



TECHNISCHE
UNIVERSITÄT
DARMSTADT

Investigations into the Universal Algebra of Hypergraph Coverings and Applications

Vom Fachbereich Mathematik
der Technischen Universität Darmstadt

zur Erlangung des Grades eines

Doktors der Naturwissenschaften
(Dr. rer. nat.)

genehmigte Dissertation

Tag der Einreichung: 18. Oktober 2018

Tag der mündlichen Prüfung: 12. Februar 2019

Referent: Prof. Dr. Martin Otto

1. Korreferent: ao. Univ.-Prof. Dr. Karl Auinger

2. Korreferent: Prof. Dr. Michael Joswig

von

M.Sc. Julian Bitterlich
aus München

Darmstadt, D17
2019

Bitterlich, Julian: *Investigations into the Universal Algebra of
Hypergraph Coverings and Applications*
Darmstadt, Technische Universität Darmstadt,
Jahr der Veröffentlichung der Dissertation auf TUPrints: 2019
URN: urn:nbn:de:tuda-tuprints-86914
Tag der mündlichen Prüfung: 12.02.2019

Veröffentlicht unter CC0 1.0 Universell
<https://creativecommons.org/licenses/>

Danksagung

Der Großteil meines Danks gebührt meinem Betreuer Prof. Dr. Martin Otto. Diese Arbeit ist durch ihn maßgeblich (zum Guten) beeinflusst worden. Außerdem gewährte er mir die notwendige Zeit diese Arbeit fertig zu stellen.

Außerdem möchte ich ao.Univ.-Prof. Dr. Karl Auinger danken. Durch unsere leider erst kürzlich begonnene Korrespondenz bekam ich nützliche Tipps in der Darstellung des zweiten Kapitels.

Ich danke auch der DFG für die Unterstützung im Rahmen des Projekts

„OT 147/6-1 Konstruktionen und Modelltheorie
für Hypergraphen kontrollierter Azyklizität“.

Deutsche Zusammenfassung

Diese Arbeit behandelt die zwei Themengebiete *azyklische Überlagerungen* und *Erweiterungsprobleme*.

Der erste Teil der Arbeit befasst sich zunächst mit unverzweigten Überlagerungen von Graphen. Es wird die generelle Theorie der unverzweigten Überlagerungen besprochen und anschliessend verallgemeinert zu *verstrickten Überlagerungen*. Überlagerungen dieses Typs erhalten festgelegte Strukturen des überlagerten Graphen. Es wird gezeigt wie unverzweigte Überlagerungen von Hypergraphen auf verstrickte Überlagerungen zurückgeführt werden können. Unter Zuhilfenahme weiterer Resultate können wir so die Klasse der Hypergraphen identifizieren die azyklische¹ unverzweigte Überlagerungen besitzen.

Der zweite Teil der Arbeit behandelt Erweiterungsprobleme. Bei Erweiterungsproblemen geht es darum, endliche Strukturen endlich so zu erweitern, dass partielle Automorphismen der Ausgangsstruktur auf der Erweiterung vervollständigt werden können. Wir besprechen klassische Resultate und formulieren diese so um, dass sie sich für eine algebraische Charakterisierung eignen. Diese können benutzt werden um neue Resultate bezüglich Erweiterungsproblemen zu erhalten.

¹Im Sinne von α -Azyklizität. Siehe [1, 8, 10]

Contents

Introduction	9
0 Preliminaries	15
0.1 Structures	15
0.2 Algebraic structures	17
1 Acyclicity and covers	21
1.1 Introduction: acyclic covers of simple graphs	21
1.2 Multidigraphs	27
1.2.1 Fundamental groupoids and fundamental groups	28
1.2.2 Spanning trees	29
1.2.3 Unbranched covers	29
1.2.4 (Local) homomorphisms of unbranched covers	31
1.2.5 The Galois connection	32
1.3 Granular and pointed granular covers	35
1.3.1 Pointed granular covers	38
1.3.2 Granular covers	40
1.4 Acyclicity in hypergraphs	45
1.4.1 Definition of acyclicity for hypergraphs	46
1.4.2 Acyclicity as absence of induced cycles	48
1.4.3 Hypercycles	49
1.4.4 Acyclicity as decomposability	54
1.5 Simply connected hypergraphs	56
1.5.1 Simple connectivity vs. hypergraph acyclicity	58
1.6 Acyclic covers of hypergraphs	61
1.6.1 Unbranched covers and granular covers	62
1.6.2 Branched covers and coset acyclic groupoids	66
2 Extension problems	77
2.1 An overview of classical results	78
2.2 Free solutions	84
2.3 Structural properties of plain extension problems	88
2.4 Inverse monoids as ‘abstract extension problems’	98
2.4.1 Inverse monoids	99
2.4.2 Freeness properties	102
2.5 Finite solutions with free behaviour	104
2.5.1 Cayley graphs and Margolis–Meakin expansions	105

Contents


2.5.2	The method of Auinger and Steinberg	108
2.5.3	The Method of Otto	112
2.5.4	The Free Extension Conjecture for extension problems	115
2.6	Applications	116
2.6.1	EPPA for tournaments and groups of odd order	116
2.6.2	F-inverse Covers of inverse monoids	118
2.7	Symmetric solutions	121
3	Conclusion	127

Introduction

Thematically, this thesis focuses on two general types of problems: obtaining ‘finite highly acyclic covers of hypergraphs’ and solving ‘extension problems’.

We exemplify each problem by a short combinatorial puzzle.

Puzzle A. Consider the following graph G with partial automorphism $\cdots\rightarrow$, i.e., $\cdots\rightarrow$ is an isomorphism of induced subgraphs:




Is it possible to extend G by adding finitely many new vertices and edges (it is not allowed to add edges between the vertices of G) such that $\cdots\rightarrow$ can be completed to an automorphism?

The answer to this puzzle is ‘yes’ (we give a possible solution at the end of this introduction). Actually, no matter which G and $\cdots\rightarrow$ are given, there is always a suitable extension for this puzzle. Even if we consider multiple partial automorphisms simultaneously:

Theorem (Hrushovski [25]). *Every finite graph G has finite extensions over which all partial automorphisms of G can be extended to automorphisms.*

This theorem is one example of classical results about ‘extension problems’. We discuss extension problems in Chapter 2.

Puzzle B. Below we give a sketch of a collection $H = \{s_0, s_1, s_2, s_3\}$ of 3-sets (sets of size 3). Intuitively, the 3-sets s_1, s_2 and s_3 form a 3-cycle.



Is it possible to find a finite collection \widehat{H} of 3-sets labelled by $\{s_0, s_1, s_2, s_3\}$ s.t. \widehat{H} has no 3-cycle with labels s_1, s_2, s_3 and it satisfies the following requirements:

- each vertex is in exactly two 3-sets and
- for distinct $i, j \in \{0, 1, 2, 3\}$, a 3-set with label s_i intersects non-trivially with exactly one 3-set with label s_j ?

The answer to this puzzle is ‘yes’ as well and we give a possible solution at the end of the introduction.

The collection H in Puzzle B is an example of a hypergraph. Generally, a hypergraph consists of a set of vertices and a set of hyperedges, which are non-empty finite subsets of the vertices. This puzzle basically asks for a finite unbranched cover of H which ‘unravels’ the cycle $s_1 s_2 s_3$.

The general question whether all finite hypergraphs have finite unbranched covers without short cycles is the underlying topic for Chapter 1.

Below we give overviews for the two parts of this thesis. In each case we outline, for purely motivational reasons, the model-theoretic setting in which these problems originate. Otherwise, the present work is combinatorial in nature and no knowledge of logic is required to understand these results. The motivation is followed by a short discussion of the major results and principal ideas.

Part 1: Acyclic covers of hypergraphs and acyclic granular covers

Characterisation theorems play an important role in model theory. A characterisation theorem gives an exact correspondence between a semantic property of formulas and a corresponding syntactic restriction. In the field of modal logic a pioneering example is the van Benthem-Rosen Theorem [38,49]. This theorem relates first order formulas that behave invariantly under so-called bisimulation to formulas of modal logic.

Theorem (van Benthem-Rosen). *A first order formula is bisimulation invariant (over finite models) if, and only if, it is equivalent to a modal logic formula (over finite models).*

This result sparked numerous variants and generalisations [12]. One such generalisation is Otto’s Theorem [34].

Theorem (Otto). *A first order formula is invariant under guarded bisimulation over finite models if, and only if, it is equivalent to a guarded fragment formula over finite models.*

A crucial step in the proof of this result is the construction of finite, highly α -acyclic branched covers of hypergraphs. The central topic of Chapter 1 is the following question: for which hypergraphs can we obtain finite, highly α -acyclic *unbranched* covers? An answer to this question could possibly result in a characterisation result for guarded logic with counting.

At this point we may address the notion of α -acyclicity. Unlike over simple graphs, there are multiple distinct notions of acyclicity for hypergraphs. For the reason given above we focus on α -acyclicity.

We now turn back to our main question. We rephrase it in a more concrete fashion: which finite hypergraphs have finite unbranched covers without α -cycles of length ℓ or shorter for arbitrary but fixed $\ell \geq 3$.

For some hypergraphs we can quickly determine that they do not have finite, highly α -acyclic unbranched covers. If the hypergraph has an α -cycle that manifests itself around some apex vertex, then every unbranched cover ‘preserves’ this cycle.

If such a configuration is not present in a hypergraph, then we call it *apex acyclic*. We are able to prove the following result (Corollary 1.6.8):

An apex acyclic hypergraph has an acyclic unbranched cover. (*)

However, we can only conjecture its finite variant, i.e., that every finite, apex acyclic hypergraph has finite, highly α -acyclic unbranched covers. So, we could not achieve a satisfactory answer to our initial question.

Still, the proof of (*) is interesting in itself. It is a non-trivial result and to establish it we obtain various side results. The proof of (*) consists of two parts: a construction step and a verification step.

The basic idea for the construction is simple: for a given hypergraph H , consider unbranched covers of its Gaifman graph G . To be able to recover a cover of H from a cover of G , we only consider covers of G that do not ‘break’ the hyperedges. This idea leads to the notion of ‘granular covers’.

Granular covers are graph covers that preserve some specified closed walks. A closed walk is preserved if all its lifts are also closed. We can specify a translation of unbranched covers of hypergraphs into suitable granular covers. This translation works for all hypergraphs.

In Proposition 1.3.11 we show that granular covers have universal objects. The corresponding result for unbranched covers of hypergraphs is Corollary 1.6.7, which says that every hypergraph has a ‘simply connected cover’.

Simple connectivity is an acyclicity notion for hypergraphs, which is topologically motivated. In general, the theory of unbranched covers of hypergraphs is a discrete formulation of basic homotopy theory.

The verification step is given in Theorem 1.5.9, which says that simple connectivity and α -acyclicity are equivalent on apex acyclic hypergraphs.

In the last part of Chapter 1 we discuss Otto’s method for constructing finite, highly α -acyclic branched covers. Providing branched covers requires different techniques than the ones we develop for unbranched covers. In order to show the existence of finite, highly α -acyclic branched covers Otto [35] defines the notion of ‘coset acyclic’ groupoids and proposes a construction of finite, coset acyclic groupoids. However, due to an error in a central argument for the validity of this construction the status of the central result of [35], stating that finite, coset acyclic groupoids exist in general, is currently in doubt. Until further clarification we treat this as Otto’s Conjecture (Conjecture 1.6.22). This error was only noticed after I submitted this thesis, thus many important applications also inherit the status of a conjecture; in particular the ‘Free Extension Conjecture’ Conjecture 2.1.11.

Part 2: Extension problems

A *homogeneous structure* is a countable relational structure \mathcal{A} such that any isomorphism between finite substructures of \mathcal{A} extends to an automorphism of \mathcal{A} . In a precise sense, homogeneous structures are determined by their finite substructures. A homogeneous structure \mathcal{A} is called the *Fraïssé limit* of a class C of finite relational structures if C consists of all finite substructures of \mathcal{A} up to isomorphism. Fraïssé’s Theorem characterises the classes which have a Fraïssé limit [14] (also see the survey article [29]).

The class of finite graphs has a Fraïssé limit which is called the random graph (other names are the Rado graph or Erdős–Rényi graph). Truss [48] formulates a condition of being a ‘generic automorphism’ of the random graph and he shows that the orbits of a generic automorphism are finite. In order to do that he proves the following, which marks the first result about extension problems.

Theorem (Truss). *Given a finite graph and an isomorphism between finite induced subgraphs, this graph can be embedded into a finite graph over which this partial automorphism can be extended to an automorphism.*

We discuss now the content of Chapter 2. An extension problem is a relational structure \mathcal{A} with a specified collection of partial automorphisms p_1, \dots, p_n . A solution of an extension problem is an extension \mathcal{A}^* of that structure with automorphisms p_1^*, \dots, p_n^* extending the p_i .

Most of the classical results about extension problems are concerned with establishing EPPA results for classes of structures: a class of structures C has the

‘extension property for partial automorphisms’ (EPPA) if every extension problem that has a solution in C (possibly infinite), also has a finite solution in C . We have seen the theorem of Hrushovski as one example. Another example is from [20] which shows EPPA for triangle free graphs.

The main contribution of Chapter 2 is that we translate these EPPA statements into conditions regarding the behaviour of the automorphisms of the solutions. The conditions we introduce are oblivious to the relational content of an extension problem so that we can consider extension problems over sets.

