) = r_d(x\gamma)r_d(\gamma)^{-1}$. So the following diagrams commute:

$$\begin{array}{ccc} x & \xrightarrow{r_d r^{-1}} & (r_d r^{-1})(x) \\ \downarrow \delta & & \downarrow (r \cdot d)(x, \delta x) \\ \delta x & \xrightarrow{r_d r^{-1}} & (r_d r^{-1})(\delta x) \end{array} \quad \begin{array}{ccc} x & \xrightarrow{h} & r_d(x\gamma)r_d(\gamma)^{-1} \\ \downarrow \delta & & \downarrow (\langle \gamma \rangle \cdot d)(x, \delta x) \\ \delta x & \xrightarrow{h} & r_d(\delta x\gamma)r_d(\gamma)^{-1} \end{array}$$

The proposition then follows by the uniqueness of r_d . \square

Corollary 2.2.6. The action of $\text{Aut}_e(G)$ on CD is free and transitive. In particular, CD admits a unique $\text{Aut}(G)$ -invariant measure.

Proof. From the computation of $r_{r \cdot d}$ we can see that, for any $d_0 \in \text{CD}$, the map $r \mapsto r \cdot d_0$ has inverse $d \mapsto r_d^{-1} r_{d_0}$. Since the group $R := \text{Aut}_e(G)$ is compact it has a unique 2-sided invariant Haar measure, h . And since the diagram $d_0(x, \gamma x) = \gamma$ is invariant under the action of Γ , we get an invariant measure

$$\mu(A) = \int_{r \in R} \mathbf{1}_{r \cdot d_0 \in A} dh.$$

Conversely, any invariant measure μ on CD gives an invariant measure \tilde{h} on R by

$$\tilde{h}(A) = \mu(A \cdot d_0)$$

which must agree with h . \square

Note that this last corollary means that checking if there is an $\text{Aut}(G)$ -FIID Cayley diagram amounts to checking if this unique measure μ is $\text{Aut}(G)$ -FIID. This puts the question squarely in the realm of Ornstein theory, i.e. the business of determining which measures are FIID.

Our first theorem on Cayley diagrams is the lifting theorem for exact solutions to LCLs. This is essentially folklore (see the comments after [Lyo16, Question 2.4]), but we record a proof here for completeness. The basic idea is to once again average over the action of $\text{Aut}_e(G)$, this time using our Cayley diagram to choose a random rotation and resample our variables.

Theorem 2.2.7. Suppose that G admits an $\text{Aut}(G)$ -FIID Cayley diagram. Then for any LCL φ , G admits an $\text{Aut}(G)$ -FIID φ -decoration if and only if G admits a Γ -f.i.i.d φ -decoration.

Proof. For clarity we will distinguish between $\gamma \in \Gamma$ and its canonical image in $\text{Aut}(G)$, $\langle \gamma \rangle$.

Let F be an $\text{Aut}(G)$ -factor map from $[0, 1]^\Gamma$ to CD, and let Φ be a Γ -factor map giving a φ -decoration. We can associate to $x \in [0, 1]^\Gamma$ an element of $\text{Aut}_e(G)$ given by $h_x := r_{F(x)}$, where r_d is as in Proposition 2.2.4. From the previous proposition $h_{r \cdot x} = h_x r^{-1}$ and $h_{\gamma \cdot x} = \langle r_d(\gamma) \rangle r_d \langle \gamma \rangle^{-1}$

We will build a factor map out of $([0, 1]^\Gamma)^\Gamma$ (this is of course isomorphic to $[0, 1]^\Gamma$). Set

$$\tilde{\Phi}(x, y) = h_x^{-1} \cdot \Phi(h_x \cdot y).$$

With probability 1, $\tilde{\Phi}$ yields a φ -decoration. We just need to show that $\tilde{\Phi}$ is $\text{Aut}(G)$ -equivariant. We check against automorphisms $\langle \gamma \rangle$ and $r \in \text{Aut}_e(G)$ separately. If $\gamma \in \Gamma$,

$$\begin{aligned} \tilde{\Phi}(\langle \gamma \rangle \cdot (x, y)) &= h_{\langle \gamma \rangle \cdot x}^{-1} \cdot \Phi(h_{\langle \gamma \rangle \cdot x} \langle \gamma \rangle \cdot y) \\ &= (\langle h_x(\gamma) \rangle h_x \langle \gamma \rangle^{-1})^{-1} \cdot \Phi(\langle h_x(\gamma) \rangle h_x \langle \gamma \rangle^{-1} \langle \gamma \rangle \cdot y) \\ &= \langle \gamma \rangle h_x^{-1} \cdot \Phi(h_x \cdot y) \\ &= \langle \gamma \rangle \cdot \tilde{\Phi}(x, y) \end{aligned}$$

and

$$h_{r \cdot x}^{-1} \cdot \Phi(h_{r \cdot x} r \cdot y) = r h_x^{-1} \cdot \Phi(h_x \cdot y).$$

□

To get an approximate version of the theorem, we will use an ultraproduct argument.

Theorem 2.2.8. Suppose that G admits an approximate $\text{Aut}(G)$ -FIID Cayley diagram. Then for any φ , G admits an approximate $\text{Aut}(G)$ -FIID φ -decoration if and only if G admits an approximate Γ -FIID φ -decoration. In fact, $\rho^\varphi(\tilde{\mathcal{S}}(G)) = \rho^\varphi(\Gamma)$.

Proof. Suppose G admits an approximate $\text{Aut}(G)$ -FIID Cayley diagram. Let \mathcal{U} be a non-principal ultrafilter and let f_1, \dots, f_n be a marking of $\tilde{\mathcal{S}}(G)$. By Łoś's theorem for weak containment (Lemma 1.4.29), the ultrapower $\mathbb{S} := (\tilde{\mathcal{S}}(G), f_1, \dots, f_n)^\mathcal{U}$ has a measurable Cayley diagram. By the Lowenheim–Skolem theorem (Theorem 1.4.30), there is a separable elementary substructure of $\mathbb{S}' < \mathbb{S}$ containing (the partition coding) this Cayley diagram.

Since \mathbb{S}' is separable and admits a Cayley diagram, \mathbb{S}' is isomorphic to the metric structure associated to a p.m.p. action of Γ on a standard probability space. By Abert–Weiss (or Hatami–Lovász–Szegedi), 1.4.24, \mathbb{S}' weakly contains $\tilde{\mathcal{S}}(\Gamma)$. So $\rho^\varphi(\mathbb{S}') \geq \rho^\varphi(\tilde{\mathcal{S}}(\Gamma))$ for any φ . But \mathbb{S}' and $\tilde{\mathcal{S}}(G)$ are elementary equivalent, so $\rho^\varphi(\tilde{\mathcal{S}}(G)) \geq \rho^\varphi(\mathcal{S}(\Gamma, E))$. The reverse inequality is trivial. □

2.2.2 Examples and applications

Given the results of the previous subsection, the obvious question is: which Cayley graphs admit an $\text{Aut}(G)$ -FIID Cayley diagram? In this section we lay out some general theorems and several examples. We start with amenable groups.

Amenable groups always admit $\text{Aut}(G)$ -FIID approximate Cayley diagrams. In fact, if Γ is amenable we can show $\tilde{\mathcal{S}}(\Gamma, E)$ is hyperfinite, all but settling the approximate theory of amenable groups.

Theorem 2.2.9. If Γ is amenable and $G = \text{Cay}(\Gamma, E)$ for some finite generating set E , then for any LCL φ the following are equivalent

1. G admits a φ -decoration
2. G admits a Γ -FIID approximate φ -decoration
3. G admits an $\text{Aut}(G)$ -FIID approximate φ -decoration.

Proof. The implication from (3) to (2) is clear. To get (1) from (2) note that if (2) holds, then for any n there are decorations meeting φ at every point in $B_n(e)$. This implies (1) by compactness.

To get (3) from (1) we will show that $\tilde{\mathcal{S}}(\Gamma, E)$ is hyperfinite, i.e. that we have increasing finite equivalence relations E_n whose union is the connectedness relation for $\tilde{\mathcal{S}}(\Gamma, E)$. That done we have

$$\bigcup_n \{x : B_i(x) \subseteq [x]_{E_n}\} = \tilde{\mathcal{S}}(\Gamma, E),$$

so there is some n with $B_i(x)$ contained in the E_n -class of x for all but ϵ -many x . By (1) and Luzin–Novikov we can choose a φ -decoration on each of the E_n -classes to get a φ -decoration off a set of measure ϵ .

To get hyperfiniteness we will show that the connectedness relation of $\tilde{\mathcal{S}}(\Gamma, E)$ is amenable (as in [KM04, Chapter 9]). Since Γ is amenable, we have amenability measures ν_x for the action of Γ on $\text{Free}(\Gamma, E)$. Abbreviate $\text{Aut}_e(G)$ as R . Define

$$\nu_{R \cdot x} = \int_{r \in R} \nu_{r \cdot x} dh(r)$$

where h is the Haar measure on R .

This is a well defined measurable assignment of finitely additive measures. We check that this is invariant. Note that if $R \cdot x$ is connected to $R \cdot y$, then by Proposition 1.4.19 we can take y to be $\gamma \cdot x$ for some $\gamma \in \Gamma$. We have

$$r\langle\gamma\rangle \cdot x = \langle r(\gamma^{-1}) \rangle^{-1} \langle r(\gamma^{-1}) \rangle r\langle\gamma\rangle \cdot x$$

and $\langle r(\gamma^{-1}) \rangle r\gamma \in R$. The map $r \mapsto \langle r(\gamma^{-1}) \rangle r\langle \gamma \rangle$ is a measure preserving bijection with inverse $r' \mapsto \langle r'(\gamma) \rangle r'\langle \gamma^{-1} \rangle$. So, if we let $R_\delta = \{r \in R : r(\gamma^{-1}) = \delta\}$, then $R = \bigsqcup_\delta R_\delta = \bigsqcup_\delta \delta R_\delta \gamma$. So, using the Γ -invariance of ν and unimodularity of R we have

$$\begin{aligned}
\int_R \nu_{r\langle \gamma \rangle \cdot x} dh &= \sum_\delta \int_{R_\delta} \nu_{\delta^{-1} \delta r \gamma \cdot x} dh \\
&= \sum_\delta \int_{R_\delta} \nu_{\delta r \gamma} dh \\
&= \sum_\delta \int_{\delta R_\delta \gamma} \nu_{r \cdot x} d(\delta \cdot h \cdot \gamma) \\
&= \int_R \nu_{r \cdot x} dh.
\end{aligned}$$

□

Finding exact diagrams is trickier. A theorem of Ornstein and Weiss implies that it suffices to understand when the unique $\text{Aut}(G)$ -invariant measure on CD is Γ -FIID [OW87, Theorem 10]. The following proposition says we can instead look for FIID measures on $\text{Aut}_e(G)$. The extra group structure here will make our analysis easier.

Proposition 2.2.10. Let h be Haar measure on $\text{Aut}_e(G)$ and let Γ act on $\text{Aut}_e(G)$ by $\gamma \cdot r = \langle r(\gamma^{-1}) \rangle^{-1} r \langle \gamma^{-1} \rangle$. Then $\Gamma \curvearrowright (\text{Aut}_e(G), h)$ is isomorphic to $\Gamma \curvearrowright (\text{CD}, \mu)$, where μ is the unique $\text{Aut}(G)$ invariant measure on CD.

Proof. The map $d \mapsto r_d$ is an isomorphism. □

The following proposition is simple, but useful in ruling out many cases.

Proposition 2.2.11. Suppose there is a subgroup $H < \text{Aut}(G)$ where $\Gamma \subseteq H$ and

$$1 < [\text{Aut}_e(G) : H \cap \text{Aut}_e(G)] < \infty.$$

Then G does not admit an $\text{Aut}(G)$ -FIID Γ -Cayley diagram.

Proof. In this case, $\text{Aut}_e(G) \cap H$ is invariant under the action of Γ on $\text{Aut}_e(G)$, and has Haar measure $1/[\text{Aut}_e(G) : H \cap \text{Aut}_e(G)]$. So the action of Γ on $(\text{Aut}_e(G), h)$ is not ergodic, and by Proposition 2.2.10 neither is the action of Γ on (CD, μ) . But since the shift action of Γ is ergodic, every Γ -FIID measure is ergodic. \square

Corollary 2.2.12. If Γ is torsion-free and nilpotent and E is any set of generators which does not induce a GRR of Γ , then $\text{Cay}(\Gamma, E)$ does not admit an $\text{Aut}(G)$ -FIID Cayley diagram.

Proof. By [MMV16, Theorem 1.2], $\text{Aut}_e(G)$ is the set of group automorphism which preserve E . In particular, $\text{Aut}_e(G)$ is finite. Thus $\Gamma \cap \text{Aut}_e(G) = \{e\}$ is finite index and the previous proposition applies with $H = \Gamma$. \square

Indeed, this proposition applies to any CI_f group (in the sense of [Mor16]). The next theorem shows that the assumption that Γ is torsion-free is necessary for the corollary above. Consider the discrete Heisenberg group

$$H_3(\mathbb{Z}) := \left\{ \begin{pmatrix} 1 & a & b \\ 0 & 1 & c \\ 0 & 0 & 1 \end{pmatrix} : a, b, c \in \mathbb{Z} \right\}.$$

This is a torsion-free nilpotent group generated by the matrices

$$A = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad B = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}$$

Theorem 2.2.13. Define F a marking of $H_3(\mathbb{Z})$, a group Γ , and E a marking of Γ as follows:

$$F := \{A, B, A^2, A^2B, B^2\} \subseteq H_3(\mathbb{Z}), \quad \Gamma = C_2 \times H_3(\mathbb{Z}), \quad \text{and } E = C_2 \times (F \cup F^{-1}).$$

Then $G = \text{Cay}(\Gamma, E)$ admits an $\text{Aut}(G)$ -FIID Cayley diagram and $|\text{CD}| = |\text{Aut}_e(G)| > 1$.

Proof. Let $G' = \text{Cay}(H_3(\mathbb{Z}), F)$. By inspection of the unit ball in G' , this is a GRR of $H_3(\mathbb{Z})$. The Cayley graph G is obtained from G' by replacing each edge in G' with a copy K_4 .

If $r \in \text{Aut}_e(G)$ is an automorphism of G , then by [MMV16, Theorem 1.11] it descends to an automorphism of G' , which must be trivial. So r decides independently for each $m \in H_3(\mathbb{Z})$ whether or not to swap $(0, m)$ and $(1, m)$, i.e. $r(a, m) = (a + x(m), m)$ for some $x \in C_2^{G'}$ with $x(I) = 0$. This is clearly FIID, we check the details below. The action of Γ on $\text{Aut}_e(G)$ translates to $C_2^{G' \setminus \{I\}}$ as $((b, \ell) \cdot x)(m) = x(\ell) + x(\ell m)$.

Consider C_2^Γ equipped with the product of counting measure (this is the same as Haar measure, viewing C_2^Γ as a compact abelian group). Define $F : C_2^\Gamma \rightarrow C_2^{G' \setminus \{I\}}$ by

$$F(y)(m) = y(0, m) + y(1, m) + y(0, I) + y(1, I).$$

Then F is a surjective a group homomorphism so pushes Haar measure onto Haar measure. It suffices to show that F is Γ equivariant, where Γ acts on C_2^Γ by shifting and on $C_2^{G' \setminus \{I\}}$ as above:

$$\begin{aligned} F((b, \ell) \cdot y)(m) &= y(1, \ell^{-1}m) + y(0, \ell^{-1}m) + y(1, \ell^{-1}) + y(0, \ell^{-1}) \\ &= ((b, \ell) \cdot F(y))(m). \end{aligned}$$

□

For solvable groups, or more general group extensions, we can sometimes lift obstructing subgroups of Proposition 2.2.11 up from the Cayley graphs of a quotient or subgroup. To illustrate, we will work out an example with the lamplighter group.

Definition 2.2.14. The **lamplighter group** is the wreath product $L = C_2 \wr \mathbb{Z}$. More concretely,

$$L = \{(n, A) \in \mathbb{Z} \times \mathcal{P}(\mathbb{Z}) : A \text{ is finite}\}$$

with the product

$$(n, A) \cdot (m, B) = (n + m, A \Delta (n + B)).$$

The **Diestel–Leader generators** for L are

$$D := \{(-1, \emptyset), (1, \emptyset), (1, \{0\}), (-1, \{-1\})\}.$$




Figure 2.2: Two copies of T_3 and the Diestel–Leader graph

Let $T_3 = (V, F)$ be the three regular tree, and let $h : T_3 \rightarrow \mathbb{Z}$ be a height function so that every vertex v has exactly one neighbour w with $h(v) = h(w) + 1$ and two neighbors with $h(v) = h(w) - 1$. Then, the **Diestel–Leader graph** is the subgraph of $T_3 \times T_3$ with vertex set $\{(v, w) : h(v) + h(w) = 0\}$ and edges $\{((v, w), (v', w')) : (v, v'), (w, w') \in F\}$.

The Cayley graph $\text{Cay}(L, D)$ is isomorphic the Diestel–Leader graph [BNW06, Corollary 3.14]. It turns out that in this case, every automorphism of $\text{Cay}(L, D)$ descends to automorphism of $\text{Cay}(\mathbb{Z}, \{1, -1\})$. The subgroup we use to apply Proposition 2.2.11 in this case is the preimage of \mathbb{Z} under this quotient. Equivalently, we look at the set of automorphisms preserving the quasi-order

$$(n, A) \leq (m, B) :\Leftrightarrow n \leq m.$$

This order translates to the product of trees presentation of the Diestel–Leader graph as

$$(v, w) \leq (v', w') :\Leftrightarrow h(v) \leq h(v') \quad (\Leftrightarrow h(w) \geq h(w')).$$

We will work with this product of trees presentation throughout.

Proposition 2.2.15. The Cayley graph $G = \text{Cay}(L, D)$ does not admit an $\text{Aut}(G)$ -FIID Cayley diagram.

Proof. Let H be the group of order preserving automorphisms. Shifting by an element of L does not affect the order, so it suffices to show $H \cap \text{Aut}_e(G)$ has index 2.

For a vertex x of G let $N(x) = B_2(x) \setminus \{x\}$. Then, $N(x)$ has two components: the vertices above x and the vertices below x . Fix $r \in \text{Aut}_e(G)$. Since r preserves distances and fixes e , r restricts to an automorphism of $N(e)$. Induction on $B_n(e)$ shows that r is order preserving if and only if it does not swap the two components of $N(e)$. Define $s \in \text{Aut}_e(G)$ by $t(x, y) = (y, x)$. The automorphism t swaps the two components of $N(e)$, so either $r \in H$ or $tr \in H$. Thus $H \cap \text{Aut}_e(G)$ has index 2.

□

A similar argument applies to other Cayley graphs for the lamplighter group and other extensions of amenable groups. However, this obstruction we found for the lamplighter group is fragile in the sense that it can be circumvented by taking a finite extension.