In Section 2.3 we introduce various properties (we call the collection of these properties approximate freeness conditions) s.t. for every approximate freeness condition Q we can establish the following result: given a finite set A with partial bijections p_1, \dots, p_n , there is a finite extension A^* of A and extensions p_1^*, \dots, p_n^* of p_1, \dots, p_n s.t. the p_1^*, \dots, p_n^* satisfy property Q .

We show how to translate the classical results into the format given above. For example, for Hrushovski’s Theorem we get a corresponding approximate freeness condition ‘parallel 2-freeness’, for which we can show that Hrushovski’s Theorem is ‘equivalent’ to the existence of finite parallel 2-free solutions.

This translation is fruitful for two reasons. First, the theorems about approximate freeness, unlike EPPA, are unconditional. This makes it easier to see potential applications.

Second, these approximate freeness conditions lend themselves to algebraic abstractions. In Section 2.4 we follow this route and define ‘abstract extension problems’ (inverse monoids with generators) and ‘solutions of extension problems’ (groups with generators). This algebraic formalisation of the theory of extension problems has the benefits that it can easily be applied to other settings and that typical tools of combinatorial group theory, such as Cayley graphs and Margolis-Meakin expansions, are at our disposal



Using the established theory we can give a reformulation of Otto’s Conjecture about ‘coset acyclic’ hypergraphs to obtain a conjecture about extension problems:² The ‘Free Extension Conjecture’ (Conjecture 2.1.11).

We show also that Otto’s Conjecture implies the Henckel-Rhodes conjecture, a long-standing open problem in the theory of inverse monoids (Lemma 2.6.8).

²To the best of my knowledge this results would generalises all previous results, except results about coherent solutions (cf. [42]). We discuss this shortly at the end of the conclusion Chapter 3.

Contents

Here are possible solution of Puzzle A and Puzzle B.



0 Preliminaries

This chapter presents definitions and results that are core the two main chapters (at least in part).

0.1 Structures

In this thesis we consider structures (as in the sense of model theory) to be purely relational. The description of a structure requires three pieces of information: its domain, its signature, and the interpretation of the relation symbols of the signature. A *signature* σ is a finite collection of relation-symbols with associated arities. A σ -*structure* $\mathcal{A} = (A, (R^A)_{R \in \sigma})$ consists of a set A , its *domain*, and relations $R^A \subseteq A^n$ for every $R \in \sigma$ with arity n . The *width* of a signature is the maximal arity of its relational symbols.

Usually we do not mention σ and assume it to be implicitly fixed in the background, for example when we say that two structures are linked by a homomorphism, it is tacitly assumed that both structures are over the same signature and if we speak of a class of structures it is understood that all structures in this class are over the same signature.

We write $\mathbf{a} = (a_1, \dots, a_n)$ for elements in A^n and extend set notation to the case of tuples e.g. $a \in \mathbf{a}$, $\mathbf{a} \cap \mathbf{b} = \emptyset$ or $\mathbf{a} \subseteq A$. For such statements we implicitly substitute $\mathbf{a} = (a_1, \dots, a_n)$ by its set of elements $\{a_1, \dots, a_n\}$. In accordance to this convention we also write $(a_i)_{i \in I} \subseteq A$ to express that the components of the sequence $(a_i)_{i \in I}$ are elements of A . For $f: A \rightarrow B$ we write $f(\mathbf{a})$ for $(f(a_1), \dots, f(a_n))$.

The Gaifman graph


A *simple graph* is a structure $G = (V, E)$ with vertex set V and irreflexive, symmetric, binary relation E . The *Gaifman graph* of a structure $\mathcal{A} = (A, (R^A)_{R \in \sigma})$ is the simple graph $G(\mathcal{A}) = (A, E)$ with edge relation E that connects distinct elements $a, b \in A$ if $a, b \in \mathbf{a}$ for some $\mathbf{a} \in R^A$ and $R \in \sigma$. Gaifman graphs give us a distance measure over structures. Two elements a, b of \mathcal{A} are at *distance* ℓ if their distance is ℓ in $G(\mathcal{A})$. The ℓ -*neighbourhood* $N^\ell(\mathcal{A}, a)$ of $a \in A$ is the set of all elements that are at most of distance ℓ of a . We also write $N^\ell(a)$ instead of $N^\ell(\mathcal{A}, a)$ if the structure \mathcal{A} is clear from the context, furthermore we also write $N(a)$ for $N^1(a)$.

Substructures, weak substructures and localisations

A *substructure* of a σ -structure \mathcal{A} is a σ -structure \mathcal{B} s.t. $B \subseteq A$ and $R^{\mathcal{B}} = R^{\mathcal{A}} \cap B^n$ for each $R \in \sigma$ with arity n . Substructures are completely determined by their domains and we write $\mathcal{A}|_B$ for the substructure with domain $B \subseteq A$. A structure \mathcal{A} is an *extension* of a structure \mathcal{B} if \mathcal{B} is a substructure of \mathcal{A} . A *weak substructure* of a σ -structure \mathcal{A} is a σ -structure \mathcal{B} s.t. $B \subseteq A$ and $R^{\mathcal{B}} \subseteq R^{\mathcal{A}} \cap B^n$ for each $R \in \sigma$ with arity n .

Definition 0.1.1. The ℓ -localisation $\mathcal{A}_{a;\ell}$ of a σ -structure \mathcal{A} at $a \in A$ is the weak substructure with domain $N^\ell(a)$ and relations $R^{\mathcal{A}_{a;\ell}} = \{\mathbf{a} \in R^{\mathcal{A}} : \mathbf{a} \cap N^{\ell-1}(a) \neq \emptyset\}$, for $R \in \sigma$ of arity n .

We write \mathcal{A}_a for the 1-localisation of \mathcal{A} at a ; in particular $R^{\mathcal{A}_a} = \{\mathbf{a} \in R^{\mathcal{A}} : a \in \mathbf{a}\}$. Below we depict a simple graph and its 2-localisation at a .



There is a discrepancy in the terminology used in model theory and graph theory. What we call a weak substructure of a graph is called a subgraph in graph theory. To avoid confusion, we use the following convention: *induced subgraphs* are substructures of graphs and *weak subgraphs* are weak substructures of graphs.

Homomorphisms, strict homomorphisms and isomorphisms

Homomorphisms, strict homomorphisms and isomorphisms between σ -structures are defined as usual: let $\mathcal{A} = (A, (R^{\mathcal{A}})_{R \in \sigma})$ and $\mathcal{B} = (B, (R^{\mathcal{B}})_{R \in \sigma})$ be two structures, and $f: A \rightarrow B$ a map. Then f is

- a *homomorphism* ($f: \mathcal{A} \xrightarrow{\text{hom}} \mathcal{B}$ for short) if for all $R \in \sigma$ of arity n and $\mathbf{a} \in A^n$

$$\mathbf{a} \in R^{\mathcal{A}} \implies f(\mathbf{a}) \in R^{\mathcal{B}}.$$

- a *strong homomorphism* if for all $R \in \sigma$ of arity n and $\mathbf{a} \in A^n$ and

$$\mathbf{a} \in R^{\mathcal{A}} \iff f(\mathbf{a}) \in R^{\mathcal{B}}.$$

- an *embedding* ($f: \mathcal{A} \hookrightarrow \mathcal{B}$ for short) if f is an injective, strong homomorphism.
- an *isomorphism* ($f: \mathcal{A} \xrightarrow{\text{iso}} \mathcal{B}$ for short) if f is a homomorphism from \mathcal{A} to \mathcal{B} and f^{-1} is a homomorphism from \mathcal{B} to \mathcal{A} .

Note that there is an embedding from \mathcal{A} into \mathcal{B} , if and only if \mathcal{A} is isomorphic to a substructure of \mathcal{B} . From a model theoretic standpoint, in which one considers structures only up to isomorphism, there is no actual difference between substructures and embedded structures. So, we may freely replace substructure by embedded structure in any context. We also note that isomorphisms can be characterised as bijective strong homomorphism.

0.2 Algebraic structures

Whereas (relational) structures are our objects of study, we use certain algebraic structures as tools to facilitate constructions on these structures. The algebraic structures in question are monoids, groups, categories and groupoids.

Monoids and groups

A *monoid* M is an algebraic structure with an associative product that has a neutral element. The product of two elements $m, n \in M$ is written as mn and the neutral element is denoted by 1 . A *group* G is a monoid in which every element $g \in G$ has an inverse g^{-1} .

A homomorphism $f: M \xrightarrow{\text{hom}} N$ of two monoids is a map $f: M \rightarrow N$ s.t. $f(1) = 1$ and $f(mn) = f(m)f(n)$ for all $m, n \in M$. A monoid homomorphism f respects inverses i.e., if m has an inverse m^{-1} , then $f(m^{-1}) = f(m)^{-1}$. Thus a monoid homomorphism between groups is also a group homomorphism.

Sometimes we exhibit over the same domain a monoid structure as well as a group structure with different products. In this case we use the following convention to distinguish the products. We denote the product in the monoid by fg and the product in the group by $f \cdot g$. We make this convention precise every time we encounter such a situation.

Categories and groupoids

Categories and groupoids can be seen as typed versions of monoids and groups.

In a *category* \mathbb{C} , an element $e \in \mathbb{C}$ (also called a morphism) has an associated *source*, $s(e)$, and a *target*, $t(e)$. The set of all sources and targets of the morphisms of \mathbb{C} is denoted by $\text{Obj}(\mathbb{C})$ and we call its elements objects (they take the role of the types). On \mathbb{C} there is a partial, binary operation that is defined for $e, f \in \mathbb{C}$ if $t(e) = s(f)$. We write ef for the product of e and f . This partial operation is associative, i.e., $(ef)g = e(fg)$ whenever this is defined, and every object $a \in \text{Obj}(\mathbb{C})$ has a neutral element 1_a , i.e., $1_a e = e$ and $f 1_a = f$ whenever this is defined.

A *groupoid* \mathbb{G} is a category in which every element $g \in \mathbb{G}$ has an inverse $g^{-1} \in \mathbb{G}$, i.e., $s(g) = t(g^{-1})$, $t(g) = s(g^{-1})$ and $gg^{-1} = 1_{s(g)}$, $g^{-1}g = 1_{t(g)}$.

Groups with generators

We fix a finite set P equipped with an involution $(\cdot)^{-1}$.

Definition 0.2.1. A P -generated group is a group G with a fixed family of generators $(g_p)_{p \in P} \subseteq G$ s.t. $g_p^{-1} = g_{p^{-1}}$.

We require that a homomorphism $f: G \xrightarrow{\text{hom}} H$ between P -generated groups G and H with generator families $(g_p)_{p \in P}$ and $(h_p)_{p \in P}$ maps g_p to h_p . Thus there is at most one homomorphism between two P -generated groups. In particular, if $f: G \xrightarrow{\text{hom}} H$ and $f': H \xrightarrow{\text{hom}} G$ are homomorphisms of P -generated groups, then $f^{-1} = f'$ as $f' \circ f$ and id_G are both endomorphisms of G and thus equal (similarly for $f \circ f' = \text{id}_H$). So two P -generated groups are isomorphic, if and only if they are mutually homomorphic.

P^* is the set of all words over P and $P^{\leq \ell}$ the set of all words of length at most ℓ . We write uv for the composition of two words $u, v \in P^*$. With this composition we can see P^* as a monoid. We define a formal inverse on P^* by $u^{-1} := p_n^{-1} \dots p_1^{-1}$ for $u = p_1 \dots p_n \in P^*$. For each P -generated group G we get a monoid homomorphism $[\cdot]: P^* \xrightarrow{\text{hom}} M$ via $u = p_1 \dots p_n \mapsto [u]_G = g_{p_1} \dots g_{p_n}$ that is also compatible with formal inverses in the sense that $[u]_G^{-1} = [u^{-1}]_G$.

A word $u = p_1 \dots p_n \in P^*$ is *reduced* if $p_i \neq p_{i+1}^{-1}$ for all $1 \leq i < n$.

Definition 0.2.2. The *reduction operator* $\text{red}: P^* \rightarrow P^*$ maps a word $u \in P^*$ to the reduced word which is obtained by successively eliminating subwords of the form pp^{-1} until there is no such substring left.

Although $\text{red}(u)$ is the result of a non-deterministic process it is uniquely defined. So, red is a well-defined operator. We also state a technical lemma for later reference.

Lemma 0.2.3. *Let $u, v \in P^*$. Then $\text{red}(\text{red}(u)\text{red}(v)) = \text{red}(uv)$.*

The set of all reduced words forms a group under the reduced product.

Definition 0.2.4. The *free group* $\text{FG}(P)$ has the reduced words over P as its elements and the *reduced product* $u \cdot v = \text{red}(uv)$ as its group operation.


The free group $\text{FG}(P)$ is a P -generated group with the canonical choice of $(p)_{p \in P}$.

The free group $\text{FG}(P)$ is universal, possessing a homomorphism into every P -generated group G . This group homomorphism is induced by the monoid homomorphism $[\cdot]: P^* \rightarrow G$.


We end this introductory chapter with a lemma about the structure of $\text{FG}(P)$. Systematically, we put it here as it is used in Chapter 1 (Lemma 1.6.18) and Chapter 2 (Lemma 2.3.12). Conceptually, we put it here as it shows us that the free group satisfies a very strong form of acyclicity. Approximating this form of acyclicity by finite means accounts for a large part of this work.