Theorem 2.2.16. There is a finite extension of the lamplighter group Γ with a finite generating set E so that $G = \text{Cay}(\Gamma, E)$ has the Diestel–Leader graph as a quotient, $|\text{Aut}_e(G)| > 1$, and G admits an $\text{Aut}(G)$ -FIID Cayley diagram.

Proof. Let H be the group of order preserving automorphisms of $\text{Cay}(L, D)$. First we show that $H \cap \text{Aut}_e(\text{Cay}(L, D))$ equipped with Haar measure is H -FIID. Then we construct an appropriate extension of L , Γ and a generating set E so that $\text{Aut}_e(G) \cong H$.

For any $g \in H$ there are $g_0, g_1 \in \text{Aut}(T_3)$ so that $g(v, w) = (g_0(v), g_1(v))$ (see [BNW06, Proposition 2.7]). Further g_0 and g_1 must preserve the height function h . So, each of g_0 and g_1 is given by deciding independently for each vertex $v \in T_3$ whether or not to swap the two subtrees below v .

Translating this to CD via Proposition 2.2.10, the Cayley diagrams corresponding to elements of $H \cap \text{Aut}_e(\text{Cay}(L, D))$ are given as follows: for each $v, w \in T_3$ with $h(v) + h(w) = 1$, assign labels from $(1, \emptyset), (1, \{0\})$ to the 4 edges of the form $((v, x), (y, w))$ in Diestel–Leader




Figure 2.3: $B_1(x)$ in G

graph so that every pair of edges incident to a common vertex receives two different labels. (One set of such edges is illustrated in bold in Figure 2.2.2.) The H invariant measure makes these assignments independently and uniformly at random. This is clearly Γ -FIID.

Now we want to build a marked group whose automorphisms correspond to order preserving automorphisms of $\text{Cay}(L, D)$. Let Γ be the subgroup of $\text{Aut}(\text{Cay}(L, D))$ generated by the lamplighter group and the swap t . More concretely,

$$\Gamma := C_2 \ltimes L$$

where C_2 has generator t which acts on the lamplighter group by $t \cdot r = r^{-1}$ for any $r \in D$ (so in Γ , $tr = r^{-1}t$). Let $E = D \cup \{t, t \cdot (1, \emptyset), t \cdot (1, \{0\})\}$. Note that every group element in $E \setminus D$ is an involution, so E is symmetric.

The Cayley graph $G = \text{Cay}(\Gamma, E)$ is two copies of $\text{Cay}(L, D)$ with edges between the two copies corresponding to t and to ascending edges in $\text{Cay}(L, D)$. Any automorphism in $H \cap \text{Aut}_e(\text{Cay}(L, D))$ extends to an automorphism of G by $h(vt) = h(v)$.

Let \tilde{H} be the image of this embedding, and note that \tilde{H} is closed in $\text{Aut}_e(G)$. By inspection of $B_1(e)$, for any $r \in \text{Aut}_e(G)$ there is some $h \in \tilde{H}$ so that hr restricts to the identity on $B_1(e)$. By induction, there is a sequence $\langle h_i : i \in \mathbb{N} \rangle$ so that $h_i r$ restricts to the identity on $B_i(e)$, meaning $\lim_i h_i = r^{-1}$. Since \tilde{H} is closed, we have $r \in \tilde{H}$. The same argument that Haar measure on $H \cap \text{Aut}_e(\text{Cay}(L, D))$ is FIID shows that Haar measure on $\text{Aut}_G(G)$ is FIID. \square

Every positive example we gave above relied on making independent choices of elements

from finite groups. Perhaps this is always the case:

Problem 2.2.17. Is there a marked torsion-free (amenable) group (Γ, E) where $G = \text{Cay}(\Gamma, E)$ admits an $\text{Aut}(G)$ -FIID Cayley diagram?

Adams's very weakly Bernoulli characterizations of Γ -FIID systems implies that the set of marked amenable groups admitting an $\text{Aut}(G)$ -FIID Cayley diagram is Borel in any reasonable coding of groups [Ada92]. This may be the best one can say, but perhaps there is a more concrete characterization of such groups.

Problem 2.2.18. Is there a marked amenable group (Γ, E) where the action of Γ on CD is mixing, but where $G = \text{Cay}(\Gamma, E)$ does not admit an $\text{Aut}(G)$ -FIID Cayley diagram?

Of course matters are much more open for nonamenable groups. Ergodicity arguments like those in the previous section can again rule out FIID Cayley diagrams in many cases. And, there are simple examples of groups which admit $\text{Aut}(G)$ -FIID Cayley diagrams. Let C_n be the cyclic group of order n .

Proposition 2.2.19. If $\Gamma = C_2 * C_3 = \langle a, b : a^2 = b^3 = e \rangle$ and $E = \{a, b, b^{-1}\}$, then $G = \text{Cay}(\Gamma, E)$ admits an $\text{Aut}(G)$ -FIID Cayley diagram.

Proof. The edge set of G decomposes into single edges and 3-cycles, and the $\text{Aut}(G)$ -invariant measure on CD corresponds to an i.i.d. random orientation of each cycle. \square

However, at the moment there is nothing like a comprehensive Ornstein theory for non-amenable groups. Indeed it could be that the groups admitting an $\text{Aut}(G)$ -FIID Cayley diagram form a strictly analytic set.

Unlike in the amenable case, nonamenable groups have an intricate approximate theory. We can show that trees admit $\text{Aut}(G)$ -FIID approximate Cayley diagrams.

Theorem 2.2.20. For any (Γ, E) , if $G = \text{Cay}(\Gamma, E)$ is an n -regular tree, then G admits an $\text{Aut}(G)$ -FIID approximate Cayley diagram.

Proof. Recall that if the Cayley graph of Γ is a tree, then $\Gamma = C_2^{*n} * F_m$ for some $n, m \in \mathbb{N}$. We proceed by induction on n . We may assume we have $n + 1$ independent labels on each vertex of G .

If $n = 0$, then we want a Cayley diagram for a free group. If $m = 1$, then Γ is amenable, and we can apply Theorem 2.2.9. Otherwise let o be a balanced orientation of $\tilde{\mathcal{S}}(\Gamma, E)$. One exists by [Tho22, Theorem 2.8]. We want to label edges so that, away from a set of measure ϵ , every vertex has one in-edge and one out-edge with each label. This is equivalent to approximately edge coloring the graph $(\tilde{\mathcal{S}}(\Gamma, E) \times \{0, 1\}, \{(x, 0), (y, 1)\} : (x, y) \in o)$. This can be done by Tóth's approximate version of König's line coloring theorem [Tot21].

If $m = 0$ and $n = 1$ or 2 then again Γ is amenable, and so so we can apply Theorem 2.2.9.

Now suppose that $n + 2m > 3$. Then Γ is nonamenable, so by the Lyons–Nazarov matching theorem [LN11, Theorem 2.4] we can use the n^{th} independent variable on each vertex to produce an $\text{Aut}(G)$ -FIID perfect matching M . Throwing away the edges in M reduces us to looking at the $(n - 1)$ case, and we're done by induction. \square

This proof shows that every $2d$ -regular shift graph admits an approximate action of a free group. Grebik has recently shown that this can be extended to all $2d$ -regular p.m.p. graphs [Gre22], and Bencs, Tóth, and Hrušková have found measurable F_d actions on the shift graph of planar lattices [BsT21].

As pointed out by Bernshteyn, the above result along with with some deep results from probability and our approximate lifting lemma give us an asymptotic calculation of the independence number of $\mathcal{S}(F_n, E)$, where E is the usual set of generators.

Proposition 2.2.21. For a p.m.p. graph H , let $i(H)$ be the supremum of measures of independent sets in H , let $\chi_\mu(H)$ be the measurable chromatic number of H , and let $\chi_\epsilon(H)$

be the approximate chromatic number of H . Then,

$$(1 - o(n)) \frac{2n}{\log(n)} \leq 1/i(\mathcal{S}(F_n, E)) \leq \chi_\epsilon(\mathcal{S}(F_n, E)) \leq \chi_\mu(\mathcal{S}(F_n, E)) \leq (1 + o(n)) \frac{2n}{\log(n)}$$

Proof. The middle two inequalities are straightforward. The right hand inequality follows from (e.g.) [Ber19]. For the left-hand inequality, Rahman and Virag [RV17] showed that

$$i(\tilde{\mathcal{S}}(T_{2n})) \leq (1 + o(n)) \frac{\log(n)}{2n}.$$

And it follows from Theorems 2.2.8 and 2.2.20 that the same holds for $\mathcal{S}(F_n, E)$. \square

On the other hand, there are groups with no $\text{Aut}(G)$ -FIID Cayley diagram. In fact, the difference between $\text{Aut}(G)$ - and Γ -FIID combinatorics can be reflected in vertex colorings. Recall that the approximate chromatic number of a marked group (Γ, E) is the minimal n so that $\mathcal{S}(\Gamma, E)$ admits an approximate n -coloring. We will construct marked groups with isomorphic Cayley graphs but different approximate chromatic numbers, answering [Wei20, Problem 2].

Consider the groups

$$\Gamma = C_2^{*3} \times C_3 = \langle a_1, a_2, a_3, b : a_i^2 = [a_i, b] = b^3 = 1 \rangle$$

$$E = \{a_1, a_2, a_3, b\}$$

and

$$\Delta = \langle a_1, a_2, a_2, b : a_i^2 = b^3 = abab = e \rangle,$$

$$F = \{a_1, a_2, a_3, b\}.$$

Note that Δ is the semidirect product $C_2^{*3} \rtimes C_3$, where C_2 acts on C_3 by inversion, and $\text{Cay}(\Gamma, E) \cong \text{Cay}(\Delta, F)$.

Theorem 2.2.22. $\text{Cay}(\Gamma, E)$ admits a Γ -FIID 3-coloring.

Proof. As in [Wei20] it suffices to find a measurable independent set A which intersects each C_3 orbit in $[0, 1]^\Gamma$. Then we can have a 3-coloring of the shift graph given by $c(v) = \min\{i : b^i \cdot x \in A\}$. We construct such an A by a measurable augmentation argument.

Any independent set meets each C_3 orbit at most once. And if an independent set cannot be extended to meet a C_3 orbit then it must meet each of the three adjacent orbits. If we have any two such missing orbits in the same component, say at distance r in G/C_3 , we can form a $3r$ augmentation of A on the path between these orbits as follows.

Given an independent set A , say that x, v_1, \dots, v_n, y is a good path of length n for A if x and v_1 are neighbours, v_n and y are neighbours, $x, y \notin A$, $v_i \in A$, and $C_3 \cdot x, C_3 \cdot v_1, \dots, C_3 \cdot y$ is a simple path in G/C_3 .

We show that if $p = (x, v_1, \dots, v_n, y)$ is a good path A , then there is an independent set A' such that $A' \Delta A \subseteq C_3 \cdot p$, and A' meets the C_3 orbit of x and every v_i .

If x has a neighbour outside of $C_3 \cdot p$ as well, then by pigeonhole there is some point of $C_3 \cdot x$ with no neighbour in A , say x' . Set $A' = A \cup \{x'\}$. Otherwise, for each i , $C_3 \cdot v_i$ contains two points not adjacent to elements of A outside of p , say v_i, v'_i . Put x in A' and swap v_1 for v'_1 . If v'_1 is adjacent to v_2 , swap v_2 for v'_2 . Continue in this manner until we no longer need to swap. This process must terminate when we reach y .

So, if A admits no $3r$ -augmentation, it admits no good paths of length n , and the C_3 orbits which avoid A_n are spaced at least n vertices apart in G/C_3 . Thus the set of vertices whose orbits miss A has measure at most $3c^n$ for some $c < 1$. Thus there is a measurable independent set which meets every orbit and a Γ -FIID 3-coloring by Lemma 1.4.11

□

Theorem 2.2.23. $\text{Cay}(\Delta, F)$ does not admit a Δ -FIID approximate 3-coloring.

Proof. Suppose f is a measurable vertex labelling of the shift graph which meets the 3-coloring constraint on the C_3 -orbit of x and all of its neighbours. Again following [Wei20], let $c_f(x) = f(b \cdot x) - f(x) \in C_3$. Then

1. c_f takes values in $\{\pm 1\}$, otherwise $f(b \cdot x) = f(x)$
2. c_f is constant on C_3 orbits, otherwise we have $f(b^2 \cdot x) = f(x) + 1 - 1 = f(x)$ for some x .
3. c_f is different on adjacent orbits, otherwise for some j , $f(b^i \cdot x) = f(x) + i$, and $f(a_j b^i \cdot x) = f(b^{-i} a_j \cdot x) = f(a_j \cdot x) - i$, which is a contradiction since $m + i = n - i$ always has a solution mod 3.

Thus an approximate 3-coloring of $\mathcal{S}(\Delta, F)$ yields an approximate 2-coloring of $\mathcal{S}(C_2^{*3}, \{a_1, a_2, a_3\})$ as follows. Define

$$g : ([0, 1]^3)^{C_2^{*3}} \rightarrow [0, 1]^\Delta$$

by

$$g(x_1, x_2, x_3)(\gamma b^i) = x_i(\gamma)$$

for $\gamma \in C_2^{*3}$. Then if $\langle f_n : n \in \mathbb{N} \rangle$ is an approximate 3-coloring of $\mathcal{S}(\Delta, F)$, we have an approximate 2-coloring of $\mathcal{S}(C_2^{*3}, \{a_1, a_2, a_3\})$ given by $\langle c_{f_n} \circ g : n \in \mathbb{N} \rangle$. But then each color set is approximately invariant under the group elements of even length, contradicting strong ergodicity (as in [KT08]). \square

Corollary 2.2.24. $\text{Cay}(\Gamma, E)$ does not admit an $\text{Aut}(\text{Cay}(\Gamma, E))$ -FIID approximate Cayley diagram.

Typically, when we show that some nonamenable group does not admit a Γ -FIID φ decoration for some φ , we can also show that it does not a Γ -FIID approximate φ -decoration. The following question seems to require more delicate tools for analyzing FIID measures for nonamenable groups:

Problem 2.2.25. Is there a torsion free nonamenable marked group (Γ, E) where $G = \text{Cay}(\Gamma, E)$ admits an $\text{Aut}(G)$ -FIID approximate Cayley diagram but no Γ -FIID Cayley diagram?

Indeed, the following is still frustratingly open:

Problem 2.2.26. Does the n -regular tree admit an $\text{Aut}(G)$ -FIID Cayley diagram for any group?

2.2.3 Complexity

We end by returning to the Borel setting. We can compute the complexity of Schreier graphs of free Z^d actions and not necessarily free C_2^{*3} actions. It turns out that \mathbb{Z}^d -Cayley diagrams admit effectivization while the set of C_2^{*3} Schreier graphs is Σ_2^1 -complete.

Theorem 2.2.27. The set of Schreier graphs of C_2^{*3} is Σ_2^1 -complete in the codes.

Proof. This is simply the class of edge 3-colorable 3-regular graphs. The classical gadget reduction from 3-coloring to edge 3-coloring also works in the Borel setting. See Appendix A for details. \square

It is tempting to conjecture that every NP-complete LCL has a Σ_2^1 -complete Borel analogue. We will see that this is true for CSPs (see Theorem 3.4.3), though the case of LCLs seems to be tied up with deep problems in computer science about what kind of reductions can capture NP-completeness

Recall that for an edge (x, y) , we write $-e = (y, x)$.

Theorem 2.2.28. If G is Δ_1^1 and induced by a free Borel action of \mathbb{Z}^d , then G is induced by a free Δ_1^1 -action of \mathbb{Z}^d .

Proof. To ease notation, we prove this in the case $d = 2$. The general case is not conceptually harder.

All components of G isomorphic the square lattice. In particular G is locally countable.

Definition 2.2.29. A **straight line** is a directed path so that no other path of the same or shorter length has the same endpoints. A **rectangle** is a simple directed cycle that can be divided into 4 straight lines.

Two directed edges e, e' are **parallel** if they lie on some straight line together or if $(-e)$ and e' are on opposite sides of a rectangle. We write $e \parallel e'$. We say e, e' are **anti-parallel** if $(-e) \parallel e'$. And we say e, e' are **perpendicular** if they are in the same component and are neither parallel nor anti parallel. We write $e \perp e'$.

Note that all of these relations are Δ_1^1 .

Definition 2.2.30. Let $\{a, b, -a, -b\}$ be the usual generating set for \mathbb{Z}^2 . For a graph G , say that a partial function $f : G \rightarrow \{a, b, -a, -b\}$ is a **partial diagram** if for all e, e' ,

1. if $e \parallel e'$, then $f(e) = f(e')$
2. if $e \parallel (-e')$ then $f(e) = -f(e')$
3. if $e \perp e'$ then $f(e)$ and $f(e')$ have different letters (i.e. $f(e) \neq \pm f(e')$).

We first check that a Cayley diagram is the same as a partial diagram whose domain is all of G . If $a : \mathbb{Z}^2 \curvearrowright \mathcal{N}$ is an action generating G , then any straight line is of the form $\{n\gamma \cdot x : n \in \mathbb{Z}\}$ for some generator γ of \mathbb{Z}^2 . It follows that any Cayley diagram satisfies (1–3) above. Conversely, given a partial diagram f with total domain, note that for any x, y $f(x, y) = -f(y, x)$ so $(\gamma - \gamma) \cdot x = x$, and $(x, a \cdot x, (b+a) \cdot x, (-a+b+a) \cdot x, (-b-a+b+a) \cdot x)$ must be a rectangle, so $x = (-b - a + b + a) \cdot x$. This shoes that f is a Cayley diagram as \mathbb{Z}^2 has finite presentation $\langle a, b : ab = ba \rangle$.

Since being a partial diagram is an independence property, by Theorem 1.4.42 it suffices to show that any sequence of partial diagrams whose domains cover G can be patched together into a Cayley diagram.

Suppose we have two Δ_1^1 partial diagrams f and g . We show how to modify and glue

them together to get a partial diagram whose domain is the union of their domains. The result then follows by induction. Define h by

1. if $e \in \text{dom}(f)$, then $h(e) = f(e)$
2. if $e \in \text{dom}(g) \setminus [\text{dom}(f)]_G$, $h(e) = g(e)$
3. else if $e \in \text{dom}(g)$ and $e \parallel e'$ for some $e' \in f$ (or $e \parallel (-e')$), then $h(e) = f(e')$ (or $-f(e')$).
4. else if $e \perp e' \in \text{dom}(f)$, $h(e)$ has the sign of $g(e)$ and the opposite letter of $f(e')$.

It is straightforward to check that h is a partial diagram. □

In fact, one can find a finite basis theorem for Schreier graphs of \mathbb{Z}^n generalizing Miller's theorem on undirectable forests of lines. The paper [Tho21] gives the details for $n = 2$.

CHAPTER 3

Constraint Satisfaction Problems

3.1 Introduction

This chapter is about Borel versions of the following general problem: for a fixed finite relational structure \mathcal{D} , when does a structure \mathcal{X} in the same language admit a homomorphism into \mathcal{D} ? We call \mathcal{D} the template and \mathcal{X} an instance of \mathcal{D} , and we sometimes refer to a homomorphism as a solution to \mathcal{X} . For example:

1. If $\mathcal{D} = K_n$, then an instance of \mathcal{D} is a directed graph and a solution is a (vertex) n -coloring.
2. If \mathcal{D} has domain $\{0, 1\}$ and relations $\left\{ (x_1, \dots, x_k) : \bigvee_{i < j} \neg x_i \vee \bigvee_{i \geq j} x_i \right\}$ then an instance of \mathcal{D} is an instance of k SAT and a solution is a satisfying assignment.
3. If \mathcal{D} is a finite field \mathbb{F} equipped with one relation for each affine subspace, then an instance of \mathcal{D} is a system of linear equations, and a homomorphism in \mathcal{D} is a solution to the system.

Following example (2) above, computer scientists refer to these as constraint satisfaction problems (CSPs). We adopt this convention:

Definition 3.1.1. For \mathcal{D} a finite relational structure, $\text{CSP}(\mathcal{D})$ is the set of finite structures which admit a homomorphism into \mathcal{D} , and $\text{CSP}_B(\mathcal{D})$ is the set of codes for Borel structures which admit Borel homomorphisms into \mathcal{D} .

To be precise, we should specify presentations and codings for our structures. For finite structures, the details will not be important in this work, but see [FV98]. For Borel structures, any standard coding should work. Appendix A gives some details of a convenient coding.

Some of the typical first questions we ask about a Borel combinatorial problem are: Is a classical solution enough to guarantee a Borel solution? Is there some kind of dichotomy theorem for this problem? Is a Borel solution equivalent to a (lightface, effective) Δ_1^1 solution for Δ_1^1 instances? As an example consider the problem of countably coloring graphs. The G_0 graph of Kechris, Solecki, and Todorcevic is an instance of the countable coloring problem with a solution but no Borel solution. The G_0 dichotomy says that G_0 admits a homomorphism into any Borel graph with no Borel countable coloring. And, the proof of the G_0 dichotomy implies that any Δ_1^1 graph with a Borel countable coloring has a Δ_1^1 countable coloring [KST99]. A positive answer to any of our first pass questions implies an upper bound on the projective complexity of the problem. So, a Σ_2^1 -completeness result rules out all of these niceties. See Subsection 1.4.7 in the introduction for more discussion.

In computer science, similar first pass questions about a problem include: Is this problem in P? Can it be solved by linear relaxation or by constraint propagation? Is there a finite list of minimal instances with no solution? Remarkably, these problems have all been solved for CSPs.

Indeed, the class of CSPs was isolated by Feder and Vardi in the 1990s as a rich class of problems where complexity questions could be settled. Despite Ladner's theorem, which says that (assuming $P \neq NP$) there must be many intermediate classes between P and the NP-complete [Lad75], most natural problems seem to fall into one of the two extremes. As a partial explanation of this phenomena Feder and Vardi conjectured that all problems of the form $CSP(\mathcal{D})$ must be either in P or NP-complete [FV98].

Not too long after Feder and Vardi made their conjecture, Jeavons showed that this conjectured dichotomy must come down to a question about polymorphism algebras [Jea98].

Definition 3.1.2. A **polymorphism** of \mathcal{D} is a homomorphism from the categorical product \mathcal{D}^n to \mathcal{D} .

So a polymorphism takes in n solutions to an instance of \mathcal{D} and returns another solution. Several examples are given after Definition 3.2.1. A classical theorem of universal algebra says that polymorphism algebras ordered by containment are in Galois correspondence with structures under a notion of simulation called pp definability. And, Jeavons showed that pp definitions yield polynomial time reductions. (This has since been generalized greatly, see Theorem 3.2.6).

About a decade later, Bulatov, Jeavons, and Krokhin conjectured an algebraic dividing for polynomial time solvability¹ [BJK05]. After decades of work in computer science, combinatorics, and universal algebra, the conjectures of Feder–Vardi and Bulatov–Jeavons–Krokhin were confirmed independently by Bulatov and Zhuk in 2017:

Theorem 3.1.3 (CSP Dichotomy Theorem, [Bul17][Zhu17]). For a finite relational structure \mathcal{D} , $\text{CSP}(\mathcal{D})$ is polynomial time solvable if there is a polymorphism f of \mathcal{D} satisfying

$$(\forall a, e, r) f(r, a, r, e) = f(a, r, e, a).$$

And $\text{CSP}(\mathcal{D})$ is NP-complete otherwise.

In this chapter, we give partial algebraic answers to some of our basic questions from Borel combinatorics.

3.1.1 Results and conjectures

Here we lay out the main results of this chapter and some related open problems. To state these results we start with a number of definitions.

¹The form of the CSP dichotomy theorem given here is somewhat anachronistic. The original conjecture was in terms of Taylor operations. See the comments after Theorem 3.2.9

Definition 3.1.4. Say that $\text{CSP}_B(\mathcal{D})$ is **essentially classical** if any time a Borel instance of \mathcal{D} admits a solution it admits a Borel solution.

Say that $\text{CSP}_B(\mathcal{D})$ is **effectivizable** if anytime a Δ_1^1 instance admits a Borel solution it admits a Δ_1^1 solution.

These properties say that classical and effective tools respectively are sufficient to understand the Borel problem. Effectivizability plays a surprisingly important role in Borel combinatorics. As discussed in the introduction, most nontrivial complexity upper bounds come with an effectivization result.

In the finitary setting, similar classes of problems are closed downward under so-called pp constructions (see Definition 3.2.3). Since pp construction can be characterized by the polymorphism algebras of \mathcal{D} and \mathcal{E} (see Theorem 3.2.6), it follows abstractly that there is some algebraic characterization of these finitary classes. The same is nearly true in the Borel setting, but for technical reasons we need to assume equality is part of our structures.

Theorem 3.1.5 (See Corollary 3.3.9). Suppose \mathcal{D} is a structure which includes equality as a relation and \mathcal{E} is pp constructible in \mathcal{D} . If $\text{CSP}_B(\mathcal{D})$ is $\mathbf{\Pi}_1^1$, effectivizable, or essentially classical then so too is $\text{CSP}_B(\mathcal{E})$.

The case of dichotomy theorems is somewhat mysterious. See the comments before Theorem 3.2.15. It remains to make the algebraic characterizations of these classes explicit and to test how far these results extend beyond structures with equality.

Definition 3.1.6. Say that an operation $f : D^n \rightarrow D$ is

1. **totally symmetric** if $f(x_1, \dots, x_n)$ only depends on $\{x_1, \dots, x_n\}$
2. **Siggers** if $n = 4$ and $f(r, a, r, e) = f(a, r, e, a)$
3. A **dual discriminator** if $n = 3$ and $f(x, y, z)$ is the repeated values among x, y , and z if there is one and x otherwise.

Our first results says that the intractable fork of the CSP dichotomy adapts to the Borel setting.

Theorem 3.1.7 (See Theorem 3.4.3). If \mathcal{D} does not admit a Siggers polymorphism, then $\text{CSP}_B(\mathcal{D})$ is Σ_2^1 -complete.

Corollary 3.1.8 ($P \neq \text{NP}$). If $\text{CSP}(\mathcal{D})$ is NP-complete, then $\text{CSP}_B(\mathcal{D})$ is Σ_2^1 -complete

This theorem gives several interesting new examples of Σ_2^1 -complete problems in Borel combinatorics. We can also use this theorem to import other finitary complexity dichotomy wholesale. For instance, Corollary 3.4.5 gives a descriptive set theoretic analog of the Hell–Nešetřil theorem on graphs.

It is natural to ask if the converse of Theorem 3.1.7 holds:

Problem 3.1.9. Is there some \mathcal{D} with $\text{CSP}_B(\mathcal{D}) \in P$ but $\text{CSP}_B(\mathcal{D})$ Σ_2^1 -complete?

More generally, we can ask if a Borel CSP complexity dichotomy holds. The answer is yes if we assume Σ_2^1 -determinacy [TV21, Remark 3.3], but it is open if this holds in ZFC:

Conjecture 3.1.10. For every \mathcal{D} , $\text{CSP}_B(\mathcal{D})$ either Π_1^1 or Σ_2^1 -complete.

Our second main result gives a complete characterization of essentially classical structures. This builds on the characterization of width 1 structures by Dalmau and Pearson 3.2.11.

Theorem 3.1.11. For any finite relational structure \mathcal{D} , $\text{CSP}_B(\mathcal{D})$ is essentially classical if and only if \mathcal{D} admits totally symmetric polymorphism of arbitrarily high arity.

One direction is Theorem 3.5.2 and the other is Theorem 3.5.7. Totally symmetric polymorphisms are quite strong, so the main content of this theorem is to rule out exotic essentially classical problems. Roughly, the only method to prove a problem is essentially classical is the reflection theorem. By examining the proof, we can get the following:

Corollary 3.1.12. If $\text{CSP}_B(\mathcal{D})$ is essentially classical, it is effectivizable.

For effectivizable CSPs in general, we have the following partial result.

Theorem 3.1.13 (See Theorem 3.6.2). If \mathcal{D} has a dual discriminator polymorphism, then $\text{CSP}_B(\mathcal{D})$ is effectivizable.

This theorem along with Theorem 3.1.7 gives a descriptive analog of the Barto–Kozik–Niven theorem on smooth digraphs, Corollary 3.6.3. The structures indicated in the above theorem are the simplest from the class of bounded width structures (see definition 3.2.12). The following question is natural:

Problem 3.1.14. Is every bounded width Borel CSP effectivizable?

These results are almost enough to recover a descriptive set theoretic analog of Schaefer’s dichotomy for Boolean CSPs [Sch78].

Definition 3.1.15. For $k \in \mathbb{N}$, $k\text{SAT}$ is the structure on $\{0, 1\}$ equipped with each of the following relations: for $0 \leq i \leq k$

$$D_i(x_1, \dots, x_k) :\Leftrightarrow \left(\bigvee_{j \leq i} \neg x_j \right) \vee \left(\bigvee_{j > i} x_j \right) = 1$$

For \mathbb{F} a finite field, we write $\mathbb{F}(k)$ with the structure whose domain is the same as \mathbb{F} equipped with one relation for each affine subset of \mathbb{F}^k , i.e. each set of the form

$$\left\{ (x_1, \dots, x_n) : \sum_i a_i x_i = a_{n+1} \right\}$$

for $a_1, \dots, a_{n+1} \in \mathbb{F}$.

Note that an instance of $k\text{SAT}$ is a Boolean formula in $k\text{CNF}$. This is abusing notation slightly as computer scientists typically refer to $\text{CSP}(k\text{SAT})$ as $k\text{SAT}$, but hopefully this will not cause confusion. And, it turns out that all of the structures $\mathbb{F}(k)$ are equivalent for all $k \geq 3$ (see the comments after Definition 3.2.3).

Corollary 3.1.16 (See Corollary 3.6.4). If \mathcal{D} is a structure on $\{0, 1\}$, one of the following holds:

1. \mathcal{D} has a totally symmetric polymorphism and $\text{CSP}(\mathcal{D})$ is essentially classical and effectivizable
2. \mathcal{D} is pp constructible in 2SAT and $\text{CSP}_B(\mathcal{D})$ is effectivizable
3. $\text{CSP}(\mathcal{D})$ is NP-complete and $\text{CSP}_B(\mathcal{D})$ is Σ_2^1 -complete, or
4. \mathcal{D} is pp constructible in $\mathbb{F}_2(3)$ and vice versa.

This leaves us with the burning question of projective complexity of linear algebra:

Problem 3.1.17. Is $\text{CSP}_B(\mathbb{F}(3))$ Σ_2^1 -complete for every \mathbb{F} ?

If the answer to Questions 3.1.17 and 3.1.14 are both positive, then Conjecture 3.1.10 is true, and further we get the following:

Conjecture 3.1.18. For any finite structure \mathcal{D} , $\text{CSP}_B(\mathcal{D})$ is Π_1^1 if and only if it is effectivizable.

3.2 Background on CSPs

In this section, we review some of the basic algebraic theory of CSPs. The material we cover here is minimal. For a more detailed survey, see [BKW17].

Definition 3.2.1. For two relational structures in the same language, $\mathcal{X} = (X, \tau)$ and (\mathcal{D}, τ) , a **homomorphism** from \mathcal{X} to \mathcal{D} is a function $f : X \rightarrow D$ so that

$$(\forall R \in \tau, x_1, \dots, x_n \in X) (x_1, \dots, x_n) \in R^{\mathcal{X}} \rightarrow (f(x_1), \dots, f(x_n)) \in R^{\mathcal{D}}.$$

Or, in words, if M is an array where all rows are in R , then applying f to the columns gives another element of R .

A **polymorphism** of \mathcal{X} is a homomorphism from \mathcal{X}^n to \mathcal{X} , where we take the so-called categorical product. That is, f is a polymorphism of \mathcal{D} if, given $\langle x_{ij} : i \leq n, j \leq k \rangle \in D^{n \times k}$ and a relation $R \in \mathcal{D}$

$$((\forall i \leq n) R(x_{i1}, \dots, R_{ik})) \Rightarrow R(f(\langle x_{i1} : i \leq n \rangle), \dots, f(\langle x_{ik} : i \leq n \rangle))$$

For any structure \mathcal{D} , $\text{Pol}(\mathcal{D})$ is the algebra of polymorphisms of \mathcal{D} .

A **clone** is an algebra equipped with every projection operation and whose collection of operations is closed under compositions. For a set of operations A , $\langle A \rangle$ is the smallest clone whose operations contain A . We say A **generates** $\langle A \rangle$.

The following examples are straightforward to verify:

1. For any structure \mathcal{D} , $\text{Pol}(\mathcal{D})$ is a clone
2. For any finite field \mathbb{F} , $\text{Pol}(\mathbb{F}(3))$ is the collection of linear functionals

$$f(x_1, \dots, x_n) = \sum_i a_i x_i$$

with $\sum_i a_i = 1$

3. The dual discriminator on $\{0, 1\}$, also called the majority function or maj, generates $\text{Pol}(2\text{SAT})$.
4. HornSAT is satisfaction problem for Horn sentences. More specifically, HornSAT has domain $\{0, 1\}$ and includes the unary predicates $\{0\}$ and $\{1\}$ and each relation of the form

$$R(x_1, \dots, x_n, y) \Leftrightarrow (x_1 \wedge \dots \wedge x_n) \rightarrow y.$$

The binary min, \wedge , generates $\text{Pol}(\text{HornSAT})$

5. NAE is the 2-coloring problem for ternary hypergraphs, i.e.

$$\text{NAE} = (\{0, 1\}, \{(x_1, x_2, x_3) : \neg(x_1 = x_2 = x_3)\}).$$

Negation generates $\text{Pol}(\text{NAE})$.

6. For the directed 3-cycle $\mathcal{D} = (\{r, p, s\}, \{(p, r), (r, s), (s, p)\})$, $f \in \text{Pol}(\mathcal{D})$ if and only if the cyclic permutation $\pi = (rps)$ is an automorphism of f . For instance, $\text{Pol}(\mathcal{D})$ contains π , the dual discriminator on $\{r, p, s\}$ and the binary rock-paper-scissors operation, (\star):

\star	r	p	s
r	r	p	r
p	p	p	s
s	r	s	s

Note that HornSAT is not finite, but it will turn out that it is pp definable in a finite number of its relations. For item (3), note that for any 3×2 matrix over \mathbb{F}_2 , applying the majority function to each column returns one of the rows. So, the majority function preserves all binary relations on $\{0, 1\}$.

Definition 3.2.2. We say \mathcal{X} is a **core** if any homomorphism $f : \mathcal{X} \rightarrow \mathcal{X}$ is an automorphism (i.e. bijective).

Two structures \mathcal{D} and \mathcal{E} are **homomorphically equivalent** if there are homomorphisms $f : \mathcal{D} \rightarrow \mathcal{E}$ and $g : \mathcal{E} \rightarrow \mathcal{D}$.

Any structure \mathcal{D} is homomorphically equivalent to a (unique up to isomorphism) core, which we refer to as the **core of \mathcal{D}** .

Note that if \mathcal{D} and \mathcal{E} are homomorphically equivalent, then $\text{CSP}(\mathcal{D}) = \text{CSP}(\mathcal{E})$. For example, K_n is a core for all n , and the core of any bipartite graph is K_2 .

Definition 3.2.3. If \mathcal{D} and \mathcal{E} are two structures with the same domain, we say \mathcal{E} is **pp definable** in \mathcal{D} if every relation R of \mathcal{E} can be written as a positive primitive formula over \mathcal{D} , i.e.

$$R(x_1, \dots, x_n) \Leftrightarrow (\exists y_1, \dots, y_k) \bigwedge_i \alpha_i(x_1, \dots, x_n, y_1, \dots, y_k)$$

where each α_i either asserts a relation from \mathcal{D} or an equality.

We say \mathcal{E} is **pp interpretable** in \mathcal{D} if it is interpretable in the model-theoretic sense with positive primitive formula, i.e. \mathcal{E} is a quotient of a pp definable structure of some power of \mathcal{D} along a pp definable equivalence relation.

And, we say \mathcal{E} is **pp constructible** in \mathcal{D} if there is a chain of structures

$$\mathcal{D} = \mathcal{E}_1, \mathcal{E}_2, \dots, \mathcal{E}_n = \mathcal{E}$$

so that for each i , one of the following holds:

1. \mathcal{E}_{i+1} is pp interpretable in \mathcal{E}_i
2. \mathcal{E}_i is homomorphically equivalent to \mathcal{E}_{i+1}
3. \mathcal{E}_i is a core and \mathcal{E}_{i+1} is \mathcal{E}_i expanded by a singleton unary relation $U(x) \Leftrightarrow x = a$

For example, given any $a_1, a_2, a_3, a_4, b \in \mathbb{F}$

$$a_1x_1 + a_2x_2 + a_3x_3 + a_4x_4 = b$$

if and only if

$$(\exists y_1, y_2) a_1x_1 + a_2x_2 = y_1 \wedge a_3x_3 + y_1 = y_2 \wedge a_4x_4 + y_2 = a.$$

So $\mathbb{F}(3)$ pp constructs $\mathbb{F}(4)$ for any finite field \mathbb{F} . A similar trick shows that $\mathbb{F}(3)$ pp constructs every $\mathbb{F}(k)$ and that HornSAT is pp definable in a finite number of its relations.

Since every Boolean formula can be converted to a formula in 3CNF by introducing new variables as above, 3SAT pp defines every structure on $\{0, 1\}$. Further, by considering binary expansions of a general relation, one can show that 3SAT pp interprets every structure. Because NAE has a nontrivial automorphism, it does not pp define 3SAT, but one can check that it pp constructs 3SAT (and thus all structures).

The basic idea of the algebraic approach to CSPs is that pp constructions preserve all complexity-theoretic information about a CSP, and that pp constructibility of structures

is captured by the (height 1) varieties of their polymorphism algebras. For completeness, we will explain this correspondence in general, though we will only refer to a handful of well-understood varieties in this work.