Lemma 0.2.5. *Let $(v_i)_{i \in \mathbb{Z}_\ell} \subseteq \text{FG}(P)$ with $[v_0 \dots v_{\ell-1}]_{\text{FG}(P)} = 1$. Then there is a v_j that has a decomposition $v_j = v_- v_+$ s.t. v_-^{-1} is a suffix of v_{j-1} and v_+^{-1} is a prefix of v_{j+1} .*

Proof. We provide an informal proof sketch with the help of some visual intuition. We write the elements $v_0, \dots, v_{\ell-1}$ as a ‘circular word’ like this



$[v_0 \dots v_{\ell-1}]_{\text{FG}(P)} = 1$ means that this ‘cycle’ reduces to the empty word by performing individual reduction steps. We fix one such possible reduction process. For this reduction process, we consider the step in which for the first time a word is completely reduced. Let v_j be that word. Then all previous cancellations of letters of v_j could only be with letters of its neighbours v_{j-1} and v_{j+1} , as indicated in the following sketch



Let v_- be the prefix of v_j that got cancelled with v_{j-1} and v_+ the suffix of v_j that got cancelled with v_{j+1} . Then v_- and v_+ is a suitable decomposition of v_j . \square

1 Acyclicity and covers

1.1 Introduction: acyclic covers of simple graphs

This section introduces, in the setting of simple graphs, all important concepts that appear in this chapter. This has two purposes. First, this serves as a gentle introduction for basic definitions and special conventions. Second, the terminology and theorems developed here form the basis for further generalisations or, in some cases, can be used to highlight the fact that certain aspects cannot be transferred to a more general setting.

The most important concepts that are presented in this introduction are: acyclicity and different ways and means to characterize it, fundamental groupoids and fundamental groups, branched and unbranched covers, the Galois connection between subgroups of the fundamental group and unbranched covers, free covers and finite locally free covers. Branched and unbranched covers do behave analogously to their continuous counterparts in homotopy theory [13,17] and also this analogy has been noted before e.g. [44]. However, the notion of locally free covers is a genuinely new notion.

Since this is an introduction, we focus on the motivation of notions and on conceptual insights and not on formal completeness. In particular, this does not contain any formal proofs. These are then provided in Section 1.2 (in the more general framework of multidigraphs). For crucial statements we also provide the forward reference to the corresponding proof.

Simple graphs

We recall the definition of a simple graph: A *simple graph* is a structure $G = (V, E)$ with vertex set V and an anti-reflexive, symmetric, binary relation E . A *walk* of length $n \geq 0$ from $a \in V$ to $b \in V$ is a succession of vertices $a_0 \dots a_n$ s.t. $a_0 = a$, $a_n = b$ and $(a_i, a_{i+1}) \in E$ for $0 \leq i < n$. A *closed walk* at a is a walk that starts and ends at a . The *distance*, $\text{dist}(a, b)$, between two vertices $a, b \in V$ is the length of a shortest connecting walk or ∞ if none exists. Given two walks $\alpha = a_0 \dots a_n$ and $\beta = b_0 \dots b_m$ s.t. $a_n = b_0$ we define the *concatenation* $\alpha\beta$ as $a_0 \dots a_{n-1}b_0 \dots b_m$. The set of all walks equipped with the concatenation operation forms a category where V is the set of objects of this category. For the object $a \in V$ the neutral element 1_a is the trivial walk a of length 0. Similarly, the set of all closed walks at $a \in V$ forms a monoid which can be understood as the restriction of the aforementioned groupoid to the object $a \in V$.

Fundamental groupoids and fundamental groups

A walk $a_0 \dots a_n$ is *reduced* if $a_{i-1} \neq a_{i+1}$ for $i = 1, \dots, n-1$. A *trail* is a reduced walk and a *tour* is a closed, reduced walk. The set of all trails in a simple graph G equipped with the *reduced product*, $\alpha \cdot \beta = \text{red}(\alpha\beta)$, forms a groupoid. Here the *reduction* $\text{red}(\alpha)$ of a walk α is the reduced walk resulting of successively replacing substrings of the form aba by a (similar to the reduction operation introduced in Definition 0.2.2). We call this groupoid the *fundamental groupoid of G* and denote it by $\pi(G)$. The inverse of $a_0 \dots a_n \in \pi(G)$ is given by $a_n \dots a_0$. The *fundamental group of G at a* , $\pi(G, a)$, is the restriction of $\pi(G)$ to the object a or, put differently, the group of all tours at a .

The fundamental groups $\pi(G, a)$ and $\pi(G, b)$ are conjugate in $\pi(G)$ if a and b are in the same connected component. In fact, $\alpha \cdot \pi(G, b) \cdot \alpha^{-1} = \pi(G, a)$ for any trail α from a to b (see Lemma 1.2.4)

Acyclicity

Acyclicity for simple graphs is defined via the prohibition of ‘cyclic configurations’. One enjoys quite some freedom in how to define these cyclic configurations in detail.

We present a couple of cyclic configurations that can be used to define acyclicity: cycles, tours, and triangles and chordless cycles. A *cycle* is a closed walk that does not repeat any vertex except the end and starting vertex. A tour, as defined above, is a closed, reduced walk. A *triangle* is a closed walk of length 3. A *chord* of a cycle $a_0 \dots a_n$ is an edge $(a_i, a_j) \in E$ with i, j non-adjacent in \mathbb{Z}_n . A cycle is *chordless* if it does not have a chord. We use these configurations to give three equivalent definitions of ℓ -acyclicity (for some acyclicity parameter $\ell \in \mathbb{N}$). A graph $G = (V, E)$ is ℓ -acyclic if one of the following, equivalent conditions is fulfilled:

(A1 $^\ell$) G has no non-trivial cycles of length at most ℓ ,

(A2 $^\ell$) G has no non-trivial tours of length at most ℓ ,

(A3 $^\ell$) G has no triangles and no chordless cycles of length at least 4 and at most ℓ .

If G is ℓ -acyclic for all $\ell \in \mathbb{N}$ then it is *acyclic*. If we do not want to specify the concrete degree of ℓ -acyclicity we use the term *local acyclicity*. This term is motivated by the observation that simple graphs are ℓ -acyclic, if and only if their ℓ -localisations are acyclic.

In Section 1.4.3 we generalize cycles, tours, and triangles and chords to suitable notions over hypergraphs. Note that (A2 $^\ell$), by definition, is equivalent to $\{\alpha \in \pi(G, a) \mid \text{the length of } \alpha \text{ is at most } \ell\} = \{a\}$ for all $a \in V$. Writing the left side more succinctly we obtain that G is ℓ -acyclic, if and only if $\pi(G, a) \cap V^{\leq \ell+1} = \{a\}$ for all $a \in V$.

We can also characterise acyclicity by decomposability. A *tree* is a connected, acyclic simple graph. Each non-trivial, finite tree has a leaf, a vertex with exactly one incident edge. Thus, a tree can be ‘trimmed’ by removing leaves until only a




Figure 1.1: An unbranched cover and branched cover. The graph in the middle with vertices $\{a, b, c\}$ is the base graph. The graph on the left is the covering graph of an unbranched cover whose covering map is indicated by the vertex labels. Similarly, the graph on the right is the covering graph of a branched cover.

single vertex is left. Only trees can be cut down in this way. In Section 1.4.4 we see how to generalize this concept for hypergraphs.

Finally, we look at the topological component of acyclicity. A connected space is simply connected if its fundamental groups are trivial. We can understand a simple graph $G = (V, E)$ as a 1-dimensional abstract simplicial complex. The fundamental group of a geometric realisation of the abstract simplicial complex at $a \in V$ is isomorphic to $\pi(G, a)$. Thus acyclicity and simple connectivity agree over connected, simple graphs. This simple relationship no longer applies to hypergraphs, but in Section 1.5 we show how acyclicity and simple connectedness can be linked for hypergraphs as well.

Covers

A cover consists of three building blocks: the *base graph* G , the *covering graph* \widehat{G} and the *covering map* $\varphi: \widehat{G} \rightarrow G$. We treat covers only for connected base and covering graphs. This stipulation simplifies the presentation.

We introduce two types of covers, branched covers and unbranched covers. For an example examine Figure 1.1. For simple graphs, unbranched covers $\varphi: \widehat{G} \rightarrow G$ can be thought of as maps that are locally isomorphisms, i.e., its restrictions to 1-localisations $\varphi: \widehat{G}_{\widehat{a}} \rightarrow G_{\varphi(\widehat{a})}$ are isomorphisms. Alternatively, an unbranched cover can be described as a homomorphism $\varphi: \widehat{G} \xrightarrow{\text{hom}} G$ that has the *unique lifting property*:

(ul): if $\varphi(\widehat{a}) = a$ and $(a, b) \in E$ then there is a unique \widehat{b} s.t. $\varphi(\widehat{b}) = b$ and $(\widehat{a}, \widehat{b}) \in \widehat{E}$.

We call $(\widehat{a}, \widehat{b})$ the *lift* of (a, b) to \widehat{a} .

For a branched cover lifts do not have to be unique. A branched cover is a homomorphism $\varphi: \widehat{G} \xrightarrow{\text{hom}} G$ that has the *lifting property*:

1 Acyclicity and covers

- (1): if $\varphi(\widehat{a}) = a$ and $(a, b) \in E$ then there is at least one \widehat{b} s.t. $\varphi(\widehat{b}) = b$ and $(\widehat{a}, \widehat{b}) \in \widehat{E}$.

We write $\varphi: \widehat{G} \xrightarrow{\text{bra}} G$ for branched covers and $\varphi: \widehat{G} \xrightarrow{\text{unb}} G$ for unbranched covers.

Sometimes it benefits the intuition to think of the covering map as a labelling of the vertices of the covering graph. From this perspective it is only natural to extend the covering map φ to walks in the covering graph. Given a walk $\widehat{\alpha} = \widehat{a}_0 \dots \widehat{a}_n$ in the covering graph, $\varphi(\widehat{\alpha})$ is the corresponding walk in the base graph given by the vertex labels, i.e. $\varphi(\widehat{\alpha}) = \varphi(\widehat{a}_0) \dots \varphi(\widehat{a}_n)$.

Another useful way of thinking about unbranched covers is to view the covering graph as some sort of unfolding/unravelling of the base graph, appealing to the idea that an unbranched cover can be constructed by some kind of unfolding process. In Figure 1.1 the unbranched cover unravels the tour $abca$. In contrast the tour $abcabca$ is not unravelled.

A *pointed cover centred at a* $\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a$ is an unbranched cover $\varphi: \widehat{G} \xrightarrow{\text{unb}} G$ with designated vertices a and \widehat{a} s.t. $\varphi(\widehat{a}) = a$. There is a tight correspondence between subgroups of $\pi(G, a)$ and pointed covers of G centred at a . We discuss this in a section on the Galois connection (see p. 25 below).

As noted, the covering map φ of a cover induces a projection of the walks in the covering graph to the walks in the base graph, which we also denote by φ . This map commutes with the composition operation and the inversion operation. However, for branched covers, φ does not necessarily map reduced walks to reduced walks. For example, the branched cover in Figure 1.1 contains a reduced walk whose projection to G is bab . Unbranched covers, on the other hand, do map reduced walks to reduced walks. Thus an unbranched cover $\varphi: \widehat{G} \xrightarrow{\text{unb}} G$ induces a groupoid homomorphism from $\pi(\widehat{G})$ to $\pi(G)$ via the mapping $\widehat{\alpha} \mapsto \varphi(\widehat{\alpha})$ and this in turn induces a group homomorphism from $\pi(\widehat{G}, \widehat{a})$ to $\pi(G, \varphi(\widehat{a}))$ (see Lemma 1.2.10).

For our purposes of finding acyclic covers of graphs there is no need to use branched covers, as any degree of acyclicity can be achieved by finite unbranched covers already (Lemma 1.3.13). The proof we provide for this statement uses the general framework of granular covers that we develop below. We want to remark that there are simpler more direct proofs of Lemma 1.3.13 (see [34]).

So in the context of simple graphs and multidigraphs (which we introduce later) we focus on unbranched covers. However, in the context of hypergraphs, branched covers are more prominent as unbranched covers can be too rigid for some notions of acyclicity (cf. Section 1.6.1).

Ramifications of the unique lifting property

Let $\varphi: \widehat{G} \rightarrow G$ be an unbranched cover. A *lift* of a walk α in G is a walk $\widehat{\alpha}$ in \widehat{G} that projects onto α i.e. $\varphi(\widehat{\alpha}) = \alpha$. The unique lifting property of unbranched covers ensures unique lifts of walks up to the choice of a starting vertex. This means that given a and a preimage $\widehat{a} \in \varphi^{-1}(a)$, every walk starting at a has a unique lift to a walk starting at \widehat{a} . This implies that $\varphi: \pi(\widehat{G}, \widehat{a}) \xrightarrow{\text{hom}} \pi(G, a)$ is injective.

1.1 Introduction: acyclic covers of simple graphs

Adopting the view of \widehat{G} as a labelled graph, the set $\varphi(\pi(\widehat{G}, \widehat{a}))$ describes those sequences of labels which, if traced starting from \widehat{a} , return to \widehat{a} . Alternatively, seeing \widehat{G} as an unravelling, $\varphi(\pi(\widehat{G}, \widehat{a}))$ describes those tours in G that are not unravelled at \widehat{a} . We can actually use those ‘un-unravelled’ trails to compare unbranched covers.