Definition 3.2.4. For an algebra \mathbb{A} , an **identity** is a true statement of the form

$$(\forall x_1, \dots, x_n, y_1, \dots, y_m) \tau(x_1, \dots, x_n) = \sigma(y_1, \dots, y_m)$$

where σ and τ are compositions of operations from \mathbb{A} . An identity is **height 1** if each of σ and τ contain exactly one function symbol (without repetition). As is somewhat standard convention, we will often leave the universal quantifier implicit when defining identities.

The **variety** generated by \mathbb{A} , $\text{Var}(\mathbb{A})$, is the class of all algebras (in the same signature) which satisfy the same identities as \mathbb{A} . The **h1 variety** generated by \mathbb{A} , $\text{Var}_{h1}(\mathbb{A})$ is the class of algebras which satisfy the same height 1 identities.

The identities defining Siggers and totally symmetric polymorphisms are height 1, but associativity and idempotence are not. The variety generated by $\text{Pol}(3\text{SAT})$ is the class of all projection algebras. On the other hand, the majority function satisfies

$$\text{maj}(x, x, y) = \text{maj}(x, y, x) = \text{maj}(y, x, x)$$

so $\text{Pol}(2\text{SAT})$ and $\text{Pol}(3\text{SAT})$ generate different h1 varieties.

We will need a more detailed analysis of $\text{Var}(\mathbb{A})$.

Definition 3.2.5. We say that \mathbb{B} is a **reduct** of \mathbb{A} if \mathbb{B} has the same domain as \mathbb{A} and is equipped with a subset of the operations from \mathbb{A} . In this case, we write $\mathbb{B} \subseteq \mathbb{A}$.

We say that \mathbb{B} is a **subalgebra** of \mathbb{A} if the domain of \mathbb{B} is a subset of the domain of \mathbb{A} , and \mathbb{B} comes equipped with all of the restrictions of operations from \mathbb{A} .

For C a class of algebras, we define:

$$H(C) := \{f(\mathbb{A}) : \mathbb{A} \in C, f \text{ a homomorphism}\}$$

$$S(C) := \{\mathbb{B} : (\exists \mathbb{A} \in C) \mathbb{B} \text{ a subalgebra of } \mathbb{A}\}$$

$$P(C) := \{\mathbb{A}^\kappa : \mathbb{A} \in C, \kappa \text{ is a cardinal}\}$$

We write $H(\mathbb{A})$ for $H(\{\mathbb{A}\})$ and likewise for $S(\mathbb{A})$ and $P(\mathbb{A})$.

The operations H and S will be important for us later.

Theorem 3.2.6. For finite relational structures \mathcal{D} and \mathcal{E}

1. \mathcal{D} pp defines \mathcal{E} iff $\text{Pol}(\mathcal{E}) \subseteq \text{Pol}(\mathcal{D})$
2. \mathcal{D} pp interprets \mathcal{E} iff a reduct of $\text{Pol}(\mathcal{E})$ is in $\text{Var}(\text{Pol}(\mathcal{D})) = \text{HSP}(\text{Pol}(\mathcal{D}))$ iff every identity satisfied by elements of $\text{Pol}(\mathcal{D})$ is satisfied by elements of $\text{Pol}(\mathcal{E})$
3. \mathcal{D} pp constructs \mathcal{E} iff a reduct of $\text{Pol}(\mathcal{E})$ is in $\text{Var}_{h1}(\text{Pol}(\mathcal{D}))$ iff every height 1 identity satisfied by elements of $\text{Pol}(\mathcal{D})$ is satisfied by elements of $\text{Pol}(\mathcal{E})$.

Proof. Statement (1) was proven independently in 60s by Geiger and Bodnarchuk, Kaluzh-nin, Kotov, and Romov [Gei68][BKK69]. The fact that $\text{Var}(\mathbb{A}) = \text{HSP}(\mathbb{A})$ is Birkhoff's HSP theorem [Bir35]. The equivalence to pp interpretation is essentially by definition and is im-plicit in Bulatov, Jeavons, and Krokhin [BJK05]. Statement (3) is due to Barto, Kozik, and Pinsker [BOP15]. □

In fact slightly more is true of item (2). For $\mathbb{A} = \text{Pol}(\mathcal{D})$, the polymorphism algebras in $H(\mathbb{A})$, $S(\mathbb{A})$, and $P(\mathbb{A})$ correspond to pp definable quotients, pp definable substructures, and powers of \mathcal{D} respectively.

Corollary 3.2.7. For structures \mathcal{D} and \mathcal{E} , \mathcal{E} is a pp definable quotient of a pp definable substructure of \mathcal{D} if and only if $\text{Pol}(\mathcal{E}) \in \text{HS}(\text{Pol}(\mathcal{D}))$.

Proof. We will show that any \mathcal{E} with $\text{Pol}(\mathcal{E}) \in \text{HS}(\text{Pol}(\mathcal{D}))$ is a quotient of a substructure. The converse is straightforward.

First note that a subalgebra of $\text{Pol}(\mathcal{D})$ is a subset of D which is closed under the operations in $\text{Pol}(\mathcal{D})$. By Theorem 3.2.6, this is the same as a pp definable subset of D .

Now suppose that $\varphi : D \rightarrow E$ is a homomorphism of algebras from $\text{Pol}(\mathcal{D})$ onto $\text{Pol}(\mathcal{E})$, and extend φ to a homomorphism $\varphi : D^n \rightarrow E^n$. For any operation $f \in \text{Pol}(\mathcal{D})$, if $\varphi(x_1, \dots, x_n) = \varphi(y_1, \dots, y_n)$ then

$$\begin{aligned} \varphi(f(x_1, \dots, x_n)) &= f(\varphi(x_1, \dots, x_n)) \\ &= f(\varphi(y_1, \dots, y_n)) \\ &= \varphi(f(y_1, \dots, y_n)) \end{aligned}$$

So the kernel of φ , i.e. the equivalence relation $\varphi(x) = \varphi(y)$, is invariant under $\text{Pol}(\mathcal{D})$ and is pp definable in \mathcal{D} . Similarly, the pullback of any relation in \mathcal{E} is pp definable. \square

As a corollary of Theorem 3.2.6, any class of structures which is closed under pp constructions (or interpretations or definitions) admits an algebraic description in terms of height 1 identities (or identities or polymorphisms). We survey several such classes and their algebraic descriptions below.

Bulatov, Jeavons, and Krokhin showed that if \mathcal{D} pp constructs \mathcal{E} then $\text{CSP}(\mathcal{E})$ is polynomial time reducible to $\text{CSP}(\mathcal{D})$ [BJK05]. The CSP dichotomy theorem gives the corresponding algebraic characterization of the polynomial time complexity classes for CSPs. We collect a few equivalent characterizations here.

Definition 3.2.8. If \mathcal{D} admits a Siggers term, we say \mathcal{D} is tractable. We say \mathcal{D} is intractable otherwise.

So the CSP dichotomy theorem says that if $\text{P} \neq \text{NP}$ then intractable is synonymous with NP-complete and tractable is synonymous with polynomial time solvable.

Theorem 3.2.9. \mathcal{D} is tractable iff any of the following hold:

1. $\text{Pol}(\mathcal{D})$ contains a Taylor operation, i.e. an operation T so that for every i there are $z_1, \dots, z_n \in \{x, y\}$ so that T satisfies an identity of the form

$$T(z_1, z_2, \dots, z_{i-1}, x, z_{i+1}, \dots, z_n) = T(z_1, z_2, \dots, z_{i-1}, y, z_{i+1}, \dots, z_n).$$

2. There is a weak near unanimity (or WNU) polymorphism $W \in \text{Pol}(\mathcal{D})$, i.e. an operation satisfying

$$W(y, x, x, \dots, x) = W(x, y, x, \dots, x) = W(x, x, y, \dots, x) = \dots = W(x, x, x, \dots, y)$$

3. For all but finitely many primes p , \mathcal{D} has cyclic polymorphism of arity p , i.e. there is $C \in \text{Pol}(\mathcal{D})$ satisfying

$$C(x_1, x_2, x_3, \dots, x_{p-1}, x_p) = C(x_2, x_3, x_4, \dots, x_p, x_1)$$

4. \mathcal{D} does not pp constructs every finite relational structure

Proof. Taylor proved an idempotent algebra \mathbb{A} admits Taylor term if and only if $\text{Var}(\mathbb{A})$ does not contain a projection algebra [Tay77]. Since the Taylor identities are height 1, this implies $\text{Pol}(\mathcal{D})$ does not contain a Taylor term if and only if \mathcal{D} pp constructs every finite relational structure. The equivalence of Taylor terms and WNU terms is shown in [MM08], the equivalence with cyclic terms in [BK12], and the equivalence with Siggers terms in [Sig10] and [KMM14]. \square

Historically, the Taylor identities were the first to appear in the literature. In the 1970s, Motivated by questions in algebraic topology, Taylor showed that these identities characterize idempotent varieties which omit projection algebras. Bulatov, Jeavons, and Krokhin originally stated the algebraic CSP dichotomy conjecture in terms of Taylor identities, and many of the equivalent forms above were motivated by computational complexity questions. We include WNU operations here as they seem to show up most often in the literature.

The second class of structures we look at is described by a simple constraint propagation algorithm. Given an instance \mathcal{X} , assign to each variable $x \in \mathcal{X}$ a unary constraint U_x , initialized to \mathcal{D} . This U_x will represent to possible values x can take. If we ever see that $R(x, y_1, \dots, y_n)$, but that $d \in U_x$ cannot be matched to elements of the U_{y_i} s to get an element of R , then remove d from U_x . Repeat until this each U_x stabilizes. If any U_x is empty there is no solution. We say that a problem is width 1 if every instance with each U_x nonempty has a solution. More formally:

Definition 3.2.10. Say that an instance \mathcal{X} of \mathcal{D} is **arc-consistent** if there are unary predicates $U_x \subseteq \mathcal{D}$ for $x \in \mathcal{X}$ which are pp definable in \mathcal{D} so that, if $(x_1, \dots, x_n) \in R^{\mathcal{X}}$, then

$$\pi_i(U_{x_1} \times \dots \times U_{x_n} \cap R^{\mathcal{D}}) = U_{x_i}.$$

A structure is **width 1** if every arc-consistent instance has a solution.

Width 1 structures will play an important role in Section 3.5. An example of a width 1 structure is HornSAT. The arc consistency algorithm in this case amounts to classical unit propagation. A Theorem of Dalmau and Pearson algebraically characterizes width 1 structures.

Theorem 3.2.11 ([DP99]). A structure \mathcal{D} is width 1 if and only if it admits a totally symmetric polymorphism of arbitrarily high arities.

Indeed, HornSAT has the n -ary \bigwedge as a polymorphism for all n .

Arc-consistency only considers unary information, We can generalize this to consider binary interactions between constraints. For instance, the usual algorithm testing 2-colorability of graphs involves testing for odd length chains of binary relations.

Definition 3.2.12. A **path** in a relational structure \mathcal{X} is a sequence

$$P = (x_1, (e_1, R_1), x_2, (e_1, R_2), \dots, x_n)$$

if each $x_i \in X$, each $e_i \in R_i$, and for each i there are j_i, k_i so that $(x_i, x_{i+1}) = \pi_{j_i k_i} e_i$. We say that P has **coordinates** $(j_1, k_1, j_2, k_2, \dots, j_{n-1}, k_{n-1})$.

A **closed path** in a relational structure \mathcal{X} is a path $(x_1, (e_1, R_1), \dots, (e_{n-1}, R_{n-1}), x_n)$ with $x_1 = x_n$.

An instance \mathcal{X} of \mathcal{D} is **cycle-consistent** if it is arc-consistent as witnessed by predicates U_x and for any closed path P with coordinates $(j_1, k_1, \dots, j_{n-1}, k_{n-1})$,

$$(\forall a_1 = a_n \in U_{x_1})(\exists a_2, a_3, \dots, a_{n-1}) \bigwedge_{1 \leq i \leq n-1} (U_{x_i}(a_i) \wedge \pi_{j_i, k_i} R_i(a_i, a_{i+1})).$$

A structure has **bounded width** if every cycle-consistent instance has a solution.

Really, this definition hides a theorem. One could generalize arc-consistency in any number of ways to consider tuples of arbitrary large arity, but these all turn out to be redundant. See [Bar14] [Bra19]. Barto and Kozik characterized the bounded width CSPs, resolving another conjecture of Feder and Vardi.

Theorem 3.2.13 ([BK09]). A structure is bounded width iff $\text{HSP}(\text{Pol}(\mathcal{D}))$ does not contain $\text{Pol}(\mathcal{F})$ for any finite field \mathbb{F} iff $\text{HS}(\text{Pol}(\mathcal{D}))$ does not contain $\text{Pol}(\mathbb{F})$ for any finite field \mathbb{F} .

Cycle-consistency will be important in Section 3.5, and the bounded width structures will play a role in Section 3.6.

Our last example is a class of structures which is not closed under pp constructions, but still has an algebraic description.

Definition 3.2.14. A **basis** for $\text{CSP}(\mathcal{D})$ is a set of structures F so that $\mathcal{X} \in \text{CSP}(\mathcal{D})$ if and only if no element of F admits a homomorphism into \mathcal{X} . Bases for $\text{CSP}_B(\mathcal{D})$ are defined *mutatis mutandis*.

Bases are also called complete obstructing sets in the computer science literature. Descriptive set theorists are very interested in the question of when $\text{CSP}_B(\mathcal{D})$ has a finite basis.

The corresponding classical question can be answered in terms of polymorphisms with extra tolerance (which require reference to the relational structure to define).

Theorem 3.2.15 ([LLT06]). For any structure \mathcal{D} , $\text{CSP}(\mathcal{D})$ has a finite basis if and only if \mathcal{D} admits a polymorphism f so that

$$f(y, x, x, \dots, x) = f(x, y, x, \dots, x) = f(x, x, y, \dots, x) = \dots = f(x, x, x, \dots, y) = x$$

and if $k - 1$ of the tuples e_1, \dots, e_k are in R for some relation R of \mathcal{D} , then

$$f(e_1, \dots, e_k) \in R.$$

We will not use this theorem in this thesis, though bases will play a role in several places. Note that the above theorem does not refer only to $\text{Pol}(\mathcal{D})$. Indeed, it is not true that a finite basis for $\text{CSP}(\mathcal{D})$ will pass through a pp definition to give a finite basis for $\text{CSP}(\mathcal{E})$.

3.3 Simple constructions

We would like a Borel analogue of the polynomial time reductions induced by pp constructions. Unfortunately, the finitary construction generally requires taking a quotient, which is not always possible in the descriptive setting (for instance, \mathbb{R}/\mathbb{Q} is not a standard Borel space). Here we introduce a stronger notion of simulation, which we call a simple construction, which does not present this difficulty. And, we prove some combinatorial lemmas relating simple and pp constructions.

Definition 3.3.1. For \mathcal{E} and \mathcal{D} structures on the same domain, say that \mathcal{D} **simply defines** \mathcal{E} if every relation in \mathcal{E} can be written as an existentially quantified conjunction of relations in \mathcal{D} . We call the formulas defining \mathcal{E} in terms of \mathcal{D} relations a **simple definition**.

Say that \mathcal{D} **simply interprets** \mathcal{E} if \mathcal{E} is a quotient of a structure simply definable in a power of \mathcal{D} . We call the defining formulas a **simple interpretation**.

And, say that \mathcal{D} **simply constructs** \mathcal{E} if \mathcal{E} can be built from \mathcal{D} by a chain of simple interpretations, homomorphic equivalences, and singleton expansions of cores. We call this sequence a **simple construction**.

That is, simple constructions are pp constructions where we do not use equality (unless our structures come equipped with equality as a relation). Our first main lemma about simple constructions comes from a careful analysis of the proof that singleton expansions of cores do not change computational complexity.

Lemma 3.3.2. For any finite core structure \mathcal{D} and $c \in D$, the following predicate is simply definable in \mathcal{D} :

$$(\ =_c) := \{(x, x) : (\exists f \in \text{Aut}(\mathcal{D}))f(c) = x\}.$$

Proof. Suppose $D = \{1, 2, \dots, n\}$ and $c = 1$. The predicate H given by

$$\begin{aligned} H(x_1, \dots, x_n) &:\Leftrightarrow d \mapsto x_d \text{ defines an endomorphism of } \mathcal{D} \\ &\Leftrightarrow \bigwedge_R \bigwedge_{e \in R} R(x_{e_1}, \dots, x_{e_n}) \end{aligned}$$

is simply definable. So, it suffices to show that

$$a =_c b \Leftrightarrow (\exists x_2, \dots, x_n) H(a, x_2, \dots, x_n) \wedge H(b, x_2, \dots, x_n).$$

If $a =_c b$, then there is some isomorphism $f : \mathcal{D} \rightarrow \mathcal{D}$ with $f(c) = a$. Then the assignment $x_d = f(d)$ satisfies $H(a, x_2, \dots, x_n) \wedge H(b, x_2, \dots, x_n)$.

Conversely, if $H(a, x_2, \dots, x_n) \wedge H(b, x_2, \dots, x_n)$, then we have homomorphisms f, g with $f(c) = a, g(c) = b$ and $f(d) = g(d)$ for $d \neq c$. These homomorphisms must be automorphisms since \mathcal{D} is a core. And if $a \neq b$, we must have that one of f or g is not onto, which is impossible. So, $a =_c b$. \square

Corollary 3.3.3. If \mathcal{D} has a transitive automorphism group and \mathcal{D} pp constructs \mathcal{E} , then \mathcal{D} simply constructs \mathcal{E} .

Proof. We can replace \mathcal{D} by its core, which will still be transitive. Then, by the above lemma \mathcal{D} simply constructs $(=)_c = (=)$. \square

Our second main lemma comes from a careful analysis of the proof that $\text{Pol}(\mathcal{D})$ invariant relations are pp definable in \mathcal{D} .

Definition 3.3.4. Say that a relation R implies an equation if $R(x_1, \dots, x_n)$ implies $x_i = x_j$ for some i, j , i.e. $\pi_{ij}(R) \subseteq (=)$.

Lemma 3.3.5. If R does not imply any equations and R is pp definable in \mathcal{D} , then R is simply definable in \mathcal{D} .