A homomorphism $f: (\psi: \widetilde{G} \xrightarrow{\text{unb}} G) \rightarrow (\varphi: \widehat{G} \xrightarrow{\text{unb}} G)$ of two unbranched covers is a graph homomorphism $f: \widetilde{G} \xrightarrow{\text{hom}} \widehat{G}$ s.t. the following triangle commutes:

$$\begin{array}{ccc} \widetilde{G} & \xrightarrow{\psi} & G \\ \downarrow f & \nearrow \varphi & \\ \widehat{G} & & \end{array}$$

Homomorphisms between pointed covers are required to map the designated vertices accordingly. Since the two covering maps ψ and φ locally are isomorphisms, the homomorphism f is also locally an isomorphism and so $f: \widetilde{G} \xrightarrow{\text{unb}} \widehat{G}$ itself is also an unbranched cover.

The unique lifting property implies that a homomorphism f of unbranched covers is completely determined if for one vertex a value is given. Thus there is at most one homomorphism between pointed covers and we can use this fact to define a partial order on pointed covers. We write $(\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a) \leq (\psi: \widetilde{G}, \widetilde{a} \xrightarrow{\text{unb}} G, a)$ if there is a homomorphism from the latter cover to the former. This is indeed a partial order as $(\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a) \leq (\psi: \widetilde{G}, \widetilde{a} \xrightarrow{\text{unb}} G, a)$ and $(\psi: \widetilde{G}, \widetilde{a} \xrightarrow{\text{unb}} G, a) \leq (\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a)$ implies that $(\psi: \widetilde{G}, \widetilde{a} \xrightarrow{\text{unb}} G, a) \simeq (\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a)$. Indeed, if $f: (\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a) \rightarrow (\psi: \widetilde{G}, \widetilde{a} \xrightarrow{\text{unb}} G, a)$ and $g: (\psi: \widetilde{G}, \widetilde{a} \xrightarrow{\text{unb}} G, a) \rightarrow (\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a)$ are homomorphisms, then $g \circ f = \text{id}_{\widehat{G}}$ as $g \circ f$ and $\text{id}_{\widehat{G}}$ are both endomorphisms of $\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a$ and thus equal, similarly for $f \circ g = \text{id}_{\widetilde{G}}$.

The Galois connection

Similar to homotopy theory [17] there is a Galois connection between the poset of pointed covers at a and the poset of subgroups of $\pi(G, a)$. This connection manifests itself as follows:

- (i) for each subgroup $N \subseteq \pi(G, a)$ there is a unique pointed cover $\varphi: \widehat{G}, \widehat{a} \rightarrow G, a$ s.t. $\varphi(\pi(\widehat{G}, \widehat{a})) = N$,
- (ii) $(\varphi: \widehat{G}, \widehat{a} \rightarrow G, a) \leq (\psi: \widetilde{G}, \widetilde{a} \rightarrow G, a)$ if, and only if, $\psi(\pi(\widetilde{G}, \widetilde{a})) \subseteq \varphi(\pi(\widehat{G}, \widehat{a}))$.

This connection can also be extended by a finiteness condition: for finite G the cover $\varphi: \widehat{G}, \widehat{a} \rightarrow G, a$ is finite, if and only if $\varphi(\pi(\widehat{G}, \widehat{a}))$ has finite index in $\pi(G, a)$. Via the Galois connection we can reduce the theory of unbranched covers to the theory of subgroups of the free group.

Establishing the Galois connection is the goal of Section 1.2.5.

The free cover

Free covers can be defined in two equivalent ways: by acyclicity or by universality. We start with acyclicity. The free cover of G is the unique unbranched cover whose covering graph is acyclic. Let $\psi: \tilde{G} \xrightarrow{\text{unb}} G$ be the free cover of G . Then, for any vertex \tilde{a} of \tilde{G} , we have $\pi(\tilde{G}, \tilde{a}) = \{\tilde{a}\}$ and consequently $\psi(\pi(\tilde{G}, \tilde{a})) = \{a\}$. Thus, by the Galois connection, $\psi: \tilde{G} \xrightarrow{\text{unb}} G$ is universal, i.e., there is a homomorphism to any unbranched cover of G . So the free cover is uniquely defined (justifying the use of the definite article). Conversely, universality also implies acyclicity. More formally, the two equivalent definitions of the free cover $\psi: \tilde{G} \xrightarrow{\text{unb}} G$ are

- (i) $\psi(\pi(\tilde{G}, \tilde{a})) = \{\psi(\tilde{a})\}$ for all vertices \tilde{a} of \tilde{G} ,
- (ii) $\psi: \tilde{G} \xrightarrow{\text{unb}} G$ is universal.

Property (i) does not change if we replace the universal quantification by an existential quantification, i.e, replace “for all vertices \tilde{a} ” by “for some vertex \tilde{a} ” (recall that we require \tilde{G} and G to be connected).

Finite approximations to the free cover

The free cover can be thought of as the ‘most general’ cover. However, for the purpose of finite combinatorics it is desirable to have finite approximations to the free cover, i.e., we want a notion of local freeness that can be achieved in finite covers. We define these *locally free covers* formally, by localising the conditions (i) and (ii) of the free cover.

Localising (i) is straightforward. For a cover $\varphi: \hat{G} \xrightarrow{\text{unb}} G$ we require the image of all tours that are contained in some localisation $\hat{G}_{\hat{a};\ell}$ to be trivial. We can rewrite this condition in order to match the original format of (i) more closely. For a given $\hat{a} \in \hat{V}$ the set of tours at \hat{a} in $\hat{G}_{\hat{a};\ell}$ is given by $\pi(\hat{G}, \hat{a}) \cap \hat{V}^{\leq 2\ell+1}$ and its image under φ is $\varphi(\pi(\hat{G}, \hat{a})) \cap V^{\leq 2\ell+1}$. Using this, we define a cover $\varphi: \hat{G} \xrightarrow{\text{unb}} G$ to be ℓ -free if

$$\varphi(\pi(\hat{G}, \hat{a})) \cap V^{\leq 2\ell+1} = \{\varphi(\hat{a})\} \text{ for all } \hat{a} \in \hat{V}.$$

For the localised version of (ii) we need a notion of local homomorphisms. For $\ell \in \mathbb{N}$ an ℓ -local homomorphism f from $\varphi: \hat{G}, \hat{a} \xrightarrow{\text{unb}} G, a$ to $\psi: \tilde{G}, \tilde{a} \xrightarrow{\text{unb}} G, a$ is a homomorphism $f: \hat{G}_{\hat{a};\ell} \xrightarrow{\text{hom}} \tilde{G}_{\tilde{a};\ell}$ s.t.

$$f(\hat{a}) = \tilde{a} \quad \text{and} \quad \varphi|_{\hat{G}_{\hat{a};\ell}} = \psi|_{\tilde{G}_{\tilde{a};\ell}} \circ f.$$

The notion of *local universality* is now defined similarly to universality: $\varphi: \hat{G}, \hat{a} \xrightarrow{\text{unb}} G, a$ is ℓ -universal if it has ℓ -local homomorphisms to every pointed cover of G centred at a . This gives the second definition of ℓ -free covers: $\varphi: \hat{G} \xrightarrow{\text{unb}} G$ is ℓ -free if for each base vertex \hat{a} of \hat{G} the induced pointed cover is ℓ -locally universal.

We sum up. A cover $\varphi: \hat{G} \xrightarrow{\text{unb}} G$ is ℓ -free if it fulfils one of the following two equivalent conditions:




Figure 1.2: The graph on the left is the base graph of the two covers indicated on the right (the covering maps are given by the vertex-labels). The cover on the top is the free cover and the cover on the bottom a 2-free cover.

- (i) $\varphi(\pi(\widehat{G}, \widehat{a})) \cap V^{2\ell+1} = \{\varphi(\widehat{a})\}$ for every $\widehat{a} \in \widehat{G}$, or
- (ii) $\varphi: \widehat{G}, \widehat{a} \rightarrow G, a$ is ℓ -universal for all $a \in V$ and $\widehat{a} \in \varphi^{-1}(a)$.

As mentioned above, locally free covers can be achieved by finite covers (Lemma 1.3.13). For example in Figure 1.2 we see the free cover of a triangle and a finite 2-locally free cover.

We discuss some properties of local homomorphisms. By and large, they behave like ordinary homomorphisms with the important difference that in general local homomorphisms do not induce unbranched covers (consider a restriction of the homomorphism of the free cover of the triangle to the 2-locally free cover in Figure 1.2). Nevertheless, every ℓ -homomorphism induces a homomorphism $f: \pi(\widehat{G}_{\widehat{a}; \ell}) \rightarrow \pi(\widetilde{G}_{\widetilde{a}; \ell})$ of the fundamental groupoids of the localised graphs, and ℓ -local homomorphisms are uniquely defined.

We also obtain some form of Galois connection for ℓ -local homomorphisms between pointed covers $\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a$ and sets of the form $N \cap V^{\leq 2\ell+1}$ for subgroups $N \subseteq \pi(G, a)$: for each subgroup $N \subseteq \pi(G, a)$ there is a pointed cover $\varphi: \widehat{G}, \widehat{a} \rightarrow G, a$ s.t. $\varphi(\pi(\widehat{G}, \widehat{a})) \cap V^{\leq 2\ell+1} = N \cap V^{\leq 2\ell+1}$. Also, if we set \leq_{ℓ} to be the partial order of pointed covers defined in terms of existence of ℓ -local homomorphism, we get that

$$(\psi: \widetilde{G}, \widetilde{a} \rightarrow G, a) \leq_{\ell} (\varphi: \widehat{G}, \widehat{a} \rightarrow G, a) \iff \varphi(\pi(\widehat{G}, \widehat{a})) \cap V^{\leq 2\ell+1} \subseteq \psi(\pi(\widetilde{G}, \widetilde{a})) \cap V^{\leq 2\ell+1}.$$

1.2 Multidigraphs

The main goal of this section is to prove the Galois connection (Theorem 1.2.17) and its generalisation for local homomorphisms. For that we revisit the claims of the introductory section and provide formal proofs for them.

Definition 1.2.1. A *multidigraph* $G = (V, E)$ is a two-sorted structure with a set of vertices V , a set of edges E , source and target functions $s: E \rightarrow V$, $t: E \rightarrow V$, and

1 Acyclicity and covers

a fixepoint-free involution $\cdot^{-1}: E \rightarrow E$ satisfying $s(e^{-1}) = t(e)$ and $t(e^{-1}) = s(e)$.¹

Multidigraphs are indeed a generalisation of simple graphs. To cast a simple graph (V, E) as a multidigraph we just put $s((u, v)) = u$, $t((u, v)) = v$ and $(u, v)^{-1} = (v, u)$. All results we obtain in this section also hold for the special case of simple graphs.

1.2.1 Fundamental groupoids and fundamental groups

Due to the change of formalism, we have to repeat some definitions. In simple graphs we describe walks, trails, tours, etc. by the vertices they traverse whereas in multidigraphs we describe them by the edges they traverse.

A *walk* in a multidigraph $G = (V, E)$ is a succession of edges $e_1 \dots e_n$ s.t. $t(e_i) = s(e_{i+1})$ for $1 \leq i < n$; its length is n and its source, $s(\alpha)$, is $s(e_1)$ and its target, $t(\alpha)$, is $t(e_n)$. The trivial walk at $a \in V$ is denoted by ε_a and has length 0.

The composition $\alpha_1 \alpha_2 = e_1 \dots e_n e'_1 \dots e'_m$ of two walks $\alpha_1 = e_1 \dots e_n$ and $\alpha_2 = e'_1 \dots e'_m$ is defined if $t(\alpha_1) = s(\alpha_2)$ and the trivial walks are the neutral elements of this operation. A walk $\alpha = e_1 \dots e_n$ is reduced (also called a trail) if $e_i \neq e_{i+1}^{-1}$ for all $1 \leq i < n$. The reduction $\text{red}(\alpha)$ of a walk α is defined as the reduction operation in Definition 0.2.2, but with the additional detail that in case all letters of $e_1 \dots e_n$ get eliminated by the reduction process we set $\text{red}(e_1 \dots e_n) = \varepsilon_a$ where a is the source of the last pair ee^{-1} that gets eliminated. One can show that $s(\alpha) = s(\text{red}(\alpha))$ and $t(\alpha) = t(\text{red}(\alpha))$. In particular, the operators s and t are well-defined.

A tour is a closed trail. A multidigraph is ℓ -acyclic if it has no non-trivial tours of length ℓ or shorter.

Definition 1.2.2. The fundamental groupoid $\pi(G)$ of $G = (V, E)$ is the groupoid on the set of trails in G with the reduced product $\alpha \cdot \beta = \text{red}(\alpha\beta)$.

The fundamental group $\pi(G, a)$ is the group of tours at a with the reduced product as group operation.

To be precise, the set of objects of $\pi(G)$ is V , ε_a is the neutral element at $a \in V$, and the inverse of $e_1 \dots e_n$ is $e_n^{-1} \dots e_1^{-1}$.

Lemma 1.2.3. $\pi(G)$ is a groupoid and $\pi(G, a)$ is a group.

Proof. Associativity is guaranteed by Lemma 0.2.3 as $(\alpha \cdot \beta) \cdot \gamma = \text{red}(\alpha\beta\gamma) = \alpha \cdot (\beta \cdot \gamma)$. The trivial walks ε_a act as neutral element by definition and the inverse of $\alpha = e_1 \dots e_n$ does its job, as $\alpha \cdot \alpha^{-1} = \text{red}(e_1 \dots e_n e_n^{-1} \dots e_1^{-1}) = \varepsilon_{s(e_1)} = \varepsilon_{s(\alpha)}$. \square

Lemma 1.2.4. Let $\alpha \in \pi(G)$ with $s(\alpha) = a$ and $t(\alpha) = b$. Then

$$\alpha \cdot \pi(G, b) \cdot \alpha^{-1} = \pi(G, a).$$

Proof. It suffices to show that $\pi(G, a) \subseteq \alpha \cdot \pi(G, b) \cdot \alpha^{-1}$ as the converse follows from $\pi(G, b) \subseteq \alpha^{-1} \cdot \pi(G, a) \cdot \alpha$. Let $r \in \pi(G, a)$. Then $\alpha^{-1} \cdot r \cdot \alpha \in \pi(G, b)$ and thus $r \in \alpha \cdot \pi(G, b) \cdot \alpha^{-1}$. \square

¹We equip our definition of multidigraphs with an involutive operation on the edges. We do so in order to be able to formulate walks that can use the reverse of an edge. This is no harm, as any ordinary multidigraph can be enriched to match our format by adding inverses of edges.

1.2.2 Spanning trees

In the following we use spanning trees to show that $\pi(G, a)$ is a free group.

A connected, acyclic multidigraph is a tree. A spanning tree of a connected multidigraph G is a maximal weak subgraph of G that is a tree.

Lemma 1.2.5. *Every connected multidigraph G has a spanning tree.*

Proof. The proof is an easy application of Zorn's Lemma. \square

A spanning tree T of a connected multidigraph G contains every vertex, otherwise it would not be maximal. Furthermore, adding an edge of G to T that is not part of T creates a cycle.

We show now that $\pi(G, a)$ is a free group. This means that for a suitable set P and choice of generators $\alpha_p \in \pi(G, a)$, $\pi(G, a)$ is isomorphic to $\text{FG}(P)$.

Lemma 1.2.6. *$\pi(G, a)$ is a free group.*

Proof. We can assume that G is connected as $\pi(G, a)$ only depends on the connected component of a . Let $T = (V, E')$ be a maximal spanning tree of $G = (V, E)$. Let $P = E \setminus E'$. We associate to each $p \in P$ the tour $\alpha_p = \beta p \beta'$ where β is the unique trail in T from a to $s(p)$ and β' equivalently from $t(p)$ to a .

Now we show that the family $(\alpha_p)_{p \in P}$ generates $\pi(G, a)$. Given $\alpha \in \pi(G, a)$ we show by induction over the number of occurrences of elements of P in α that $\alpha \in \langle \alpha_p \rangle_{p \in P}$. If no element of P occurs in α , then α is a tour in T and thus trivial. Otherwise we decompose α in $\gamma p \gamma'$ with γ being a trail in T and $p \in P$. Since trails in T are uniquely defined by their source and target we get that $\gamma = \beta$ where $\alpha_p = \beta p \beta'$. Hence

$$\alpha_p^{-1} \cdot \alpha = (\beta'^{-1} p^{-1} \beta^{-1}) \cdot (\gamma p \gamma') = \beta'^{-1} \cdot \gamma'.$$

The tour $\beta'^{-1} \cdot \gamma'$ has fewer occurrences of elements in P and thus is in $\langle \alpha_p \rangle_{p \in P}$. Consequently $\alpha = \alpha_p \cdot (\alpha_p^{-1} \cdot \alpha) \in \langle \alpha_p \rangle_{p \in P}$.

$\pi(G, a)$ as a P -generated group is isomorphic to $\text{FG}(P)$: if $\text{red}(p_1 \dots p_n) \neq \emptyset$ then we have also $\text{red}(\alpha_{p_1} \dots \alpha_{p_n}) \neq \varepsilon_a$ as the elements that do not cancel in $\text{red}(p_1 \dots p_n)$ also appear in the same order in $\text{red}(\alpha_{p_1} \dots \alpha_{p_n}) \neq \varepsilon_a$ and thus do not cancel there as well. \square

A corollary to the proof is the following.

Corollary 1.2.7. *If G is finite then $\pi(G, a)$ is finitely generated.*

1.2.3 Unbranched covers

We define unbranched covers as homomorphisms with the unique lifting property, like in the setting of simple graphs.

1 Acyclicity and covers

Definition 1.2.8. A homomorphism $f: \widehat{G} \xrightarrow{\text{hom}} G$ from a multidigraph $\widehat{G} = (\widehat{V}, \widehat{E})$ to a multidigraph $G = (V, E)$ is a two-sorted map $f = (f_V: \widehat{V} \rightarrow V, f_E: \widehat{E} \rightarrow E)$ that commutes with the source and target functions as well as with inversion, i.e.,

$$s(f_E(\widehat{e})) = f_V(s(\widehat{e})), \quad t(f_E(\widehat{e})) = f_V(t(\widehat{e})), \quad f_E(\widehat{e}^{-1}) = f_E(\widehat{e})^{-1}.$$

Usually we simply write f to denote both f_E and f_V .

A homomorphism $f: \widehat{G} \xrightarrow{\text{hom}} G$ can be extended to walks in \widehat{G} by $f(\widehat{e}_1 \dots \widehat{e}_n) := f(\widehat{e}_1) \dots f(\widehat{e}_n)$ and $f(\varepsilon_{\widehat{a}}) := \varepsilon_{f(\widehat{a})}$. The extended f provides a functor between the walks in \widehat{G} and G , in particular: $s(f(\widehat{\alpha})) = f(s(\widehat{\alpha}))$, $t(f(\widehat{\alpha})) = f(t(\widehat{\alpha}))$ and $f(\widehat{\alpha}^{-1}) = f(\widehat{\alpha})^{-1}$. Furthermore $f(\widehat{\alpha})$ is closed/non-reduced if $\widehat{\alpha}$ is closed/non-reduced.

Definition 1.2.9. An *unbranched cover* is a homomorphism $\varphi: \widehat{G} \xrightarrow{\text{hom}} G$ of connected multidigraphs that has the unique lifting property. The *unique lifting property* consists of *unique source-lifting* (usl) and *unique target-lifting* (utl):

(usl) if $\varphi(\widehat{a}) = a$ and $s(e) = a$ then there is a unique \widehat{e} s.t. $s(\widehat{e}) = \widehat{a}$ and $\varphi(\widehat{e}) = e$.

(utl) if $\varphi(\widehat{a}) = a$ and $t(e) = a$ then there is a unique \widehat{e} s.t. $t(\widehat{e}) = \widehat{a}$ and $\varphi(\widehat{e}) = e$.

We write $\varphi: \widehat{G} \xrightarrow{\text{unb}} G$ to denote an unbranched cover of G by \widehat{G} . A *pointed unbranched cover* $\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a$ is an unbranched cover $\varphi: \widehat{G} \xrightarrow{\text{unb}} G$ with designated vertices \widehat{a} and a s.t. $\varphi(\widehat{a}) = a$.

For a homomorphism $\varphi: \widehat{G} \xrightarrow{\text{hom}} G$ the properties (usl) and (utl) are equivalent. So in order to show that φ is an unbranched cover it suffices to establish one of the two properties.

We can understand the uniqueness part of the unique lifting property also as local injectivity of φ . If $\varphi(\widehat{e}_1) = \varphi(\widehat{e}_2)$ and $s(\widehat{e}_1) = s(\widehat{e}_2)$ (or $t(\widehat{e}_1) = t(\widehat{e}_2)$) then $\widehat{e}_1 = \widehat{e}_2$.

Lemma 1.2.10. *The covering map φ of an unbranched cover $\varphi: \widehat{G} \xrightarrow{\text{iso}} G$ induces a homomorphism $\varphi: \pi(\widehat{G}) \rightarrow \pi(G)$ via $\varphi(\widehat{e}_1 \dots \widehat{e}_n) = \varphi(\widehat{e}_1) \dots \varphi(\widehat{e}_n)$.*

Proof. First we show that φ is well-defined on trails, i.e., $\varphi(\pi(\widehat{G})) \subseteq \pi(G)$. Assume that $\varphi(\widehat{\alpha})$ is not reduced for some walk $\widehat{\alpha} = \widehat{e}_1 \dots \widehat{e}_n$ in \widehat{G} . Then $\varphi(\widehat{e}_i) = \varphi(\widehat{e}_{i+1})^{-1}$ for some $i \in \{1, \dots, n\}$. Since $t(\widehat{e}_i) = s(\widehat{e}_{i+1}) = t(\widehat{e}_{i+1}^{-1})$ we get by (utl) that $\widehat{e}_i = \widehat{e}_{i+1}^{-1}$. Hence $\widehat{\alpha}$ is not reduced.

It remains to show that φ is compatible with the reduced product. For this we use that

$$\varphi(\text{red}_{\widehat{G}}(\widehat{\alpha})) = \text{red}_G(\varphi(\widehat{\alpha})), \quad (*)$$

where $\text{red}_{\widehat{G}}$ and red_G are the reduction operations in the respective multidigraphs. This can be shown easily by induction on the reduction process. Now

$$\varphi(\widehat{\alpha} \cdot \widehat{\beta}) = \varphi(\text{red}_{\widehat{G}}(\widehat{\alpha}\widehat{\beta})) \stackrel{(*)}{=} \text{red}_G(\varphi(\widehat{\alpha}\widehat{\beta})) = \text{red}_G(\varphi(\widehat{\alpha})\varphi(\widehat{\beta})) = \varphi(\widehat{\alpha}) \cdot \varphi(\widehat{\beta}). \quad \square$$

A lift of a walk α to \hat{a} is a walk $\hat{\alpha}$ s.t. $s(\hat{\alpha}) = \hat{a}$ and $\varphi(\hat{\alpha}) = \alpha$. The next lemma shows us that the unique lifting property generalises to walks.

Lemma 1.2.11. *Let α be a walk starting at a and $\varphi: \hat{G}, \hat{a} \xrightarrow{\text{unb}} G, a$. Then there is a unique lift of α to \hat{a} .*

Proof. Let $G = (V, E)$ and $\hat{G} = (\hat{V}, \hat{E})$. If $\alpha = \varepsilon_a$ then only $\varepsilon_{\hat{a}}$ has \hat{a} as source and projects to ε_a . Uniqueness for longer walks follows by induction: if α is a walk of length $n + 1$, then it decomposes into $\alpha = \beta e$, where β is a walk of length n and e an edge. Any lift $\hat{\alpha}$ of α also decomposes similarly into $\hat{\alpha} = \hat{\beta} \hat{e}$. Then $\hat{\beta}$ is a lift of β and uniquely determined by induction hypothesis and so then is \hat{e} by (usl). \square

Corollary 1.2.12. *A pointed cover $\varphi: \hat{G}, \hat{a} \xrightarrow{\text{unb}} G, a$ induces an injective group homomorphism $\varphi: \pi(\hat{G}, \hat{a}) \xrightarrow{\text{hom}} \pi(G, a)$.*

1.2.4 (Local) homomorphisms of unbranched covers

We introduce homomorphisms and local homomorphisms between unbranched covers and show that they induce homomorphisms between fundamental groupoids.

Definition 1.2.13. A homomorphism from the unbranched cover $\psi: \tilde{G} \xrightarrow{\text{unb}} G$ to the unbranched cover $\varphi: \hat{G} \xrightarrow{\text{unb}} G$ is a homomorphism of multidigraphs $f: \tilde{G} \xrightarrow{\text{hom}} \hat{G}$ s.t. $\psi = \varphi \circ f$.

A homomorphism between pointed covers $\psi: \tilde{G}, \tilde{a} \xrightarrow{\text{unb}} G, a$ and $\varphi: \hat{G}, \hat{a} \xrightarrow{\text{unb}} G, a$ is a homomorphism between $\psi: \tilde{G} \xrightarrow{\text{unb}} G$ and $\varphi: \hat{G} \xrightarrow{\text{unb}} G$ that maps \tilde{a} to \hat{a} .

Lemma 1.2.14. *Let $f: \tilde{G} \xrightarrow{\text{hom}} \hat{G}$ be a homomorphism between $\psi: \tilde{G} \xrightarrow{\text{unb}} G$ and $\varphi: \hat{G} \xrightarrow{\text{unb}} G$. Then this homomorphism is also a cover $f: \tilde{G} \xrightarrow{\text{unb}} \hat{G}$.*

Proof. We check that f satisfies (usl). Let $f(\tilde{a}) = \hat{a}$ and $s(\tilde{e}) = \hat{a}$. We have to show that there is a unique \tilde{e} satisfying $(*)$: $s(\tilde{e}) = \tilde{a}$ and $f(\tilde{e}) = \hat{e}$.

Uniqueness: if \tilde{e}_1 and \tilde{e}_2 satisfy $(*)$ then $s(\tilde{e}_1) = s(\tilde{e}_2)$ and $\psi(\tilde{e}_1) = \varphi(f(\tilde{e}_1)) = \varphi(f(\tilde{e}_2)) = \psi(\tilde{e}_2)$ and so $\tilde{e}_1 = \tilde{e}_2$ by (usl) for ψ .