Proof. Let M be a matrix whose rows are the tuples in R , and let σ_i be the i^{th} column of M . Note that R not implying any equations means that the σ_i s are all distinct. Since $R^{\mathcal{D}^n}(\sigma_1, \dots, \sigma_k)$ holds in the product, if there is a polymorphism $f \in \text{Pol}(\mathcal{D})$ so that $f(\sigma_i) = x_i$, then $R(x_1, \dots, x_k)$. Conversely, if $R(x_1, \dots, x_k)$, say (x_1, \dots, x_k) is the j^{th} row of M , then the projection π_j is a polymorphism with $\pi_j(\sigma_i) = x_i$. Thus,

$$R(x_1, \dots, x_k) \Leftrightarrow (\exists f \in \text{Pol}(\mathcal{D})) f(\sigma_i) = x_i.$$

Since the σ_i are all distinct, we can eliminate equality in the above definition by replacing each instance of σ_i with x_i . More formally, Suppose $\mathcal{D}^k = \{\sigma_1, \dots, \sigma_k, \sigma_{k+1}, \dots, \sigma_m\}$, and let

$$P(x_1, \dots, x_m) := \Leftrightarrow \sigma_i \mapsto x_i \text{ is a polymorphism of } \mathcal{D}.$$

Then P is simply definable, and we have a simple definition of R given by

$$R(x_1, \dots, x_k) \Leftrightarrow (\exists x_{k+1}, \dots, x_m) P(x_1, \dots, x_m).$$

\square

Corollary 3.3.6. If no relation in \mathcal{E} implies an equation and $\text{Pol}(\mathcal{E}) \in \text{HS}(\text{Pol}(\mathcal{D}))$, then \mathcal{D} simply constructs \mathcal{E} .

Proof. If \mathcal{E} is in $\text{HS}(\text{Pol}(\mathcal{D}))$, then \mathcal{E} is a pp definable quotient of a pp definable substructure of \mathcal{D} , i.e. there is some $U \subseteq \mathcal{D}$ and some $f : U \rightarrow \mathcal{E}$ with U and $f^{-1}(R)$ pp definable for each relation R in \mathcal{E} . No unary predicate can imply an equation, and since no R in \mathcal{E} implies an equation neither does any $f^{-1}(R)$. So by our lemma, all of these are in fact simply definable in \mathcal{D} . \square

Corollary 3.3.7. If \mathcal{D} is not bounded width, then \mathcal{D} simply constructs $\mathbb{F}(3)$ for some finite field \mathbb{F} .

We're ready to prove our descriptive analog of Bulatov and Jeavons's theorem about polynomial time reductions. Since we have not specified our coding, we will leave it to the reader to verify that the construction below is Δ_1^1 in the codes, though this is straightforward using the coding described in the appendix and Lemma A.0.7.

Theorem 3.3.8. If \mathcal{D} simply constructs \mathcal{E} , then $\text{CSP}_B(\mathcal{E})$ Borel reduces to $\text{CSP}_B(\mathcal{D})$. In fact, there are maps F, G , and H which are Δ_1^1 in the codes so that:

1. If \mathcal{X} is an instance of \mathcal{E} then $F(\mathcal{X})$ is an instance of \mathcal{D}
2. If g is a solution to \mathcal{X} , then $G(g)$ is a solution to $F(\mathcal{X})$
3. If h is a solution to $F(\mathcal{X})$, then $H(h)$ is a solution to \mathcal{X} .

And there is some finite N so that, if each $x \in \mathcal{X}$ appear in fewer than κ tuples in relations \mathcal{X} , then each $y \in F(\mathcal{X})$ appears in fewer than $N \times \kappa$ tuples in $F(\mathcal{X})$.

This last clause means that F sends bounded degree, locally finite, and locally countable instances to the same.

Proof. It suffices to consider two cases: (1) \mathcal{D} is a core and \mathcal{E} is an expansion of \mathcal{D} by singleton unary predicates, and (2) \mathcal{E} is simply interpretable in \mathcal{D} .

For (1) the main difficulty is controlling degrees, i.e. meeting the last clause of the theorem. We may assume $(=_c)$ is in the signature of \mathcal{D} for each $c \in D$ and $\mathcal{E} = \mathcal{D} \cup \{U_d\}$, where $U_d(x) \Leftrightarrow x = d$. If we did not want to worry about degree, we could let $F(\mathcal{X})$ be \mathcal{X} along with a copy of \mathcal{D} and extra relations saying $x =_d d$ whenever $U_d^{\mathcal{X}}(x)$. Then any solution to \mathcal{X} , say g , extends to a solution $G(g)$ of $F(\mathcal{X})$ by setting $G(g)(c) = c$ for $c \in \mathcal{D}$. And if h is a solution to $F(\mathcal{X})$, then h restricts to an automorphism f on the copy of \mathcal{D} . So, we have a solution to \mathcal{X} given by $H(h) = f^{-1} \circ h$ on \mathcal{X} .

In the construction we just sketched, the copy of \mathcal{D} ends up having quite large degree. To correct this, we give each variable in \mathcal{X} its own copy of \mathcal{D} . So, define $F(\mathcal{X})$ as follows: $F(\mathcal{X})$ has domain $X \sqcup (X \times D)$ and the following relations

- For R a relation in \mathcal{D} with $R \neq (=_c)$ for any c

$$R^{F(\mathcal{X})} = R^{\mathcal{X}} \sqcup \{((x, a_1), \dots, (x, a_k)) : x \in X, R^{\mathcal{D}}(a_1, \dots, a_k)\}$$

- For $c \neq d$

$$({}_c)_{F(\mathcal{X})} = ({}_c)^{\mathcal{X}} \sqcup \{((x, c), (y, c)) : (x, y) \text{ share some relation in } \mathcal{X}\}$$

- And,

$$({}_d)_{F(\mathcal{X})} = ({}_d)^{\mathcal{X}} \sqcup \{(x, (x, d)) : U_d^{\mathcal{X}}(x)\}$$

Similar to the above, any solution g to \mathcal{X} gives a solution $G(g)$ to $F(\mathcal{X})$ equal to the identity on any copy of \mathcal{D} . And if h is a solution to $F(\mathcal{X})$, $\tilde{h}_x(d) := h(x, d)$ defines an automorphism of \mathcal{D} for each x . Further $\tilde{h}_x = \tilde{h}_y$ whenever x, y share some relation in \mathcal{X} , so we get a solution to \mathcal{X} $H(h)(x) = (\tilde{h}_x^{-1} \circ h)(x)$.

For (2), the construction of $F(\mathcal{X})$ is identical to the Bulatov and Jeavons construction, but we use the Luzin–Novikov theorem to find G . Suppose we are given the following data from the definition of simple interpretation:

1. An onto function $c : A \rightarrow E$ with a simple definition for $A \subseteq D^n$:

$$\bar{x} \in A \Leftrightarrow (\exists \bar{z}) \bigwedge_i \alpha_{A,i}(\bar{x}, \bar{z})$$

2. For each relation $R \in \mathcal{E}$ a simple definition for $c^{-1}(R)$:

$$\bar{x}_1, \dots, \bar{x}_k \in c^{-1}(R) \Leftrightarrow (\exists \bar{z}) \bigwedge_i \alpha_{R,i}(\bar{x}_1, \dots, \bar{x}_k, \bar{z}).$$

we construct an instance $F(\mathcal{X})$ of \mathcal{D} as follows. For each variable $x \in X$ introduce tuples of variables \bar{y}_x (with arity n as in item 1 above) and $\bar{z}_{A,x}$ (with the same arity as \bar{z} in the definition of item 1), and for each tuple $(x_1, \dots, x_n) \in R^{\mathcal{X}}$ introduce tuples $\bar{z}_{R,x_1, \dots, x_n}$ (with the same arity as in the definition in item 2). For each variable x add in the relations

$$\alpha_{A,i}(\bar{y}_x, \bar{z}_{A,x})$$

and for each $(x_1, \dots, x_n) \in R^{\mathcal{X}}$ add in relations

$$\alpha_{R,i}(\bar{y}_{x_1}, \dots, \bar{y}_{x_n}, \bar{z}_{R,x_1, \dots, x_n}).$$

Given a solution g to $F(\mathcal{X})$ we get a solution $G(g)$ to \mathcal{X} given by $G(g)(x) = c(g(\bar{y}_x))$ (where we apply g coordinatewise to \bar{y}_x). And if we have a solution g to \mathcal{X} , we can get a solution to $F(\mathcal{X})$ by using the Luzin–Novikov theorem to choose values from $c^{-1}(g(x))$ for each \bar{y}_z and to choose values for each $\bar{z}_{A,x}$ and $\bar{z}_{R,\bar{x}}$ from the witnesses to

$$(\exists \bar{z}) \bigwedge_i \alpha_{A,i}(\bar{y}_x, \bar{z}) \quad \text{and} \quad (\exists \bar{z}) \bigwedge_i \alpha_{R,i}(\bar{y}_{x_1}, \dots, \bar{y}_{x_k}, \bar{z}).$$

□

If \mathcal{D} has equality in its signature, then the distinction between simple and pp constructions collapses and we get the following:

Corollary 3.3.9. Suppose \mathcal{D} is a structure which includes equality as a relation and \mathcal{E} is pp constructible in \mathcal{D} . If $\text{CSP}_B(\mathcal{D})$ is $\mathbf{\Pi}_1^1$, effectivizable, or essentially classical then so too is $\text{CSP}_B(\mathcal{E})$.

It follows that these classes all have algebraic characterizations when restricted to structures with equality, though the identities involved may be quite complex.

3.4 Intractable CSPs

One of the consequences of the CSP dichotomy theorem is that (assuming $P \neq NP$) any structure with $\text{CSP}(\mathcal{D})$ complete for polynomial time reductions is in fact maximal for pp constructions. Using a lemma from the previous section and a lemma of Bulatov and Jeavons, we can show that, in fact, any such structure is maximal for simple constructions. This, along with Todorćević and Vidnyánszky's Σ_2^1 -completeness theorem for 3-coloring [TV21], gives a number of new projective complexity bounds.

Lemma 3.4.1 (Bulatov–Jeavons [BJ01]). If \mathcal{D} is intractable and includes all singleton unary predicates, then $\text{HS}(\text{Pol}(\mathcal{D}))$ contains a projection algebra.

Proof. This lemma appears in a seemingly unpublished technical report, so we give a sketch of their proof here.

Let $\mathbb{A} = \text{Pol}(\mathcal{D})$ and let $\mathbb{B} \leq \mathbb{A}^n$ be a subalgebra with some homomorphism $f : \mathbb{B} \rightarrow \mathbb{X}$ onto a nontrivial projection algebra. For $I \subseteq \{1, \dots, n\}$ and $\bar{a} \in A^I$, let $\mathbb{B} \upharpoonright (I, \bar{a}) = \pi_I = \bar{a}$. Fix I maximal with f not constant on $\mathbb{B} \upharpoonright (I, \bar{a})$. Without loss of generality, we may assume $I = \{1, \dots, k\}$ and $\bar{a} = (a_1, \dots, a_k)$.

Since \mathcal{D} contains all unary predicates, \mathbb{A} is idempotent. Thus, $\mathbb{B}' = \mathbb{B} \upharpoonright (I, \bar{a})$ is a subalgebra of \mathbb{B} . And we have that $f(\mathbb{B}')$ is a nontrivial projection algebra.

Let $\mathbb{A}' = \pi_{k+1}^{-1}(\mathbb{B}')$. If there is some $a \in \mathbb{A}'$ with f is not constant on $\pi_{k+1}^{-1}(a) \cap \mathbb{B}'$, then I is not maximal. So, f factors through π_{k+1} and a nontrivial subalgebra of \mathbb{X} is an image of \mathbb{A}' as well. \square

Corollary 3.4.2. If \mathcal{D} is intractable, \mathcal{D} simply constructs all structures.

Proof. We may assume \mathcal{D} is a core with all singleton unary relations. Then, by the lemma above, $\text{Pol}(3\text{SAT}) \in \text{HS}(\text{Pol}(\mathcal{D}))$. Since 3SAT does not imply equations, by lemma 3.3.5, \mathcal{D} simply constructs 3SAT. And since 3SAT simply defines equality and is maximal for pp constructions, \mathcal{E} simply constructs all structures. \square

Since there is a Σ_2^1 -complete Borel CSP, we get the following:

Theorem 3.4.3. If \mathcal{D} is intractable, then $\text{CSP}_B(\mathcal{D})$ is Σ_2^1 -complete.

Proof. Todorcevic and Vidnyánszky showed that $\text{CSP}_B(K_3)$ is Σ_2^1 -complete [TV21]. By the above theorem, if \mathcal{D} is intractable, then $\text{CSP}_B(K_3)$ Borel reduces to $\text{CSP}_B(\mathcal{D})$. \square

Corollary 3.4.4 ($P \neq \text{NP}$). If $\text{CSP}(\mathcal{D})$ is NP-complete, then $\text{CSP}_B(\mathcal{D})$ is Σ_2^1 -complete.

In fact, since the instances from the Todorcevic–Vidnyánszky theorem are locally finite, these problems are Σ_2^1 -complete even when we restrict to locally finite instances. This can be improved to bounded degree by recent work of Brandt, Chang, Grebík, Grunau, Rozhoň, and Vidnyánszky [BCG21a]. We can use this theorem to lift results from finitary complexity theory wholesale. For instance, the Hell–Nešetřil theorem [HN90] yields:

Theorem 3.4.5. For a simple graph G , the following are equivalent

1. G is bipartite
2. G is tractable
3. $\text{CSP}_B(G)$ is effectivizable
4. $\text{CSP}_B(G)$ is Π_1^1
5. $\text{CSP}_B(G)$ is not Σ_2^1 -complete

Proof. (1) \Rightarrow (2) is an unpublished result of Louveau and also follows from Theorem 3.6.2; (3) \Rightarrow (4) \Rightarrow (5) are all clear; (5) \Rightarrow (2) is the theorem above. The equivalence of (1) and

(2) is essentially the classical Hell–Nešetřil theorem. For completeness, we sketch a proof below.

If G is bipartite, the core of G is K_2 , so G is tractable. If G is tractable, then it admits a cyclic polymorphism of all large enough prime arities. If G is not bipartite, then there is some closed walk of large prime length p , say (x_1, \dots, x_p) . If c is a cyclic polymorphism of arity p , then there is a loop at $c(x_1, \dots, x_p) = c(x_2, \dots, x_p, x_1)$, but G is supposed to be a simple graph. \square

A more general version for so-called smooth digraphs is given in Corollary 3.6.3. We can also generalize Todorčević and Vidnyánszky’s theorem on graph coloring to hypergraphs

Corollary 3.4.6. For any arity k and any number $n \geq 2$, the problem of Borel n -coloring k -ary hypergraphs is Σ_2^1 -complete.

Proof. The hypergraph coloring problem is equivalent to the CSP(\mathcal{D}) for $\mathcal{D} = (\{1, \dots, n\}, \{(x_1, \dots, x_k) : \neg(x_1 = \dots = x_k)\})$. These are known to be intractable. \square

And, we can find some exotic examples of Σ_2^1 -complete problems.

Corollary 3.4.7. The directed graph shown in Figure 3.1 has a Σ_2^1 -complete Borel CSP.

Proof. Barto, Kozik, Maróti, and Niven showed that this structure is intractable [BKM09]. \square

We give a last example relating to the theory of LCLs.

Definition 3.4.8. A **graph with parallels** is a graph G with a specified set of pairs of directed edges P . An **orientation** of a graph with parallels (G, P) is an orientation o of G so that $e \in o$ iff $f \in o$ whenever $(e, f) \in P$. An orientation is **balanced** if the indegree of every vertex is the same as the outdegree.




Figure 3.1: An intractable directed graph

Corollary 3.4.9. The set of Borel 4-regular graphs with parallels which admit a Borel balanced orientation is Σ_2^1 -complete.

Proof. Given a vertex v incident to (undirected) edges e_1, e_2, e_3, e_4 and an orientation o , let $f(e_i, v, o)$ be 1 if e_i is oriented toward v and -1 otherwise. Then o is balanced if and only if $\sum_i f(e_i, v, o) = 0$ for all v , and this is true even if the sum is taken mod 3. This gives us a way to code 3 variable linear equations with nonzero solutions, a problem which turns out to be NP-complete.

Consider $\mathbb{F}_3^\times(3)$ the structure on $\{-1, 1\}$ equipped with all relations of the form

$$a_1 z_1 + a_2 z_2 + a_3 z_3 + a_4 = 0$$

for $a_1, a_2, a_3, a_4 \in \{-1, 1\}$. It is straightforward to check that $\mathbb{F}_3^\times(3)$ is intractable (using, for instance, Schaefer's theorem). Write R_{a_1, a_2, a_3, a_4} for the relation $\{(x, y, z) : a_1 x + a_2 y + a_3 z + a_4 = 0\}$.

Given an instance \mathcal{X} of $\mathbb{F}_3^\times(3)$, define a graph with parallels $f(\mathcal{X})$ as follows. For a directed edge (x, y) , write $-1 \cdot (x, y)$ for (y, x) . Let T_4 be the infinite 4-regular tree, and fix some vertex $v_0 \in T_4$ incident to (directed) edges $e_i = (v_0, v_i)$ for $i = 1, 2, 3, 4$. Then the

graph of $f(\mathcal{X})$ is a large disjoint union of copies of T_4 :

$$V = \{(v, x) : v \in T_4, x \in X \text{ or } x \in R^{\mathcal{X}} \text{ for some } R\}$$

$$(v, x)E(v', x') : \Leftrightarrow vEv' \text{ and } x = x'.$$

And, the parallels of $f(\mathcal{X})$ are as follows:

1. For $e = (x_1, x_2, x_3) \in R_{a_1, \dots, a_4}^{\mathcal{X}}$, $((v_0, x_i), (v_1, x_i))$ is parallel to $a_i \cdot ((v_0, e), (v_i, e))$,
2. For $e \in R_{a_1, \dots, a_4}^{\mathcal{X}}$ and $e' \in R_{b_1, \dots, b_4}^{\mathcal{X}}$, $((v_0, e), (v_4, e))$ is parallel to $(a_4 b_4) \cdot ((v_0, e'), (v_4, e'))$

Suppose o is a balanced orientation of $f(\mathcal{X})$. Because of the parallels in item (2), for $e \in R_{a_1, \dots, a_4}^{\mathcal{X}}$, whether $((v_0, e), (v_4, e))$ is in o or not only depends on a_4 . Possibly replacing o with its opposite orientation, we may assume $((v_0, e), (v_4, e)) \in o$ if and only if $a_4 = 1$. For any edge e in the graph of $f(\mathcal{X})$ define

$$\tilde{g}(e) = \begin{cases} 1 & e \in o \\ -1 & \text{otherwise} \end{cases}$$

Then, I claim that \mathcal{X} has a solution given by $g(x) = \tilde{g}((v_0, x), (v_1, x))$. Indeed, if $e = (x_1, x_2, x_3) \in R_{a_1, a_2, a_3, a_4}^{\mathcal{X}}$, then by the parallels in item (1) above $g(x_i) = a_i \cdot \tilde{g}((v_0, e), (v_i, e))$. By our normalization, $\tilde{g}((v_0, e), (v_4, e)) = a_4$. And, since o is balanced, $\sum_i g((v_0, e), (v_i, e)) = 0$.