Existence: let \tilde{e} be the source-lift of $\varphi(\hat{e})$ to \tilde{a} . Then $s(f(\tilde{e})) = f(s(\tilde{e})) = \hat{a} = s(\hat{e})$ and $\varphi(f(\tilde{e})) = \psi(\tilde{e}) = \varphi(\hat{e})$. Thus $f(\tilde{e}) = \hat{e}$ by (usl) for φ . Hence \tilde{e} satisfies $(*)$. \square

The ℓ -localisation $G_{a;\ell}$ of a multidigraph $G = (V, E)$ is the weak subgraph

$$G_{a;\ell} = (N^\ell(a), E') \text{ with } E' = \{e \in E : s(e) \in N^{\ell-1}(a) \text{ or } t(e) \in N^{\ell-1}(a)\}.$$

Definition 1.2.15. An ℓ -local homomorphism from $\varphi: \hat{G}, \hat{a} \xrightarrow{\text{unb}} G, a$ to $\psi: \tilde{G}, \tilde{a} \xrightarrow{\text{unb}} G, a$ for $\ell \in \mathbb{N}$ is a homomorphism $f: \hat{G}_{\hat{a};\ell} \xrightarrow{\text{hom}} \tilde{G}_{\tilde{a};\ell}$ s.t.

$$f(\hat{a}) = \tilde{a} \quad \text{and} \quad \varphi|_{\hat{G}_{\hat{a};\ell}} = \psi|_{\tilde{G}_{\tilde{a};\ell}} \circ f.$$

1 Acyclicity and covers

Local homomorphisms are uniquely defined and they induce homomorphisms between fundamental groupoids and fundamental groups.

Lemma 1.2.16. *There is at most one ℓ -local homomorphism f from $\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a$ to $\psi: \widetilde{G}, \widetilde{a} \xrightarrow{\text{unb}} G, a$. Such a homomorphism f induces a homomorphism $f: \pi(\widehat{G}_{\widehat{a}; \ell}) \xrightarrow{\text{hom}} \pi(\widetilde{G}_{\widetilde{a}; \ell})$ and an injective homomorphism $f: \pi(\widehat{G}_{\widehat{a}; \ell}, \widehat{a}) \xrightarrow{\text{hom}} \pi(\widetilde{G}_{\widetilde{a}; \ell}, \widetilde{a})$.*

Proof. Observe that $\varphi(\widehat{\alpha}) = \psi(f(\widehat{\alpha}))$ for walks $\widehat{\alpha}$ in $\widehat{G}_{\widehat{a}; \ell}$.

Injectivity of f as a map between the fundamental groups $\pi(\widehat{G}_{\widehat{a}; \ell}, \widehat{a})$ and $\pi(\widetilde{G}_{\widetilde{a}; \ell}, \widetilde{a})$ basically follows from injectivity of φ on $\pi(\widehat{G}, \widehat{a})$.

To show that f induces a homomorphism between $\pi(\widehat{G}_{\widehat{a}; \ell}, \widehat{a})$ and $\pi(\widetilde{G}_{\widetilde{a}; \ell}, \widetilde{a})$ we have to prove that f maps reduced walks to reduced walks and also that f respects the groupoid operations. The first point is given by the following chain of implications

$$\widehat{\alpha} \text{ reduced} \implies \varphi(\widehat{\alpha}) \text{ reduced} \implies \psi(f(\widehat{\alpha})) \text{ reduced} \implies f(\widehat{\alpha}) \text{ reduced.}$$

For the second point we observe that

$$\psi(f(\widehat{\alpha} \cdot \widehat{\beta})) = \varphi(\widehat{\alpha} \cdot \widehat{\beta}) = \varphi(\widehat{\alpha}) \cdot \varphi(\widehat{\beta}) = \psi(f(\widehat{\alpha})) \cdot \psi(f(\widehat{\beta})) = \psi(f(\widehat{\alpha}) \cdot f(\widehat{\beta}))$$

Whence $f(\widehat{\alpha} \cdot \widehat{\beta}) = f(\widehat{\alpha}) \cdot f(\widehat{\beta})$ as $s(f(\widehat{\alpha} \cdot \widehat{\beta})) = s(f(\widehat{\alpha}) \cdot f(\widehat{\beta}))$.

We are left to show the uniqueness of f . Let g be another ℓ -local homomorphism. For $\widehat{b} \in \mathbb{N}^\ell(\widehat{a})$ let $\widehat{\alpha}$ be some trail from \widehat{a} to \widehat{b} in $\widehat{G}_{\widehat{a}; \ell}$. Then

$$s(f(\widehat{\alpha})) = f(s(\widehat{\alpha})) = \widetilde{a} = g(s(\widehat{\alpha})) = s(g(\widehat{\alpha})) \quad \text{and} \quad \psi(f(\widehat{\alpha})) = \varphi(\widehat{\alpha}) = \psi(g(\widehat{\alpha}))$$

and hence $f(\widehat{\alpha}) = g(\widehat{\alpha})$. So f and g agree on all trails in $\widehat{G}_{\widehat{a}; \ell}$. Hence they also agree on the vertices and edges of $\widehat{G}_{\widehat{a}; \ell}$. \square

Note that by the uniqueness of local homomorphisms any ℓ -local homomorphism between $\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a$ and $\psi: \widetilde{G}, \widetilde{a} \xrightarrow{\text{unb}} G, a$ is an extension of the ℓ' -local homomorphism between $\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a$ and $\psi: \widetilde{G}, \widetilde{a} \xrightarrow{\text{unb}} G, a$ for $\ell' \leq \ell$. In particular, if we have a sequence of i -local homomorphisms $f_i: (\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a) \rightarrow (\psi: \widetilde{G}, \widetilde{a} \xrightarrow{\text{unb}} G, a)$ for $i \in \mathbb{N}$ then $f = \bigcup_{i \in \mathbb{N}} f_i$ is a homomorphism.

1.2.5 The Galois connection

The Galois connection describes a fundamental connection between unbranched covers and subgroups of fundamental groups. This connection allows us to reduce the questions about unbranched covers to problems in group theory.

We write $(\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a) \leq (\psi: \widetilde{G}, \widetilde{a} \xrightarrow{\text{unb}} G, a)$ if there is a homomorphism from the pointed cover $(\psi: \widetilde{G}, \widetilde{a} \xrightarrow{\text{unb}} G, a)$ to the pointed cover $(\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a)$.

As we discussed in the introductory section, this defines a preorder on the class of pointed covers and in this preorder two covers are mutually comparable if, and

only if, they are isomorphic (see the argument for this in ‘Ramifications of the unique lifting property’ in Section 1.1). Similarly, we write $(\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a) \leq_\ell (\psi: \widetilde{G}, \widetilde{a} \xrightarrow{\text{unb}} G, a)$ if there is an ℓ -local homomorphism from the latter to the former. Note that \leq is the limit of the \leq_ℓ .

Theorem 1.2.17. *Let $G = (V, E)$ be a multidigraph and $a \in V$. Then:*

- (i) *For each subgroup $N \subseteq \pi(G, a)$ there is, up to isomorphism, a unique pointed cover $\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a$ s.t. $\varphi(\pi(\widehat{G}, \widehat{a})) = N$.*
- (ii) *$(\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a) \leq (\psi: \widetilde{G}, \widetilde{a} \xrightarrow{\text{unb}} G, a)$ if, and only if, $\psi(\pi(\widetilde{G}, \widetilde{a})) \subseteq \varphi(\pi(\widehat{G}, \widehat{a}))$.*
- (iii) *For each subgroup $N \subseteq \pi(G, a)$ there is, up to ℓ -local isomorphism, a unique pointed cover $\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a$ s.t. $\varphi(\pi(\widehat{G}, \widehat{a})) \cap E^{\leq 2\ell} = N \cap E^{\leq 2\ell}$.²*
- (iv) *$(\psi: \widetilde{G}, \widetilde{a} \xrightarrow{\text{unb}} G, a) \leq_\ell (\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a)$ if, and only if, $\varphi(\pi(\widehat{G}, \widehat{a})) \cap E^{\leq 2\ell} \subseteq \psi(\pi(\widetilde{G}, \widetilde{a})) \cap E^{\leq 2\ell}$.*

Proof. It suffices to prove (iv) and the existential claim of (i): that there is a cover $\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a$ with $\varphi(\pi(\widehat{G}, \widehat{a})) = N$. Then (ii) follows from (iv) and the fact that \leq is the limit of the \leq_ℓ . The uniqueness claim of (i) then follows by (ii), and the uniqueness claim of (iii) by (iv). Finally the existence claim for (iii) follows by (i). We prove the existence claim for (i) in Lemma 1.2.19 and we prove (iv) in Lemma 1.2.20 \square

Before we finish the proof of the Galois connection by giving the corresponding lemmas we give an account on how the Galois connection translates finiteness conditions. We say that a cover $\varphi: \widehat{G} \rightarrow G$ is finite if \widehat{G} is finite.

Lemma 1.2.18. *Let G be a finite, connected multidigraph. Then a pointed cover $\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a$ is finite if, and only if, $\varphi(\pi(\widehat{G}, \widehat{a}))$ has finite index in $\pi(G, a)$.*

Proof. We let $\widehat{G} = (\widehat{V}, \widehat{E})$, $N = \varphi(\pi(\widehat{G}, \widehat{a}))$ and define

$$N \backslash \pi(G) := \{ N \cdot \alpha \mid \alpha \in \pi(G), s(\alpha) = a \}$$

as the right coset of N in $\pi(G)$. Then $N \cdot \alpha \mapsto t(\widehat{\alpha})$, where $\widehat{\alpha}$ is the lift of α to \widehat{a} , provides a bijection between \widehat{V} and $N \backslash \pi(G)$.

It remains to show that $N \backslash \pi(G)$ is finite if $N \backslash \pi(G, a)$ is finite. Define

$$N \backslash \pi(G)[b] := \{ N\alpha \mid \alpha \in \pi(G), s(\alpha) = a, t(\alpha) = b \}.$$

Then $|N \backslash \pi(G)[b]| = |N \backslash \pi(G, a)|$ and thus $N \backslash \pi(G) = \bigcup_{b \in V} N \backslash \pi(G)[b]$ is finite. \square

²Note that we can see the fundamental group $\pi(G, a)$ of a multidigraph $G = (V, E)$ as a subset of E^* by identifying ε_a with ε . So we can denote the set of all tours at a of length at most ℓ by $\pi(G, a) \cap E^{\leq \ell}$.

1 Acyclicity and covers

We now give the crucial lemmas for the proof of the Galois connection.

Lemma 1.2.19. *Let N be a subgroup of the fundamental group $\pi(G, a)$. Then there exists an unbranched cover $\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a$ s.t. $\varphi(\pi(\widehat{G}, \widehat{a})) = N$.*

Proof. Given a multidigraph $G = (V, E)$ with vertex $a \in V$ and a subgroup $N \subseteq \pi(G, a)$. Let $\widehat{G} = (\widehat{V}, \widehat{E})$ be the multidigraph with the vertex set and edge set

$$\begin{aligned}\widehat{V} &:= \{ (t(\alpha), N \cdot \alpha) \mid \alpha \in \pi(G), s(\alpha) = a \} \\ \widehat{E} &:= \{ (e, N \cdot \alpha \mid \alpha \in \pi(G), e \in E, s(\alpha) = a, t(\alpha) = s(e)) \}\end{aligned}$$

whose source, target and involution maps are given as follows

$$\begin{aligned}s((e, N \cdot \alpha)) &:= (s(e), N \cdot \alpha), & t((e, N \cdot \alpha)) &:= (t(e), N \cdot (\alpha \cdot e)), \\ (e, N \cdot \alpha)^{-1} &:= (e^{-1}, N \cdot (\alpha \cdot e)).\end{aligned}$$

Let $\widehat{a} = (a, N) \in \widehat{V}$ and let φ be the projection to the first component. It is easy to check that $\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a$ is a pointed cover.

Let $\widehat{\alpha}$ be a trail at \widehat{a} and $\alpha = \varphi(\widehat{\alpha})$ its projection. Then $t(\widehat{\alpha}) = (t(\alpha), N \cdot \alpha)$. So

$$\widehat{\alpha} \text{ is a tour} \iff (a, N) = (t(\alpha), N \cdot \alpha) \iff \alpha \in N. \quad \square$$

For the next lemma we note that the projection of walks by a homomorphism $\varphi: \widehat{G} \xrightarrow{\text{unb}} G$ of multidigraphs is length preserving. In particular, $\varphi(\pi(\widehat{G}, \widehat{a})) \cap E^{\leq \ell} = \varphi(\{ \widehat{\alpha} \in \pi(\widehat{G}, \widehat{a}) \mid |\widehat{\alpha}| \leq \ell \})$.

Lemma 1.2.20. *Let $\psi: \widetilde{G}, \widetilde{a} \xrightarrow{\text{unb}} G, a$ and $\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a$ be two pointed covers and $\ell \in \mathbb{N}$. Then*

$$(\psi: \widetilde{G}, \widetilde{a} \xrightarrow{\text{unb}} G, a) \leq_{\ell} (\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a)$$

if, and only if,

$$\varphi(\pi(\widehat{G}, \widehat{a})) \cap E^{\leq 2\ell} \subseteq \psi(\pi(\widetilde{G}, \widetilde{a})) \cap E^{\leq 2\ell}.$$

Proof. We start with “ \implies ”. Let f be an ℓ -local homomorphism from $\psi: \widetilde{G}, \widetilde{a} \xrightarrow{\text{unb}} G, a$ to $\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a$. If $\varphi(\widehat{\alpha})$ is a tour in $\varphi(\pi(\widehat{G}, \widehat{a}))$ of length at most 2ℓ , then $\widehat{\alpha} \in \pi(\widehat{G}_{\widehat{a}; \ell}, \widehat{a})$. Thus $f(\widehat{\alpha}) \in \pi(\widetilde{G}_{\widetilde{a}; \ell}, \widetilde{a}) \subseteq \pi(\widetilde{G}, \widetilde{a})$, and so $\varphi(\widehat{\alpha}) = \psi(f(\widehat{\alpha})) \in \psi(\pi(\widetilde{G}, \widetilde{a})) \cap E^{\leq 2\ell}$.