Conversely, if \mathcal{X} has a solution g , we get a partial balanced orientation o of $f(\mathcal{X})$ by setting

$$((v_0, x), (v_1, x)) \in o \Leftrightarrow g(x) = 1$$

for $x \in \mathcal{X}$, and for $e = (x_1, x_2, x_3) \in R_{a_1, \dots, a_4}^{\mathcal{X}}$ and $i = 1, 2, 3$,

$$((v_0, e), (v_i, e)) \in o \Leftrightarrow a_i g(x_i) = 1$$

and

$$((v_0, e), (v_4, e)) \in o \Leftrightarrow a_4 = 1.$$

Then o assigns an orientation to every edge which occurs in the set of parallels of $f(\mathcal{X})$, and every vertex is incident to either 1 or 4 edges assigned an orientation by 0. Then since the graph of $f(\mathcal{X})$ is acyclic and has a smooth connectedness relation, o extends to a balanced orientation on all of $f(\mathcal{X})$. \square

Of course all of the complexity in the above construction is carried by the parallels. We would like to compute the complexity of graphs with balanced orientations, or of Borel LCLs in general, but these are not in general exactly equivalent to CSPs. Indeed every NP-language is equivalent to an LCL.

Problem 3.4.10. Is the Borel version of every NP-complete LCL Σ_2^1 -complete?

3.5 Essentially classical CSPs

From the point of view of descriptive set theory, essentially classical structures are trivial. Interestingly, though, these structures turn out to be exactly the width 1 structures from computer science, i.e. those solved by arc-consistency. We first give a characterization of arc-consistency which is convenient for reflection arguments.

Definition 3.5.1. For an instance \mathcal{X} of \mathcal{D} , a say $f : X \rightarrow \mathcal{P}(D)$ is **closed** if,

$$(\forall (x_1, \dots, x_n) \in R^{\mathcal{X}}, i \leq n, a \in D) \quad \bigwedge_{e \in \pi_i^{-1}(a)} e \notin R \setminus \Pi_j f(x_j) \Rightarrow a \in f(x_i).$$

The **closure** of f is function \bar{f} obtained by iteratively adding points to each $f(x)$ to satisfy the above condition.

A **good witness** for \mathcal{X} is a closed function f so that, for all x , $f(x) \neq D$.

So, a structure is width 1 iff every structure with a good witness has a solution. And, \mathcal{X} has a good witness iff the closure of \emptyset is a good witness. Recall Theorem 3.2.11 says that a structure is width 1 if and only if it has a totally symmetric polymorphism of arbitrarily high arity.

Theorem 3.5.2. If \mathcal{D} is a finite width 1 relational structure, then $\text{CSP}_B(\mathcal{D})$ is essentially classical and effectivizable.

Proof. Suppose \mathcal{X} is an arc-consistent Δ_1^1 instance of \mathcal{X} . We first show that \mathcal{X} has a Δ_1^1 1-minimal witness, then we mimic the proof that \mathcal{X} has a solution to get a Δ_1^1 solution.

Write $\Phi(f)$ to mean that \bar{f} is a good witness for \mathcal{X} . Note that, if f is Σ_1^1 , so is \bar{f} . Also Φ is Π_1^1 on Σ_1^1 , so the first reflection [Mar19, Theorem 2.27] theorem says any Σ_1^1 set satisfying Φ is contained in a Δ_1^1 set satisfying Φ . Since taking the closure and applying the reflection theorem can be done uniformly in the codes we have an increasing Δ_1^1 sequence of sets

$$f_0 = \emptyset \subseteq \bar{f}_0 \subseteq f_1 \subseteq \bar{f}_1 \subseteq f_2 \dots$$

where each A_i is Δ_1^1 and satisfies Φ . Then $\tilde{f} := \bigcup_i f_i$ is a Δ_1^1 good witness for \mathcal{X} .

Let N be the largest arity of a relation in \mathcal{D} and let T be a totally symmetric polymorphism of \mathcal{D} of arity $n \geq |\mathcal{D}| \times N$. By the Lusin–Novikov theorem, there is some $f : \mathcal{X} \rightarrow \mathcal{D}^n$ so that

$$D \setminus \tilde{f}(x) = \{f(x)_1, \dots, f(x)_n\}.$$

I claim that $T \circ f$ is a solution to \mathcal{X} . To see this, suppose $(x_1, \dots, x_k) \in R^{\mathcal{X}}$, let M be an $k \times n$ matrix whose columns are in R and whose rows enumerate U_{x_i} , and let σ_i be the i^{th} row of M . This possible since $n \geq |D| \times k$ and \tilde{f} is a good witness. Then by total symmetry $T \circ f(x_i) = T(\sigma_i)$ and since T is a polymorphism, $(T \circ f(x_1), \dots, T \circ f(x_k)) \in R$. \square

The converse requires a little more work. A theorem of Feder and Vardi says that a structure is width 1 if and only if it admits a basis of acyclic structures. So, if \mathcal{D} is not width 1, there is some unsolvable instance \mathcal{X} where every acyclic lift of \mathcal{X} has a solution. We sharpen this a bit to find an instance \mathcal{X} with a point x so that any acyclic lift of \mathcal{X} has a solution, but there is an acyclic lift which has a solution that is constant on the fiber over x . Then a modification of the G_0 construction gives an acyclic lift of \mathcal{X} with no Borel solution.

Definition 3.5.3. A **simple path** in \mathcal{X} is a path $(x_1, (e_1, R_1), \dots, (e_{n-1}, R_{n-1}), x_n)$ so that $(x_1, (e_1, R_1), x_2, (e_2, R_2), \dots, x_{n-1}, (e_n, R_n))$ is an injective sequence.

A **cycle** is a simple path $(x_1, (e_1, R_1), x_2, (e_2, R_2), \dots, x_n)$ with $x_1 = x_n$. A structure is **acyclic** if it does not contain any cycles.

A **lift** of \mathcal{X} is a structure \mathcal{Y} with a homomorphism $f : \mathcal{Y} \rightarrow \mathcal{X}$.

Note that if we impose any unary constraints on an acyclic structure it remains acyclic. The following lemma is essentially equivalent to tree duality for width 1 structures [FV98]. We include a proof for completeness

Lemma 3.5.4. An instance \mathcal{X} of \mathcal{D} is arc-consistent if and only if any acyclic lift of \mathcal{X} has a solution.

Proof. First suppose that \mathcal{X} is arc-consistent as witnessed by predicates $U_x, U_{x_0}(a), f : \mathcal{Y} \rightarrow \mathcal{X}$ is an acyclic lift of \mathcal{X} . If $A \subseteq Y$ is connected, i.e. there is a simple path between any two points in A , and $R(y_1, \dots, y_n)$ holds with some $y_i \in A$ and $y_j \notin A$, then in fact $\{y_1, \dots, y_n\}$ contains only one element of A and no other relations of \mathcal{Y} meet $A \cup \{y_1, \dots, y_n\}$ at some element of $\{y_1, \dots, y_n\}$. Otherwise would be able to find a cycle in \mathcal{Y} . So, by arc-consistency, any partial solution to \mathcal{Y} with connected domain has an extension and a total solution must exist.

For the other direction, we build acyclic structures which encode the steps in the obvious algorithm for checking arc-consistency. Fix a linear order on pairs (x, e, R) where x is a coordinate of e and $e \in R^{\mathcal{X}}$. Say (x, e, R) is bad for $\langle U_x : x \in \mathcal{X} \rangle$ if x is the i^{th} coordinate of $e = (x_1, \dots, x_n)$ and $U_x \neq \pi_i(R \cap U_{x_1} \times \dots \times U_{x_n})$. For $i \in \mathbb{N}$ and $x \in \mathcal{X}$ define U_x^i inductively as follows:

1. $U_x^0 = D$ for all x
2. If there is no (x, e, R) bad for $\langle U_x^i : x \in \mathcal{X} \rangle$ set $U_x^{i+1} = U_x^i$ for all x

3. If (x, e, R) is the least bad triple set $U_{x'}^{i+1} = U_x^i$ for $x \neq x'$ and set $U_x^{i+1} = \pi_i(R \cap U_{x_1} \times \dots \times U_{x_n})$.

By construction, \mathcal{X} is not arc-consistent if and only if $U_x^i = \emptyset$ for some i and some x . We will show by induction that there are acyclic lifts $f_x^i : \mathcal{Y}_x^i \rightarrow \mathcal{X}$ and points $y_x^i \in (f_x^i)^{-1}(x)$ so that

$$\{g(y_x^i) : g \text{ is a solution to } \mathcal{Y}_x^i\} = U_x^i.$$

For the base case, let \mathcal{Y}_x^0 and have domain $\{y_x^i\}$, have $R^{\mathcal{Y}_x^0}$ empty if R has arity greater than 1, and $U(y_x^0) \Leftrightarrow U(x)$ for all unary relations, and set $f_x^i(y_x^i) = x$. If i is as in case (2) of the induction or if x is not in the least bad triple at step i , set $\mathcal{Y}_x^{i+1} = \mathcal{Y}_x^i$. Otherwise, suppose $(x, (x_1, \dots, x_n), R)$ is the least bad tuple at step i , $x = x_k$, and define

$$\mathcal{Y}_x^{i+1} = \bigsqcup_{j=1}^n \mathcal{Y}_{x_j}^i$$

with relations

$$R^{\mathcal{Y}_x^{i+1}} = \{(y_{x_1}^i, \dots, y_{x_n}^i)\} \cup \bigsqcup_{j=1}^n \mathcal{R}^{\mathcal{Y}_{x_j}^i}$$

and $S^{\mathcal{Y}_x^{i+1}}$ is just the union of the S relations from each $\mathcal{Y}_{x_j}^i$ for any relation S besides R . Also, put $f_x^{i+1}(y) = f_{x_j}^i(y)$ if $y \in \mathcal{Y}_{x_j}^i$ and put $y_x^{i+1} = y_x^i$. Then, g is a solution \mathcal{Y}_x^{i+1} if and only if $g \upharpoonright \mathcal{Y}_{x_j}^i$ is a solution to $\mathcal{Y}_{x_j}^i$ for each j and $(g(y_{x_1}^i), \dots, g(y_{x_n}^i)) \in R$. So by the inductive hypothesis

$$\begin{aligned} \{g(y_x^{i+1}) : g \text{ is a solution to } \mathcal{Y}_x^{i+1}\} &= \pi_k(R \cap \prod_j U_{x_j}^i) \\ &= U_x^{i+1} \end{aligned}$$

In particular, if \mathcal{X} is not arc-consistent, then $U_x^i = \emptyset$ for some i, x , so some acyclic lift of \mathcal{X} has no solution. \square

To make sure the lift we construct in the next lemma is acyclic, we will need to introduce some new, simply definable relations. This is no issue since, whenever \mathcal{D} is width 1 or essentially classical, so is any structure which is simply definable in \mathcal{D} .

Definition 3.5.5. For a structure \mathcal{D} , let $\langle\langle\mathcal{D}\rangle\rangle$ be the structure with the same domain as \mathcal{D} equipped with every relation simply definable in \mathcal{D} .

Note that $\langle\langle\mathcal{D}\rangle\rangle$ will not be finite even if \mathcal{D} is finite. But since it is simply definable in \mathcal{D} , most theorems about finite structures will apply.

Lemma 3.5.6. If there is an instance of \mathcal{D} which is arc-consistent but not cycle-consistent, then there is a finite arc-consistent instance \mathcal{X} of $\langle\langle\mathcal{D}\rangle\rangle$ with a finite acyclic lift $f : \mathcal{Y} \rightarrow \mathcal{X}$ and some $x \in \mathcal{X}$ so that \mathcal{Y} has no solution which is constant on $f^{-1}(x)$.

Proof. Similar to the arc-consistency algorithm sketched above, we can test cycle-consistency of an instance \mathcal{X} with arc-consistency witnesses U_x by iteratively going through each closed path $P = (x_1, (e_1, R_1), \dots, x_n, (e_n, R_n), x_1)$ with coordinates $(j_1, k_1, \dots, j_n, k_n)$, imposing a new unary constraint $U_P(x_1)$ on x_1 with

$$U_P(a_1) := (\exists a_2, a_3, \dots, a_n) \bigwedge_k (U_{x_k}(a_k) \wedge \pi_{i_k, j_k} R_k(a_k, a_{k+1})),$$

and then refining the witnesses to arc-consistency. Stopping this process one step early we can assume \mathcal{X} is arc-consistent but $\mathcal{X} \cup \{U_P\}$ is not arc-consistent.

By the previous lemma, there is an acyclic lift $f : \mathcal{Y} \rightarrow \mathcal{X} \cup U_P$ with no solution. We can convert \mathcal{Y} into an acyclic lift $f' : \mathcal{Y}' \rightarrow \mathcal{X}$ with no solution which is constant on $f^{-1}(x_1)$ as follows. Note that, since $U_P^{\mathcal{X}} = \{x_1\}$, $U_P^{\mathcal{Y}} \subseteq f^{-1}(x)$. For each $y \in U_P^{\mathcal{Y}}$ introduce new variables, $z_{2,y}, z_{3,y}, \dots, z_{n,y}$, and impose constraints $\pi_{i_k, j_k} R_k(z_{k,y}, z_{k+1,y})$, $U_{x_1}(y)$, and $U_{x_k}(z_k)$. And, extend f to f' by setting $f(z_{k,y}) = x_k$ for each new variable $z_{k,y}$. Then \mathcal{Y}' is an acyclic lift of \mathcal{X} . If g is a solution to \mathcal{Y}' which is constant on $f^{-1}(x)$, then whenever $y \in U_P^{\mathcal{Y}}$, $g(z_{2,y}), \dots, g(z_{n,y})$ witness that $g(y) \in U_P$, so g restricts to a solution to \mathcal{Y} , which is a contradiction. \square

Now with \mathcal{X} and \mathcal{Y} as above, a simple modification of the G_0 construction gives an acyclic lift \mathcal{Y}_0 of \mathcal{X} which contains many copies of \mathcal{Y} so that any Borel (in fact any Baire measurable) map $g : \mathcal{Y} \rightarrow \mathcal{D}$ must be constant on the fiber of y in some copy of \mathcal{Y} .

Theorem 3.5.7. If \mathcal{D} is finite and not width 1, then there is a Borel instance of \mathcal{D} with a solution but no Baire measurable solution.

Proof. We may assume $\mathcal{D} = \langle\langle\mathcal{D}\rangle\rangle$. Either \mathcal{D} is bounded width or, by Theorem 3.2.13, \mathcal{D} simply constructs $\mathbb{F}(3)$ for some finite field \mathbb{F} . In either case, there is an arc-consistent but not cycle-consistent instance of \mathcal{D} . Then by the previous lemma, we can find \mathcal{X} arc-consistent, $f : \mathcal{Y} \rightarrow \mathcal{X}$ a finite acyclic lift, and $x \in \mathcal{X}$ so that \mathcal{Y} has no solution which is constant on $f^{-1}(x)$.

Let $f^{-1}(x) = \{x_1, \dots, x_n\}$ and define an n -ary relation R by

$$R(a_1, \dots, a_n) :\Leftrightarrow x_i \mapsto a_i \text{ extends to a solution to } \mathcal{Y}.$$

Then, R is simply definable in \mathcal{D} , any acyclic instance of R has a solution, and there is no solution to any instance of R which constant on any tuple of R .

Fix a sequence $\langle\sigma_i : i \in \mathbb{N}\rangle$ with $\sigma_i \in \{1, \dots, n\}^i$ so that every string in $\{1, \dots, n\}$ extends to some σ_i . Let \mathcal{Y}_0 be the instance of R with domain $\{1, \dots, n\}^\omega$ and

$$R^{\mathcal{Y}_0}(s_1, \dots, s_n) :\Leftrightarrow (\exists 1 \leq i \leq n, t \in \{1, \dots, n\}^\omega) \bigwedge_{j=1}^n s_j = \sigma_i \hat{\wedge} j \hat{\wedge} t.$$

Since \mathcal{Y}_0 is acyclic it has a solution. But suppose $g : \mathcal{Y}_0 \rightarrow \mathcal{D}$ is Baire measurable. There is some $a \in \mathcal{D}$ with $g^{-1}(a)$ nonmeager. Then $g^{-1}(a)$ is comeager in some basic neighborhood N_{σ_i} . Since the maps which cycle the $(i+1)^{th}$ coordinate of a sequence are self-homeomorphisms of N_{σ_i} , there is some t so that $(\sigma_i \hat{\wedge} j \hat{\wedge} t) \in g^{-1}(a)$ for $j = 1, \dots, n$. But then g is constant on some tuple which is a contradiction. \square

Corollary 3.5.8. The set of essentially classical structures is decidable.

Corollary 3.5.9. If \mathcal{D} is essentially classical it is effectivizable.

Proof. If \mathcal{D} is essentially classical, then by the above theorem \mathcal{D} admits a totally symmetric polymorphism. So by Theorem 3.5.2, \mathcal{D} is effectivizable. \square

3.6 Effectivizable CSPs

In this section, we give a number of examples of effectivizable structures. In particular we show that any structure with a dual discriminator polymorphism is effectivizable. This is modest progress, but it is enough to compute the complexity of any so-called smooth directed graph and any Boolean structure except $\mathbb{F}_2(3)$.

Proposition 3.6.1 (Folklore). Suppose that \mathcal{E} has domain D and a dual discriminator polymorphism. Then \mathcal{E} is simply definable in the structure \mathcal{D} with the following relations:

- Every unary predicate
- For $a, b \in A$, each predicate

$$R_{a,b}(x, y) :\Leftrightarrow x = a \vee y = b$$

- For $f \in \text{Sym}(A)$, each predicate

$$R_f(x, y) :\Leftrightarrow y = f(x)$$

Proof. Let $T : D^3 \rightarrow D$ be the dual discriminator. It is straightforward to check these relations are all preserved by T . Suppose that $R \subseteq D^n$ is preserved by T . We first show that R is a conjunction of binary predicates. Suppose that $\pi_{i,j}(a) \in \pi_{i,j}(R)$. We check by induction on $|J|$ that $\pi_J(a) \in \pi_J(R)$ for any $J \subseteq \{1, \dots, n\}$. Pick some J with $|J| = \ell + 1$ and pick $J_1, J_2, J_3 \subseteq J$ distinct with $|J_1| = |J_2| = |J_3| = \ell$. By induction, there are b_1, b_2 , and $b_3 \in R$ with $\pi_{J_i}(b_i) = \pi_{J_i}(a)$. Then, $b := T(b_1, b_2, b_3) \in R$ and for any $j \in J$, there are at least 2 values of i with $j \in J_i$, meaning the majority of b_1, b_2, b_3 agree with a in coordinate j . Thus $\pi_J(a) = \pi_J(b)$.

Now suppose that R is a binary relation which is invariant under T . We show that it is a conjunction of relations of the above form. Let $A = \pi_2(R)$ and $B = \pi_1(R)$. Note that if

$a \in D$ can be paired with two different elements b, b' so that $(a, b), (a, b') \in R$, then for any $(c, d) \in R$, we have

$$d((c, d), (a, b), (a, b')) = (a, d) \in R.$$

So if a can be paired with two different elements of B , then a can be paired with anything in B . Thus R is either $A \times B$ or the intersection of $A \times B$ with a relation of the form R_π or $R_{a,b}$. \square

Effectivization for $\text{CSP}_B(\mathcal{D})$ follows fairly easily from the main theorem of [Tho21].