Now we show “ \impliedby ”. We have to provide an ℓ -local homomorphism f witnessing $(\psi: \widetilde{G}, \widetilde{a} \xrightarrow{\text{unb}} G, a) \leq_{\ell} (\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a)$. We give f as a map that operates on trails. For this we choose f such that the following diagram commutes:


$$\begin{array}{ccc} \pi(\widetilde{G}, \widetilde{a}) \cap \widetilde{E}^{\leq 2\ell} & & \\ \uparrow f & \searrow \psi & \\ & & \pi(G, a) \\ & \nearrow \varphi & \\ \pi(\widehat{G}, \widehat{a}) \cap \widehat{E}^{\leq 2\ell} & & \end{array}$$

Such an f exists since $\varphi(\pi(\widehat{G}, \widehat{a})) \cap E^{\leq 2\ell} \subseteq \psi(\pi(\widetilde{G}, \widetilde{a})) \cap E^{\leq 2\ell}$ and is uniquely defined (as φ is injective). This $f: \pi(\widehat{G}, \widehat{a}) \cap \widehat{E}^{\leq 2\ell} \rightarrow \pi(\widetilde{G}, \widetilde{a}) \cap \widetilde{E}^{\leq 2\ell}$ can then be used to define the required homomorphism $f: \widehat{G}_{\widehat{a}, \ell} \xrightarrow{\text{hom}} \widetilde{G}_{\widetilde{a}, \ell}$. \square


1.3 Granular and pointed granular covers

In this section we introduce and discuss granular covers. Granular covers are unbranched covers that do not unravel certain designated tours in the base graph. We provide two settings here: ‘pointed granular covers’ and ‘granular covers’. For both settings we discuss freeness conditions and local freeness conditions. Before we get to the nitty-gritty details, let us go through these notions with an example.

Let G be the graph depicted below consisting of one vertex a and two edges r and s that loop at a .



Before considering granular covers we look at $\varphi^{(1)}: G_1 \rightarrow G$ the free cover of G , which is depicted below. The image of the fundamental group $\varphi^{(1)}(\pi(G_1, a_1))$ at any vertex of G_1 is the trivial subgroup $\{\varepsilon_a\}$ of $\pi(G, a)$. Notice that the unique lifting property holds: every vertex in G_1 has exactly one incoming and outgoing edge with label r respectively s . We can also describe the free cover as a special granular cover which does not require to preserve any tour. In the terminology that has yet to be introduced, the free cover can be called the free \emptyset -granular cover.



Now we consider *pointed $\{rs\}$ -granular covers*. These are pointed covers $\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a$ for which the lift of rs to \widehat{a} is a tour. We say that a pointed cover $\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a$ *preserves* a tour α if the lift of α to \widehat{a} is a tour, otherwise we say that the cover *unravels* α . Using this terminology, we can say that a pointed $\{rs\}$ -granular cover is a pointed cover that preserves rs .


Below we depict the free pointed $\{rs\}$ -granular cover $\varphi^{(2)}: G_2, a_2 \rightarrow G, a$. This is the pointed cover that unravels as many tours as possible while preserving rs . We

1 Acyclicity and covers

see that

$$\varphi^{(2)}(\pi(G_2, a_2)) = \{\dots, (rs)^{-1}, \varepsilon_a, (rs)^1, \dots\} = \text{cl}(\{rs\}),$$

where $\text{cl}(\{rs\})$ is the closure of $\{rs\}$ in $\pi(G, a)$. As we will see, every pointed $\{rs\}$ -granular cover preserves the tours in $\text{cl}(\{rs\})$ and thus the free pointed $\{rs\}$ -granular cover is universal for the $\{rs\}$ -granular covers.³




The free pointed $\{rs\}$ -granular cover is infinite but there are finite approximations to it, locally free pointed granular covers. The cover $\varphi^{(3)}: G_3, a_3 \xrightarrow{\text{unb}} G, a$ depicted below is one such. It is a finite, 2-free $\{rs\}$ -granular cover, i.e., $\varphi^{(2)}: G_2, a_2 \xrightarrow{\text{unb}} G, a \leq_2 \varphi^{(3)}: G_3, a_3 \xrightarrow{\text{unb}} G, a$ (Recall that \leq_2 refers to the existence of a 2-local homomorphism. Be aware that the loops at vertices of distance 2 of a_3 are not part of the 2-localisation $(G_3)_{a_3;2}$). Using the Galois connection we can express this fact by

$$\varphi^{(3)}(\pi(G_3, a_3)) \cap \{r, s\}^{\leq 4} = \varphi^{(2)}(\pi(G_2, a_2)) \cap \{r, s\}^{\leq 4} = \text{cl}(\{rs\}) \cap \{r, s\}^{\leq 4}.$$

In other words, $\varphi^{(3)}: G_3, a_3 \xrightarrow{\text{unb}} G, a$ preserves exactly those tours of length at most 4 that are preserved by the free pointed $\{rs\}$ -granular cover. Note that $\varphi^{(3)}: G_3, a_3 \xrightarrow{\text{unb}} G, a$ is not 3-free as e.g. $ssssss$ is in $\varphi^{(3)}(\pi(G_3, a_3))$ but not in $\varphi^{(2)}(\pi(G_2, a_2))$.

³Further examples of free pointed granular covers of G for other choices of tours that have to be preserved can be found in [17, p. 58].




Now we look at $\{rs\}$ -granular covers of G . An $\{rs\}$ -granular cover preserves the tour rs globally, i.e., every lift of rs is a tour. We say an unbranched cover *preserves* a tour if all lifts of that tour are tours and we say that the cover *unravels* a tour if every lift of that tour is not a tour. We introduce $\pi_{\text{cl}}(G)$, the set of all tours in a multidigraph G . We can express succinctly that an unbranched cover $\varphi: \widehat{G} \xrightarrow{\text{unb}} G$ preserves a tour α : every lift of α is a tour, if and only if $\varphi^{-1}(\alpha) \subseteq \pi_{\text{cl}}(\widehat{G})$. Similarly, $\varphi: \widehat{G} \xrightarrow{\text{unb}} G$ unravels α , if and only if $\varphi^{-1}(\alpha) \cap \pi_{\text{cl}}(\widehat{G}) = \emptyset$.

Below we depict the free $\{rs\}$ -granular cover $\varphi^{(4)}: G_4 \xrightarrow{\text{unb}} G$, that is the unbranched cover that unravels as many tours as possible while preserving rs . We can see that $\varphi^{(4)}: G_4 \xrightarrow{\text{unb}} G$ preserves all the tours in $\text{ncl}(\{rs\})$, the normal closure of $\{rs\}$ in $\pi(G, a)$, and unravels all others. In other ‘words’

$$(\varphi^{(4)})^{-1}(\text{ncl}(\{rs\})) = \pi_{\text{cl}}(G_4).$$

We will see that the inclusion $\varphi^{-1}(\text{ncl}(\{rs\})) \subseteq \pi_{\text{cl}}(\widehat{G})$ is true for all $\{rs\}$ -granular covers $\varphi: \widehat{G} \xrightarrow{\text{unb}} G$ and that this implies that the free $\{rs\}$ -granular cover is universal among $\{rs\}$ -granular covers.




Similarly to pointed granular covers we are interested in finite approximations of the free cover. Below we depict $\varphi^{(5)}: G_5 \rightarrow G$ a finite, 2-free $\{rs\}$ -granular cover. 2-freeness means that $\varphi^{(4)}: G_4, a_4 \xrightarrow{\text{unb}} G, a \leq_2 \varphi^{(5)}: G_5, a_5 \xrightarrow{\text{unb}} G, a$ for any choice

1 Acyclicity and covers

of a_4 and a_5 . So it preserves and unravels the same tours as the free $\{rs\}$ -granular cover for tours up to length 4. We can express the last condition more succinctly as

$$(\varphi^{(5)})^{-1}(\text{ncl}(\{rs\})) \cap \{E_5\}^{\leq 4} = \pi_{\text{cl}}(G_5) \cap \{E_5\}^{\leq 4},$$

where E_5 is the set of edges of G_5 .



If we compare G_3 and G_5 , the covering graphs of the 2-free pointed $\{rs\}$ -granular cover and the 2-free $\{rs\}$ -granular cover we see that G_3 is much more irregular than G_5 . The reason is that a locally free pointed granular cover is only required to show free behaviour in a neighbourhood around one distinguished vertex of the covering graph whereas in a locally free granular cover the neighbourhood of every vertex has to behave free. So, locally free granular covers are more restrictive than locally free pointed granular covers and thus are also more regular. Indeed, we will see that finite, locally free pointed granular covers do exist in general whereas finite, locally free granular covers do not always exist.

1.3.1 Pointed granular covers

A pointed cover $\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{umb}} G, a$ preserves a tour $\alpha \in \pi(G, a)$ if the lift of α to \widehat{a} is also a tour; otherwise we say that the cover *unravels* α . For a set $R \subseteq \pi(G, a)$ we say that $\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{umb}} G, a$ preserves R if all elements $\alpha \in R$ are preserved.

Definition 1.3.1. A *pointed R -granular cover* of a multidigraph G with designated vertex a and set of tours $R \subseteq \pi(G, a)$ is a pointed cover $\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{umb}} G, a$ s.t.

$$R \subseteq \varphi(\pi(\widehat{G}, \widehat{a})).$$

It is easy to see that a pointed cover $\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{umb}} G, a$ that preserves α and β also preserves α^{-1} and $\alpha\beta$. From this observation we get the following lemma ($\text{cl}(R)$ stands for the closure of R in $\pi(G, a)$).

Lemma 1.3.2. A *pointed R -granular cover* is also a *pointed $\text{cl}(R)$ -granular cover*.

In the following we show that the ‘free pointed R -granular cover’ exactly preserves $\text{cl}(R)$ and unravels $\pi(G, a) \setminus \text{cl}(R)$.

Free and locally free pointed granular covers

A pointed R -granular cover $\psi: \tilde{G}, \tilde{a} \xrightarrow{\text{unb}} G, a$ is universal if $(\varphi: \hat{G}, \hat{a} \xrightarrow{\text{unb}} G, a) \leq (\psi: \tilde{G}, \tilde{a} \xrightarrow{\text{unb}} G, a)$ for all pointed R -granular covers $\varphi: \hat{G}, \hat{a} \xrightarrow{\text{unb}} G, a$.

Proposition 1.3.3. *Let $\psi: \tilde{G}, \tilde{a} \xrightarrow{\text{unb}} G, a$ be a pointed cover and $R \subseteq \pi(G, a)$. Then the following are equivalent*

- (i) $\psi(\pi(\tilde{G}, \tilde{a})) = \text{cl}(R)$.
- (ii) $\psi: \tilde{G}, \tilde{a} \xrightarrow{\text{unb}} G, a$ is a universal pointed R -granular cover.

Furthermore such a universal cover exists and is unique up to isomorphism.

Proof. By Theorem 1.2.17 (i) there exists a unique pointed cover $\psi: \tilde{G}, \tilde{a} \xrightarrow{\text{unb}} G, a$ with $\psi(\pi(\tilde{G}, \tilde{a})) = \text{cl}(R)$. We are left to show that (i) and (ii) are equivalent. By Theorem 1.2.17 (ii) a cover satisfying (i) is also universal for R -granular covers. This shows us (i) \implies (ii).

For (ii) \implies (i) we observe that for a universal R -granular cover $\varphi: \tilde{G}, \tilde{a} \xrightarrow{\text{unb}} G, a$ there is a homomorphism to a cover $\varphi: \hat{G}, \hat{a} \xrightarrow{\text{unb}} G, a$ with property (i) and thus by Theorem 1.2.17 (ii)

$$\text{cl}(R) \subseteq \psi(\pi(\tilde{G}, \tilde{a})) \subseteq \varphi(\pi(\hat{G}, \hat{a})) = \text{cl}(R). \quad \square$$

We call the pointed R -granular cover described in Proposition 1.3.3 the *free pointed R -granular cover*.

We can derive a similar result for locally free pointed granular covers. A pointed R -granular cover $\varphi: \hat{G}, \hat{a} \xrightarrow{\text{unb}} G, a$ is ℓ -universal if $(\psi: \tilde{G}, \tilde{a} \xrightarrow{\text{unb}} G, a) \leq_{\ell} (\varphi: \hat{G}, \hat{a} \xrightarrow{\text{unb}} G, a)$ for every pointed R -granular cover $\psi: \tilde{G}, \tilde{a} \rightarrow G, a$. The following proposition can be proved similar to Proposition 1.3.3 using the results of Theorem 1.2.17.