Theorem 3.6.2. If \mathcal{D} has a dual discriminator polymorphisms, then $\text{CSP}_B(\mathcal{D})$ is effectivizable.

Proof. We may assume that \mathcal{D} is of the form indicated in the previous proposition with domain $D = \{1, \dots, n\}$. Fix a Δ_1^1 instance \mathcal{X} of \mathcal{D} . For $f \subseteq \mathcal{N} \times D$, write $\Phi(f)$ to mean

1. f is a partial function, i.e. $(\forall x, y, z) \neg((x, y), (x, z) \in f, y \neq z)$
2. For all $x \in \mathcal{X}$, unary predicates U with $U^{\mathcal{X}}(x)$, and $d \notin U$, $(x, d) \notin f$
3. For all $x, y \in \mathcal{X}$ and $a, b \in D$ with $(x, y) \in R_{a,b}^{\mathcal{X}}$, if there is $c \neq a$ with $(x, c) \in f$ then $(y, b) \in f$ (and likewise if $(y, c) \in f$ for some $c \neq b$ then $(x, a) \in f$)
4. For all $x, y \in X$ and $g \in \text{Sym}(D)$ with $(x, y) \in R_f^{\mathcal{X}}$, if $(x, a) \in f$, then $(y, g(a)) \in f$ (and likewise, if $(y, a) \in f$, then $(x, g^{-1}(a)) \in f$).

I claim that \mathcal{X} has a (Δ_1^1) solution if and only if there is a sequence of (Δ_1^1) sets $\langle f_i : i \in \mathbb{N} \rangle$ such that $\Phi(f_i)$ for all i and $\mathcal{N} = \bigcup_i \text{dom}(f_i)$. Then since the properties above are all closure and independence properties the theorem follows by [Tho21, Theorem 3.6].

If f is a solution to \mathcal{X} then f satisfies all of the above properties and $\text{dom}(f) = \mathcal{N}$. Conversely, suppose $\langle f_i : i \in \mathbb{N} \rangle$ is such a sequence and define

$$n(x) := \min\{i : x \in \text{dom}(f_i)\}, \quad f(x) := f_{n(x)}(x).$$

We check that f preserves all of the relations in \mathcal{D} .

- If $U^x(x)$ for some unary U , then by property (2) $U(f_i(x))$ holds for all i with $x \in \text{dom}(f_i)$. Thus $U(f(x))$.
- If $R_{a,b}^x(x, y)$, then suppose without loss of generality $n(x) \leq n(y)$. If $f_{n(x)}(x) = a$, then $f(x) = a$. If $f_{n(x)}(x) \neq a$, then by property (3) $(y, b) \in f_{n(x)}$, so $n(x) = n(y)$ and $f(y) = b$. In either case $R_{a,b}(f(x), f(y))$.
- If $R_g^x(x, y)$, then suppose without loss of generality $n(x) \leq n(y)$. If $(x, a) \in f_i$, then by property (4), $(y, g(a)) \in f_i$, so $n(y) = n(x)$ and $R_g(f(x), f(y))$.

□

We can now generalize Corollary 3.4.5 to so-called smooth digraphs. If we could generalize this to all directed graph, then we would have Conjecture 3.1.10 [BDJ15].

Corollary 3.6.3. If \mathcal{D} is a directed graph with no sources or sinks (these are sometimes called smooth digraphs), then $\text{CSP}_B(\mathcal{D})$ is effectivizable if and only if it is Π_1^1 if and only if \mathcal{D} is tractable.

Proof. By a theorem of Barto, Kozik, and Niven [BKN09], for such graphs either \mathcal{D} is intractable or the core of \mathcal{D} is a disjoint union of directed cycles. A disjoint union of cycles is the graph of a permutation, so any tractable smooth digraph admits a dual discriminator polymorphism. □

We can also compute the complexity of most Boolean structures.

Corollary 3.6.4 (c.f. Schaefer's theorem [Sch78]). If \mathcal{D} is a structure on $\{0, 1\}$, one of the following holds:

1. \mathcal{D} has a totally symmetric polymorphism and $\text{CSP}(\mathcal{D})$ is essentially classical

2. \mathcal{D} is pp constructible in 2SAT and $\text{CSP}_B(\mathcal{D})$ is effectivizable
3. \mathcal{D} is intractable and $\text{CSP}_B(\mathcal{D})$ is Σ_2^1 -complete, or
4. \mathcal{D} is pp constructible in \mathbb{F}_2 and vice versa.

Proof. We'll take the opportunity to use a bit of a sledgehammer here. In the 1940's Post classified all clones on $\{0, 1\}$, see [Pos41]. By inspecting the minimal elements in the lattice of clones, one can see that \mathcal{D} falls into one of the following cases:

1. $\text{Pol}(\mathcal{D})$ contains a constant function, \vee , or \wedge
2. $\text{Pol}(\mathcal{D})$ contains the majority function
3. $\text{Pol}(\mathcal{D}) \subseteq \langle \neg \rangle$ (the algebra generated by the negation function)
4. $\text{Pol}(\mathcal{D})$ is one of the following: $\langle x \oplus y \oplus z \rangle$ or $\langle x \oplus y \oplus z, \neg \rangle$.

All of the operations in the first case are totally symmetric. In the second case, $\text{CSP}(\mathcal{D})$ is in fact pp definable in 2SAT. Note that $\langle \neg \rangle = \text{Pol}(N)$, where N is the not-all-equal predicate. So in the third case, \mathcal{D} is intractable. And either of the algebras in the last class correspond to structures which are equivalent \mathbb{F}_2 . □

And, we have one last example.

Theorem 3.6.5. Let \mathcal{D} be the structure on $\{r, p, s\}$ equipped with all relations which are preserved by the rock-paper-scissors operation, (\star) (see item (6) after Definition 3.2.1). Then Borel solutions to locally countable instances of \mathcal{D} are effectivizable.

Proof. \mathcal{D} is generated under pp definitions by the following relations:

- R_π , the graph of the cyclic permutation $\pi = (rps)$.
- $R_\star(x, y, z) :\Leftrightarrow x \in \{p, s\} \wedge (x = s \vee y = z)$.

Let \mathcal{X} be a locally countable Δ_1^1 instance of \mathcal{D} with a Borel solution \tilde{f} . Let \mathcal{A} be the 1-minimal closure of \emptyset . Since \mathcal{X} is locally countable, \mathcal{A} is Δ_1^1 . Let $U_x = \{d : (x, d) \notin \mathcal{A}\}$. Note that U_x is a witness to arc-consistency for \mathcal{X} .

Say that x is fixed if $|U_x| = 1$ and critical if $|U_x| = 2$ and free otherwise. Let G be the weighted directed Borel graph with edges (x, y) and (y, x) of weight 0 whenever $R(z, x, y)$ with z fixed and $p \in U_z$, edges (x, y) of weight 1 whenever $R_\pi(x, y)$, and edges (y, x) of weight -1 whenever (x, y) is an edge of weight 1. Then, let $d(x, y)$ be the sum of weights along a directed path from x to y modulo 3 if such a path exists and ∞ otherwise. This is well defined since \mathcal{X} must be cycle-consistent. Note that if $d(x, y) < \infty$ and x is free (or critical or fixed) then so is y . If $d(x, y) = i < \infty$ and f is a solution to \mathcal{X} , we must have $f(y) = \pi^i(x)$.

Define f as follows:

- for x fixed, set $f(x) = a$ if $a \in U_x$
- for x critical with $U_x = \{a, b\}$, set $f(x) = a \star b$

If $d(x, y) = i$ and $f(x) = a$ (in particular x is not free), then $U_y = \pi^i(U_x)$, so $f(y) = \pi^i(x)$. And, if $R_\star(x, y, z)$, we must have that x is not free. If x is critical, then $U_x = \{p, s\}$, so $f(x) = s$ and this instance of R_\star is satisfied by any extension of f (in particular f is a partial homomorphism). If x is fixed, then this instance of R_\star will be satisfied by an extension g of f if and only if $g(y) = g(z)$.

So, we want to find a Δ_1^1 function g on the free variables so that $d(x, y) = i$ implies $g(y) = \pi^i(g(x))$. We know that we have a Borel such function (namely the restriction of \tilde{f} to the free variables). And, such functions are effectivizable by the previous theorem.

□

This last example is archetypal of bounded width structures, which can be solved by an intricate greedy algorithm. However Problem 3.1.14 remains open even if we restrict to

locally countable instances. We end with a conditional result:

Proposition 3.6.6. If every bounded width structure is effectivizable and every structure of the form $\mathbb{F}(3)$ is Σ_2^1 -complete, then the following are equivalent:

1. $\text{CSP}_B(\mathcal{D})$ is effectivizable
2. $\text{CSP}_B(\mathcal{D})$ is Π_1^1
3. $\text{CSP}_B(\mathcal{D})$ is not Σ_2^1 -complete

Proof. (1) \Rightarrow (2) \Rightarrow (3) is clear. If (3) holds, then $\text{CSP}_B(\mathcal{D})$ must not pp construct $\mathbb{F}(3)$ for any finite field \mathbb{F} . In particular, $\text{Pol}(\mathbb{F}(3)) \notin HS(\text{Pol}(\mathcal{D}))$, but then \mathcal{D} is bounded width, and (1) holds. \square

It is not clear how plausible the assumptions in this proposition are, but the resulting equivalences are appealing.

APPENDIX A

Codings and edge coloring

In this appendix, we prove that the set of Borel edge 3-colorable graphs is Σ_2^1 -complete. Since edge coloring involves a restricted class of instances, this does not follow from Theorem 3.4.3, but the classical proof that edge 3-coloring is NP-complete still adapts to the Borel setting. We also take this as an opportunity to show in detail how to verify that a construction is Δ_1^1 in the codes.

First, we recall the classical NP-completeness proof. Roughly, we will reduce 3SAT to 3 edge coloring by coding variables as pairs of edges, and the values that the variable can take will be coded into whether the corresponding edges can receive the same color. We need some lemmas.

Lemma A.0.1 (Inverter lemma). There is a graph I with distinguished edges a, b, c, d, e so that a 3-coloring f of a, b, c, d, e extends to a coloring of H if and only if one of the following holds:

$$f(a) = f(b) \quad \text{and} \quad f(e) \neq f(c) \neq f(d) \neq f(e)$$

or

$$f(c) = f(d) \quad \text{and} \quad f(e) \neq f(a) \neq f(b) \neq f(e)$$

Proof. Such a graph is pictured below in Figure A.1.

□

We will represent H diagrammatically as in the right image of Figure A.1. In the following, we will refer to a, b, c, d , and e as **coding edges**. And, we will not include the degree 1




Figure A.1: An inverter component and diagram representation

vertices when we refer to vertices of H . (The coding edges will connect to vertices in other components).

Lemma A.0.2 (Variable setting lemma). For any n , there is a graph V_n with n pairs of distinguished edges $e_{i,0}, e_{i,1}$ so that a 3 coloring f of $e_{1,0}, e_{1,1}, \dots, e_{n,0}, e_{n,1}$ extends to an edge 3-coloring of V_n if and only if one of the following holds

$$(\forall i) f(e_{i,0}) = f(e_{i,1})$$

or

$$(\forall i) f(e_{i,0}) \neq f(e_{i,1})$$

Proof. An example of S_4 is drawn in Figure A.2, with the $e_{i,j}$ s being the eight edges connected to vertices of degree one. One can build S_n for general n by following a similar pattern (or chaining together copies of S_4 .) \square

As above, we will refer to the $e_{i,j}$ s as coding edges and ignore the degrees one vertices in the following.

Lemma A.0.3 (Or gate lemma). There is a graph O with three distinguished pairs of edges $e_{i,j}$ for $i = 1, 2, 3$ and $j = 0, 1$ so that a 3 coloring f of the $e_{i,j}$ extends to a coloring of O if and only if $f(e_{i,0}) = f(e_{i,1})$ for some i .




Figure A.2: A variable setting component

Proof. Such a graph is shown in Figure A.3, where the $e_{i,j}$ s are the 6 edges along the bottom of the image.

□




Figure A.3: An or gate

Once again, we will refer to the $e_{i,j}$ s as coding edges and ignore the degree one vertices incident to coding edges.

Combining these lemmas, the reduction from 3SAT is straightforward

Theorem A.0.4. There is a polynomial time reduction from 3SAT to 3 edge coloring

Proof. Given an instance \mathcal{X} of 3SAT in CNF, build a graph $g(\mathcal{X})$ with one copy of S_n for each variable which appears in n disjunctions, one copy of O for each disjunction of \mathcal{X} . Wire each of the outgoing pairs of edge of each copy of S_n into the appropriate copy of O , passing

them through a copy of H if the corresponding variable appears negated in the disjunction. See Figure A.4 for an example.

By construction, an edge 3-coloring of $g(\mathcal{X})$ yields a satisfying assignment of \mathcal{X} by assigning a variable True if the coloring agree on the outgoing pairs of edges in its copy of S_n and False otherwise. And any satisfying assignment gives a partial coloring of the edges by the same scheme, which then extends an edge 3-coloring.

□




Figure A.4: The reduction applied to $(\neg v_1 \vee v_2 \vee \neg v_3) \wedge (v_1 \vee \neg v_2 \vee v_3)$

We will verify that the same construction can be carried out for locally finite instances of Borel 3SAT. First, we fix a coding for Borel sets. The important point is Lemma A.0.7, which lets us easily check that a construction is Δ_1^1 in the codes.

Theorem A.0.5. There is a good ω -parameterization of Π_1^1 , i.e. a Π_1^1 set $U \subseteq \omega \times \mathcal{N}$ so that

1. For every Π_1^1 $P \subseteq \mathcal{N}$ there is an e so that $P = U_e = \{x \in \mathcal{N} : (e, x) \in U\}$.

2. For every Π_1^1 $P \subseteq \omega^k \times \mathcal{N}$ there is a recursive function $S : \omega^k \rightarrow \omega$ so that, for $n \in \omega^k$,

$$(n, x) \in P \Leftrightarrow (S(n), x) \in U$$

Fix such a good ω -parameterization U .

Definition A.0.6. Fix a Δ_1^1 set B . A **simple code** for B is a pair $\langle e, i \rangle \in \omega^2$ so that

$$U_e = \mathcal{N} \setminus U_i = B$$

for some good U .

A **nice coding** is a triple $(\mathcal{C}, \mathcal{D}^\Pi, \mathcal{D}^\Sigma)$ where

1. $\mathcal{C} \subseteq \omega$ is a Π_1^1 set, referred to as the codes
2. $\mathcal{D}^\Pi, \mathcal{D}^\Sigma \subseteq \omega \times \mathcal{N}$, \mathcal{D}^Π is Π_1^1 , and \mathcal{D}^Σ is Σ_1^1
3. For every $e \in \mathcal{C}$, $\mathcal{D}_e^\Pi = \mathcal{D}_e^\Sigma$
4. For every Δ_1^1 set $B \subseteq \mathcal{N}$ there is a code $e \in \mathcal{C}$ with $B = D_e^\pi = D_e^\Sigma$
5. There are recursive functions f, g so that, for every B ,

$$e \text{ is a nice code for } B \rightarrow f(e) \text{ is a simple code for } B$$

$$\langle e, i \rangle \text{ is a simple code for } B \rightarrow g(e, i) \text{ is a nice code for } B$$

One can obtain nice codes from the simple codes by a uniform application of separation.

See [Mos09, Section 3.3]

We'll fix simple and nice codings for $\Delta_1^1(\mathcal{N}^k)$ for all k , and for $e \in \mathcal{C}$ write \mathcal{D}_e for \mathcal{D}_e^Π .

Lemma A.0.7. Fix a Δ_1^1 linear order \preceq of \mathcal{N} . The following maps are Δ_1^1 in the codes:

1. $(A, B) \mapsto A \cup B$

2. $(A, B) \mapsto A \cap B$
3. $(A, B) \mapsto A \times B$
4. $A \mapsto \text{dom}(A)$, where A is a relation with countable sections.
5. $(A, f) \mapsto f(A)$, where f is a countable-to-one function
6. $A \mapsto f$ where A is a relation with finite sections, and

$$f(x, i) = y \Leftrightarrow y \text{ is the } i^{\text{th}} \text{ element of } R_x \text{ according to } \preceq$$

That is, there is a function $f : \omega^2 \rightarrow \omega$ so that, if e, i are nice codes for A and B , then $f(e, i)$ is a nice code for $A \cup B$ (and likewise for the other items).

Proof. We'll prove (1) and (4) and sketch (6), the other items are easy once you've seen the general idea.

Using item (5) of the definition of nice codes, it suffices to show that, for our good ω -parameterization U , there are Δ_1^1 functions P, S so that, for any $e, i, j, k \in \omega$, if

$$A = U_e = \mathcal{N} \setminus U_i \quad B = U_j = \mathcal{N} \setminus U_k$$

then

$$U_{P(e,i,j,k)} = \mathcal{N} \setminus U_{S(e,i,j,k)} = A \cup B.$$

Define

$$R_P(e, i, j, k, x) :\Leftrightarrow (e, x) \in U \text{ or } (j, x) \in U$$

$$R_S(e, i, j, k, x) :\Leftrightarrow (i, x) \in \omega \times \mathcal{N} \setminus U \text{ or } (k, x) \in \omega \times \mathcal{N} \setminus U.$$

Then R_P is Π_1^1 and R_S is Σ_1^1 . By item (2) of the definition of good ω -parameterizations, there are recursive functions P, S so that, for any $e, i, j, k \in \omega$ and $x \in \mathcal{N}$

$$R_P(e, i, j, k, x) \Leftrightarrow P(e, i, j, k) \in U$$

$$R_S(e, i, j, k, x) \Leftrightarrow S(e, i, j, k) \notin U$$

as desired.

Similarly, for item (4) we want Δ_1^1 functions S, P so that, if $\langle e, i \rangle$ is a simple code for a relation $A \subseteq \mathcal{N}^2$ with countable sections, then $\langle S(e, i), P(e, i) \rangle$ is a simple code for the domain of A . Define

$$R_P(e, i, x) \Leftrightarrow (\exists y \in \Delta_1^1(x)) (e, x, y) \in U_e$$

$$R_S(e, i, x) \Leftrightarrow (\exists y) (e, x, y) \in U_e$$

Note that R_P is Π_1^1 and R_S is Σ_1^1 . By the effective perfect set theorem, if A is relation with countable sections, and $\langle e, i \rangle$ is a simple code for A , then for any x , $R_P(e, i, x) \Leftrightarrow R_S(e, i, x) \Leftrightarrow x \in \text{dom}(A)$. So, again we can find recursive functions as desired.

For (6), we use the following definition of f : $f(x, i) = y$ if and only if

$$(\exists y_1, \dots, y_i \in R_x) [y_1 \prec \dots \prec y_i = y]$$

and

$$(\forall y_1, \dots, y_{i+1} \in R_x) [y_1 \prec \dots \prec y_{i+1} \Rightarrow y_{i+1} \neq y].$$

And, either quantifier can be taken to range over $\Delta_1^1(x)$. □

This lemma can be usually be used a black box without the need to delve into the details of the coding.

Definition A.0.8. \mathbf{E} is the set of (nice) codes for Borel graphs with Borel edge 3-colorings, where we view a graph as a vertex set $V \subseteq \mathcal{N}$ and a symmetric subset of edges $E \subseteq V^2$. That is \mathbf{E} is

$$\{\langle v, e \rangle \in \mathcal{C}^2 : D_e \subseteq D_v^2 \text{ and } D_e \text{ is symmetric and Borel edge 3-colorable}\}$$

$\text{CSP}_B^{lf}(3\text{SAT})$ is the set of codes for locally finite Borel instances of 3SAT with Borel satisfying assignments.

It follows from the comments following 3.4.3 that $\text{CSP}_B^{lf}(\text{3SAT})$ is Σ_2^1 -complete.

Theorem A.0.9. $\text{CSP}_B^{lf}(\text{3SAT}) \leq_B \mathbf{E}$.

Proof. The construction in Theorem A.0.4 works in the Borel setting. We will describe this construction formally so that it is clear that $G_{\mathcal{X}}$ is generated from \mathcal{X} using operations from Lemma A.0.7, and then use the Luzin–Novikov theorem to verify that this gives a reduction.

Fix an locally finite Borel instance \mathcal{X} of 3SAT in CNF with variables V and constraints $C = \{c_1, \dots, c_m\}$. We may assume no variable shows up twice in any c_i . We define the following parameters:

- for a variable v , $r(v)$ is the number of constraints v or $\neg v$ appears in
- for any variable v , $c_1^v, \dots, c_{r(v)}^v$ lists the constraints v or $\neg v$ appears in
- for any constraint c , v_1^c, v_2^c, v_3^c lists the variables which appear in c
- $N = \{(v, c) \in V \times C : \neg v \text{ appears in } c\}$

The vertex set of $G_{\mathcal{X}}$ is

$$V_{\mathcal{X}} := \bigcup_{v \in V} \{v\} \times S_{r(v)} \cup \bigcup_{c \in C} \{c\} \times O \cup \bigcup_{(v,c) \in N} \{(v, c)\} \times H.$$

Note that $V_{\mathcal{X}}$ can be built from \mathcal{X} and the components of O, H , and $S_{r(n)}$ using products, intersections, and unions. Since O, H , and $S_{r(n)}$ all have computable codes, $\mathcal{X} \mapsto V_{\mathcal{X}}$ is Δ_1^1 in the codes.

The edge set of $G_{\mathcal{X}}$, $E_{\mathcal{X}}$, includes the following edges:

1. Any $\{x\} \times e \in V_{\mathcal{X}}^2$ with e a non-coding edge
2. $\{(v, u), (c, w)\}$ where v appears in c and u, w are on corresponding coding edges, i.e. where $c = c_i^v$, $v = v_j^c$, $(v, c) \notin N$, and u is on $e_{i,k}$ in $S_{r(v)}$ and w is on $e_{j,k}$ in O for some $k \in \{0, 1\}$.

3. $\{(v, u), (v, c, w)\}$ where $(v, c) \in N$, $c = c_i^v$, and either u is on $e_{i,0}$ in $S_{r(v)}$ and w is on a in H or u is on $e_{i,1}$ and w is on b .
4. $\{(v, c, w), (c, u)\}$ where $v = v_j^c$, $(v, c) \in N$, and either u is on $e_{j,0}$ in O and w is on c in H or u is on $e_{j,1}$ and w is on d .

Again, $E_{\mathcal{X}}$ can be built from codes for \mathcal{X} , O , H , and $S_{r(n)}$ using operations from Lemma A.0.7, so $\mathcal{X} \mapsto G_{\mathcal{X}} = (V_{\mathcal{X}}, E_{\mathcal{X}})$ is Δ_1^1 in the codes.

As before, any Borel edge 3-coloring of \mathcal{G}_x induces a Borel satisfying assignment of \mathcal{X} by setting a variable to True if the coloring agrees on corresponding pairs of edges and False otherwise.

For the converse, suppose we have a Borel satisfying assignment of \mathcal{X} . Then there is a Borel partial coloring f of the corresponding edges in each copy of S_n which extends to a not-necessarily Borel edge 3-coloring. On each copy of S_n , H , and O , there are finitely many edge colorings consistent with f . Using the Luzin–Novikov theorem, we can select a such coloring on each of these components and get a Borel edge 3-coloring of $G_{\mathcal{X}}$. \square

REFERENCES

- [Ada92] Scot Adams. “Very weak Bernoulli for amenable groups.” *Israel Journal of Mathematics*, **78**:145–176, 1992.
- [AK00] Scot Adams and Alexander S. Kechris. “Linear algebraic groups and countable Borel equivalence relations.” *Journal of the American Mathematical Society*, **13**:909–944, 2000.
- [AW11] Miklós Abért and Benjamin Weiss. “Bernoulli actions are weakly contained in any free action.” *Ergodic Theory and Dynamical Systems*, **33**, 03 2011.
- [AZ18] Francis Adams and Jindřich Zapletal. “Cardinal invariants of closed graphs.” *Israel Journal of Mathematics*, **227**:861–888, 2018.
- [Bar14] Libor Barto. “The collapse of the bounded width hierarchy.” *Journal of Logic and Computation*, **26**, 11 2014.
- [BCG21a] Sebastian Brandt, Yi-Jun Chang, Jan Grebík, Christoph Grunau, Václav Rozhon, and Zoltán Vidnyánszky. “On Homomorphism Graphs.” *arXiv: 2111.03683*, 2021.
- [BCG21b] Sebastian Brandt, Yi-Jun Chang, Jan Grebík, Christoph Grunau, Václav Rozhoň, and Zoltán Vidnyánszky. “Local Problems on Trees from the Perspectives of Distributed Algorithms, Finitary Factors, and Descriptive Combinatorics.” *arXiv: 2106.02066*, 2021.
- [BDJ15] Jakub Bulín, Dejan Delić, Marcel Jackson, and Todd Niven. “A finer reduction of constraint problems to digraphs.” *Log. Methods Comput. Sci.*, **11**(4), 2015.
- [Ber19] Anton Bernshteyn. “Measurable versions of the Lovász Local Lemma and measurable graph colorings.” *Advances in Mathematics*, **353**:153–223, 2019.
- [Ber21] Anton Bernshteyn. “Probabilistic constructions in continuous combinatorics and a bridge to distributed algorithms.” *arXiv*, **abs/2102.08797**, 2021.
- [BGR20] Sebastian Brandt, Christoph Grunau, and Václav Rozhoň. “Generalizing the Sharp Threshold Phenomenon for the Distributed Complexity of the Lovász Local Lemma.” *Proceedings of the 39th Symposium on Principles of Distributed Computing*, 2020.
- [BHU06] Alexander Berenstein, C. Henson, and Alexander Usyatsov. “Model Theory for Metric Structures.” *Model Theory with Applications to Algebra and Analysis*, **2**, 01 2006.

- [Bir35] Garret Birkhoff. “The structure of abstract algebras.” *Proceedings of the Cambridge Philosophical Society*, **31**:433–454, 1935.
- [BJ01] Andrei Bulatov and Peter Jeavons. “Algebraic structures in combinatorial problems.” *Technical report, Technische Universitat Dresden*, 2001.
- [BJK05] Andrei Bulatov, Peter Jeavons, and Andrei Krokhin. “Classifying the Complexity of Constraints Using Finite Algebras.” *SIAM J. Comput.*, **34**:720–742, 04 2005.
- [BK09] Libor Barto and Marcin Kozik. “Constraint Satisfaction Problems of Bounded Width.” In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pp. 595–603, 2009.
- [BK12] Libor Barto and Marcin Kozik. “Absorbing Subalgebras, Cyclic Terms, and the Constraint Satisfaction Problem.” *Log. Methods Comput. Sci.*, **8**, 2012.
- [BKK69] V. G. Bodnarchuk, L. A. Kaluzhnin, V. N. Kotov, and Boris A. Romov. “Galois theory for Post algebras. II.” *Cybernetics*, **5**:531–539, 1969.
- [BKM09] Libor Barto, Marcin Kozik, Miklós Maróti, and Todd Niven. “CSP dichotomy for special triads.” *Proceedings of The American Mathematical Society*, **137**, 09 2009.
- [BKN09] Libor Barto, Marcin Kozik, and Todd Niven. “The CSP Dichotomy Holds for Digraphs with No Sources and No Sinks (A Positive Answer to a Conjecture of Bang-Jensen and Hell).” *SIAM J. Comput.*, **38**:1782–1802, 01 2009.
- [BKW17] Libor Barto, Andrei Krokhin, and Ross Willard. “Polymorphisms, and How to Use Them.” In Andrei Krokhin and Stanislav Zivny, editors, *The Constraint Satisfaction Problem: Complexity and Approximability*, volume 7 of *Dagstuhl Follow-Ups*, pp. 1–44. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2017.
- [BNW06] Laurent Bartholdi, Markus Neuhauser, and Wolfgang Woess. “Horocyclic products of trees.” *Journal of the European Mathematical Society*, **10**:771–816, 2006.
- [BOP15] Libor Barto, Jakub Opršal, and Michael Pinsker. “The wonderland of reflections.” *Israel Journal of Mathematics*, **223**, 10 2015.
- [Bow19] Lewis Bowen. “Finitary random interlacements and the Gaboriau–Lyons problem.” *Geometric and Functional Analysis*, **29**:659–689, 2019.
- [Bra19] Zarathustra Brady. “Examples, counterexamples, and structure in bounded width algebras.” *arXiv: 1909.05901*, 2019.
- [BsT21] Ferenc Bencs, Aranka Hrušková, and László Márton Tóth. “Factor-of-iid Schreier decorations of lattices in Euclidean spaces.” *arXiv: 2101.12577*, 2021.

- [Bul17] Andrei Bulatov. “A Dichotomy Theorem for Nonuniform CSPs.” *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 319–330, 2017.
- [CKT12] Clinton T. Conley, Alexander S. Kechris, and Robin D. Tucker-Drob. “Ultraproducts of measure preserving actions and graph combinatorics.” *Ergodic Theory and Dynamical Systems*, **33**:334 – 374, 2012.
- [CLP16] Endre Csóka, Gábor Lippner, and Oleg Pikhurko. “König’s line coloring and Vizing’s theorems for graphings.” *Forum of Mathematics, Sigma*, **4**, 2016.
- [CP17] Yi-Jun Chang and Seth Pettie. “A Time Hierarchy Theorem for the LOCAL Model.” *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 156–167, 2017.
- [CT20] Clinton T. Conley and Omer Tamuz. “Unfriendly colorings of graphs with finite average degree.” *Proceedings of the London Mathematical Society*, **122**, 2020.
- [DP99] Víctor Dalmau and Justin Pearson. “Closure Functions and Width 1 Problems.” In *In CP 1999*, pp. 159–173. Springer-Verlag, 1999.
- [FV98] Tomás Feder and Moshe Vardi. “The Computational Structure of Monotone Monadic SNP and Constraint Satisfaction: A Study Through Datalog and Group Theory.” *SIAM J. Comput.*, **28**:57–104, 01 1998.
- [Gei68] David Geiger. “Closed systems of functions and predicates.” *Pacific Journal of Mathematics*, **27**(1):95 – 100, 1968.
- [GP20] Jan Grebik and Oleg Pikhurko. “Measurable versions of Vizing’s theorem.” *Advances in Mathematics*, **374**:107378, 08 2020.
- [GR21] Jan Grebík and Václav Rozhoň. “Local Problems on Grids from the Perspective of Distributed Algorithms, Finitary Factors, and Descriptive Combinatorics.” *arXiv: 2103.08394*, 2021.
- [Gre22] Jan Grebík. “Approximate Schreier decorations and approximate König’s line coloring Theorem.” *Annales Henri Lebesgue*, **5**:303–315, 2022.
- [GS17] Mohsen Ghaffari and Hsin-Hao Su. “Distributed Degree Splitting, Edge Coloring, and Orientations.” In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 2505–2523, 01 2017.
- [Hjo08] Greg Hjorth. “Selection Theorems and Treeability.” *Proceedings of the American Mathematical Society*, **136**(10):3647–3653, 2008.

- [HKL90] Leo Harrington, Alexander S. Kechris, and Alain Louveau. “A Glimm-Effros dichotomy for Borel equivalence relations.” *Journal of the American Mathematical Society*, **3**:903–928, 1990.
- [HLS14] Hamed Hatami, László Miklós Lovász, and Balázs Szegedy. “Limits of locally–globally convergent graph sequences.” *Geometric and Functional Analysis*, **24**:269–296, 2014.
- [HM09] Greg Hjorth and Benjamin Miller. “Ends of graphed equivalence relations, II.” *Israel Journal of Mathematics*, **169**:393–415, 01 2009.
- [HN90] Pavol Hell and Jaroslav Nešetřil. “On the complexity of H-coloring.” *Journal of Combinatorial Theory, Series B*, **48**(1):92–110, 1990.
- [Jea98] Peter Jeavons. “On the Algebraic Structure of Combinatorial Problems.” *Theor. Comput. Sci.*, **200**:185–204, 1998.
- [KM04] Alexander S. Kechris and Benjamin Miller. *Topics in Orbit Equivalence*. Springer Berlin Heidelberg, 2004.
- [KMM14] Keith Kearnes, Petar Marković, and Ralph Mckenzie. “Optimal strong Mal’cev conditions for omitting type 1 in locally finite varieties.” *Algebra universalis*, **72**, 08 2014.
- [KST99] Alexander S. Kechris, Slawomir Solecki, and Stevo Todorcevic. “Borel Chromatic Numbers.” *Advances in Mathematics*, **141**(1):1–44, 1999.
- [KT08] Alexander S. Kechris and Todor Tsankov. “Amenable actions and almost invariant sets.” *Proceedings of the AMS*, **136**:687 – 697, 2008.
- [Kun22] Gabor Kun. “The measurable Hall theorem fails for treeings.” *arXiv: 2106.02013*, 2022.
- [Lad75] Richard E. Ladner. “On the Structure of Polynomial Time Reducibility.” *J. ACM*, **22**(1):155–171, 1975.
- [LLT06] Benoît Larose, Cynthia Loten, and Claude Tardif. “A Characterisation of First-Order Constraint Satisfaction Problems.” *21st Annual IEEE Symposium on Logic in Computer Science (LICS’06)*, pp. 201–210, 2006.
- [LN11] Russell Lyons and Fedor Nazarov. “Perfect matchings as IID factors on non-amenable groups.” *Eur. J. Comb.*, **32**:1115–1125, 2011.
- [Lyo16] Russell Lyons. “Factors of IID on Trees.” *Combinatorics, Probability and Computing*, **26**:285 – 300, 2016.

- [Mar13] Andrew S. Marks. “A determinacy approach to Borel combinatorics.” *Journal of the American Mathematical Society*, **29**:579–600, 2013.
- [Mar19] Andrew S. Marks. “Effective Descriptive Set Theory.” *Preprint: https://www.math.ucla.edu/marks/notes/edst_notes.pdf*, 2019.
- [MM08] Miklós Maróti and Ralph McKenzie. “Existence theorems for weakly symmetric operations.” *Algebra Universalis*, **59**:463–489, 12 2008.
- [MMV16] Dave Witte Morris, Joy Morris, and Gabriel Verret. “Isomorphisms of Cayley graphs on nilpotent groups.” *New York Journal of Mathematics*, **22**:453–467, 2016.
- [Mor16] Joy Morris. “The CI Problem for Infinite Groups.” *Electron. J. Comb.*, **23**:P4.37, 2016.
- [Mos09] Y.N. Moschovakis. *Descriptive Set Theory*. Mathematical surveys and monographs. American Mathematical Society, 2009.
- [MT10] Robin A. Moser and Gábor Tardos. “A Constructive Proof of the General Lovász Local Lemma.” *Journal of the ACM*, **57**(2), 2010.
- [OW87] Donald Ornstein and Benjamin Weiss. “Entropy and isomorphism theorems for actions of amenable groups.” *Journal d’Analyse Mathématique*, **48**:1–141, 1987.
- [Pik21] Oleg Pikhurko. “Borel combinatorics of locally finite graphs.” *arXiv: 2009.09113*, 2021.
- [Pos41] Emil Leon Post. *The Two-Valued Iterative Systems of Mathematical Logic*. London: Oxford University Press, 1941.
- [RV17] Mustazee Rahman and Bálint Virág. “Local algorithms for independent sets are half-optimal.” *Annals of probability.*, **45**(3):1543–1577, 2017.
- [Sch78] Thomas J. Schaefer. “The Complexity of Satisfiability Problems.” In *Proceedings of the Tenth Annual ACM Symposium on Theory of Computing*, STOC ’78, p. 216–226, New York, NY, USA, 1978. Association for Computing Machinery.
- [She84] Saharon Shelah. “Can you take Solovay’s inaccessible away?” *Israel Journal of Mathematics*, **48**:1–47, 1984.
- [Sig10] Mark Siggers. “A strong Mal’cev condition for locally finite varieties omitting the unary type.” *Algebra Universalis*, **64**:15–20, 10 2010.
- [Sol98] Slawomir Solecki. “Decomposing Borel sets and functions and the structure of Baire class 1 functions.” *Journal of the American Mathematical Society*, **11**, 07 1998.

- [ST16] Brandon Seward and Robin D. Tucker-Drob. “Borel structurability on the 2-shift of a countable group.” *Ann. Pure Appl. Log.*, **167**:1–21, 2016.
- [SU13] Edward R. Scheinerman and Daniel H. Ullman. *Fractional Graph Theory: a Rational Approach to the Theory of Graphs*. Dover Publications, Minola, N.Y., 2013.
- [Tay77] Walter Taylor. “Varieties Obeying Homotopy Laws.” *Canadian Journal of Mathematics*, **29**(3):498–527, 1977.
- [Tho21] Riley Thornton. “ Δ_1^1 Effectivization in Borel Combinatorics.” *arXiv:2105.04063*, 2021.
- [Tho22] Riley Thornton. “Orienting Borel graphs.” *Proceedings of the American Mathematics Society*, **150**:1779–1793, 2022.
- [Tot21] László Márton Tóth. “Invariant Schreier decorations of unimodular random networks.” *Annales Henri Lebesgue*, **4**:1705–1726, 2021.
- [TV21] Stevo Todorčević and Zoltán Vidnyánszky. “A complexity problem for Borel graphs.” *Inventiones mathematicae*, **226**, 10 2021.
- [Wei20] Felix Weilacher. “Marked groups with isomorphic Cayley graphs but different Borel combinatorics.” *Fundamenta Mathematicae*, **251**:69–86, 2020.
- [Zhu17] Dmitriy Zhuk. “A Proof of CSP Dichotomy Conjecture.” *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 331–342, 2017.