Proposition 1.3.4. *Let $\varphi: \hat{G}, \hat{a} \xrightarrow{\text{unb}} G, a$ be a pointed R -granular cover and $R \subseteq \pi(G, a)$. Then the following are equivalent*

- (i) $\varphi(\pi(\hat{G}, \hat{a})) \cap E^{\leq 2\ell} = \text{cl}(R) \cap E^{\leq 2\ell}$.
- (ii) $\varphi: \hat{G}, \hat{a} \xrightarrow{\text{unb}} G, a$ is ℓ -universal.

We say that a pointed R -granular cover is ℓ -free if it fulfils one of the conditions (i) or (ii) of Proposition 1.3.4.

Finite, locally free pointed granular covers and the Theorem of M. Hall

Proposition 1.3.3 and Proposition 1.3.4 are basically exercises in applying the Galois connection. Their purpose is mainly to show that our definitions of freeness and local freeness for pointed granular covers are sensible notions.

1 Acyclicity and covers

The interesting, novel part is that we can show the existence of finite, locally free pointed granular covers provided that the base graph $G = (V, E)$ and $R \subseteq \pi(G, a)$ are finite.

Let N be a subgroup of $\pi(G, a)$ and $\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a$ its associated pointed cover. Then we have the following correspondences:

$$\begin{aligned} \varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a \text{ is a pointed } R\text{-granular cover} &\iff R \subseteq N, \\ \varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a \text{ is } \ell\text{-locally free} &\iff N \cap E^{\leq 2\ell} = \text{cl}(R) \cap E^{\leq 2\ell}, \\ \varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a \text{ is finite} &\stackrel{1,2,18}{\iff} N \text{ has finite index in } \pi(G, a). \end{aligned}$$

So we can reduce the question of the existence of finite, ℓ -free pointed covers to the group theoretic question whether there exists a subgroup $N \subseteq \pi(G, a)$ with the properties: $R \subseteq N$, $N \cap E^{\leq 2\ell} = \text{cl}(R) \cap E^{\leq 2\ell}$, and $|\pi(G, a) \setminus N| < \infty$.

To show the existence of such an N the following separation result of M. Hall suits perfectly. We give the original statement here [16, Theorem 5.1].

Theorem 1.3.5 (M. Hall). *Given a free group F with an arbitrary number of generators, and a finite number of elements $\alpha_1, \alpha_2, \dots, \alpha_m$ of F . Suppose we are also given a finite number of elements $\beta_1, \beta_2, \dots, \beta_n$ such that no β belongs to the subgroup H generated by $\alpha_1, \alpha_2, \dots, \alpha_m$. Then we may construct a subgroup \overline{H} of finite index in F containing $\alpha_1, \alpha_2, \dots, \alpha_m$ (and hence H) but no one of $\beta_1, \beta_2, \dots, \beta_n$*

We can use M. Hall's Theorem to show the existence of finite, locally free pointed granular covers.

Proposition 1.3.6. *Every finite multidigraph G with given finite set of tours $R \subseteq \pi(G, a)$ has finite, ℓ -free pointed R -granular covers for any $\ell \in \mathbb{N}$.*

Proof. By the discussion above we need to find a subgroups $N \subseteq \pi(G, a)$ with properties: $R \subseteq N$, $N \cap E^{\leq 2\ell} = \text{cl}(R) \cap E^{\leq 2\ell}$, and $|\pi(G, a) \setminus N| \leq \infty$.

According to Corollary 1.2.7, $\pi(G, a)$ is a free group. Let $\alpha_1, \dots, \alpha_m$ be an enumeration of R and β_1, \dots, β_n an enumeration of $(\pi(G, a) \setminus \text{cl}(R)) \cap E^{\leq 2\ell}$. By M. Hall's Theorem we obtain a subgroup $N \subseteq \pi(G, a)$ of finite index s.t. $R \subseteq N$ and $(\pi(G, a) \setminus \text{cl}(R)) \cap E^{\leq 2\ell} \cap N = \emptyset$. Thus $N \cap E^{\leq 2\ell} = \text{cl}(R) \cap E^{\leq 2\ell}$. \square

1.3.2 Granular covers

Granular covers preserve tours globally. We say that a cover $\varphi: \widehat{G} \xrightarrow{\text{unb}} G$ preserves a tour α if all lifts of α are also tours. If no lift of α is a tour then we say that the cover *unravels* α .

Definition 1.3.7. The set of all tours of a multidigraph G is denoted by $\pi_{\text{cl}}(G)$.

Note that $\pi_{\text{cl}}(G)$ is the disjoint union of the fundamental groups of $G = (V, E)$, i.e., $\pi_{\text{cl}}(G) = \bigcup_{a \in V} \pi(G, a)$. To express that α is preserved by $\varphi: \widehat{G} \xrightarrow{\text{unb}} G$ we can also just state that $\varphi^{-1}(\alpha) \subseteq \pi_{\text{cl}}(\widehat{G})$.

Definition 1.3.8. An R -granular cover of a multidigraph G with a set of tours $R \subseteq \pi_{\text{cl}}(G)$ is a cover $\varphi: \widehat{G} \xrightarrow{\text{unb}} G$ s.t.

$$\varphi^{-1}(R) \subseteq \pi_{\text{cl}}(\widehat{G}).$$

We give a characterisation of R -granular covers that might be more intuitive as it is more in line with the definition of pointed granular covers. A cover $\varphi: \widehat{G} \xrightarrow{\text{unb}} G$ is an R -granular cover, if and only if for all $a \in V$ and $\widehat{a} \in \varphi^{-1}(a)$,

$$R \cap \pi(G, a) \subseteq \varphi(\pi(\widehat{G}, \widehat{a})).$$

One can easily show this equivalence in a straightforward fashion. Simply assume one inclusion and then show the other and vice versa. However, we can also transform these two statements into each other by a sequence of simple equivalences. For this we have to subdivide the function $\varphi: \pi(\widehat{G}) \rightarrow \pi(G)$ into its injective components. For each $\widehat{a} \in \widehat{V}$ we let $\varphi_{\widehat{a}}$ be the restriction of φ to trails that start at \widehat{a} . Note that $\varphi_{\widehat{a}}$ is injective and that $\varphi^{-1}(a) = \bigcup_{\widehat{a} \in \varphi^{-1}(a)} \varphi_{\widehat{a}}^{-1}(a)$. We can now show the characterisation:

$$\begin{aligned} \varphi^{-1}(R) \subseteq \pi_{\text{cl}}(\widehat{G}) &\iff \varphi^{-1}(R \cap \pi_{\text{cl}}(G)) \subseteq \bigcup_{\widehat{a} \in \widehat{V}} \pi(\widehat{G}, \widehat{a}) \\ &\iff \varphi^{-1}(R \cap \bigcup_{a \in V} \pi(G, a)) \subseteq \bigcup_{a \in V} \bigcup_{\widehat{a} \in \varphi^{-1}(a)} \pi(\widehat{G}, \widehat{a}) \\ &\iff \bigcup_{a \in V} \varphi^{-1}(R \cap \pi(G, a)) \subseteq \bigcup_{a \in V} \bigcup_{\widehat{a} \in \varphi^{-1}(a)} \pi(\widehat{G}, \widehat{a}) \\ &\iff \varphi^{-1}(R \cap \pi(G, a)) \subseteq \bigcup_{\widehat{a} \in \varphi^{-1}(a)} \pi(\widehat{G}, \widehat{a}) \text{ for all } a \in V \\ &\iff \bigcup_{\widehat{a} \in \varphi^{-1}(a)} \varphi_{\widehat{a}}^{-1}(R \cap \pi(G, a)) \subseteq \bigcup_{\widehat{a} \in \varphi^{-1}(a)} \pi(\widehat{G}, \widehat{a}) \text{ for all } a \in V \\ &\iff \varphi_{\widehat{a}}^{-1}(R \cap \pi(G, a)) \subseteq \pi(\widehat{G}, \widehat{a}) \text{ for all } a \in V \text{ and } \widehat{a} \in \varphi^{-1}(a) \\ &\iff R \cap \pi(G, a) \subseteq \varphi_{\widehat{a}}(\pi(\widehat{G}, \widehat{a})) \text{ for all } a \in V \text{ and } \widehat{a} \in \varphi^{-1}(a) \\ &\iff R \cap \pi(G, a) \subseteq \varphi(\pi(\widehat{G}, \widehat{a})) \text{ for all } a \in V \text{ and } \widehat{a} \in \varphi^{-1}(a). \end{aligned}$$

Typically an R -granular covers preserves more tours than just the ones in R . We show that all R -granular covers necessarily preserve $\text{ncl}_G(R)$, the ‘normal subgroupoid of $\pi(G)$ generated by the R ’.

Definition 1.3.9. The *normal closure* $\text{ncl}_G(R)$ of a set $R \subseteq \pi_{\text{cl}}(G)$ is the smallest subset of $\pi(G)$ that contains the identities $\{\varepsilon_a \mid a \in V\}$ and R , and is closed under products and conjugation, i.e., if $\beta \in \pi(G)$ and $\alpha \in \text{ncl}_G(R)$ with $t(\beta) = s(\alpha)$, then $\beta \cdot \alpha \cdot \beta^{-1} \in \text{ncl}_G(R)$.

1 Acyclicity and covers

Note that the normal closure of G stays in $\pi_{\text{cl}}(G)$. An explicit description of $\text{ncl}_G(R)$ can be given by

$$\text{ncl}_G(R) = \{ \alpha_1 \cdot r_1 \cdot \alpha_1^{-1} \cdot \dots \cdot \alpha_n \cdot r_n \cdot \alpha_n^{-1} \mid \\ n \in \mathbb{N}, \alpha_i \in \pi(G), r_i \in R, \text{ s.t. } s(\alpha_i) = s(\alpha_j), t(\alpha_i) = s(r_i) \}.$$

Lemma 1.3.10. *Every R -granular cover is also an $\text{ncl}_G(R)$ -granular cover.*

Proof. Let $\varphi: \widehat{G} \xrightarrow{\text{unb}} G$ be an R -granular cover and let X be the set of tours that are preserved by this cover. Clearly $R \subseteq X$ and also $\varepsilon_a \in X$ for all vertices a of G . For the closure conditions we show that X is closed under conjugation; closure under products is shown similarly. Let $\alpha \in X$ and $\beta \in \pi(G)$ with $t(\beta) = s(\alpha)$. Consider a lift $\widehat{\beta}_0 \widehat{\alpha} \widehat{\beta}_1$ of $\beta \alpha \beta^{-1}$. Then $\widehat{\alpha}$ is a tour as it is a lift of α . Thus $\widehat{\beta}_1 = \widehat{\beta}_0^{-1}$. Thus $\widehat{\beta}_0 \widehat{\alpha} \widehat{\beta}_1 = \widehat{\beta}_0 \widehat{\alpha} \widehat{\beta}_0^{-1}$ is a tour as well. \square

Free and locally free granular covers

The basic results about free and locally free granular covers are the same as for pointed granular covers: the free granular cover is unique and can be defined in terms of universality or minimality of the set of preserved tours, and similarly, locally free granular covers can be defined in these two ways. However, as we show below, finite, locally free granular covers do not exist in some cases.

An R -granular cover $\psi: \widetilde{G} \xrightarrow{\text{unb}} G$ is universal if $(\varphi: \widehat{G}, \widehat{a} \xrightarrow{\text{unb}} G, a) \leq (\psi: \widetilde{G}, \widetilde{a} \xrightarrow{\text{unb}} G, a)$ for all R -granular covers $\varphi: \widehat{G} \xrightarrow{\text{unb}} G$ and choices of vertices \widetilde{a} and \widehat{a} .

Proposition 1.3.11. *Let $\psi: \widetilde{G} \xrightarrow{\text{unb}} G$ be an unbranched cover and $R \subseteq \pi_{\text{cl}}(G)$. Then the following are equivalent:*

(i) $\psi^{-1}(\text{ncl}_G(R)) = \pi_{\text{cl}}(\widetilde{G})$.

(ii) $\psi: \widetilde{G} \rightarrow G$ is a universal R -granular cover.

Furthermore such a universal cover exists and is unique, up to isomorphism.

Proof. By the discussion after Definition 1.3.8, (i) is equivalent to

$$\psi(\pi(\widetilde{G}, \widetilde{a})) = \text{ncl}_G(R) \cap \pi(G, a) \quad \text{for all } a \in V \text{ and } \widehat{a} \in \psi^{-1}(a). \quad (*)$$

We prove that a cover with (*) exists and that it is uniquely defined. If $\varphi: \widehat{G} \xrightarrow{\text{unb}} G$ is another cover with property (*) then $\varphi(\pi(\widehat{G}, \widehat{a})) = \psi(\pi(\widetilde{G}, \widetilde{a}))$ for any choice of \widehat{a} and \widetilde{a} with $\varphi(\widehat{a}) = \psi(\widetilde{a})$. Thus, by the Galois connection, $\varphi: \widehat{G} \xrightarrow{\text{unb}} G$ and $\psi: \widetilde{G} \xrightarrow{\text{unb}} G$ are isomorphic. This settles the uniqueness claim. Let c be any vertex in G . Then, by the Galois connection, there is a pointed cover $\psi: \widetilde{G}, \widetilde{c} \xrightarrow{\text{unb}} G, c$

with $\psi(\pi(\tilde{G}, \tilde{c})) = R \cap \pi(G, c)$. We show that $\psi: \tilde{G} \xrightarrow{\text{unb}} G$ has (*): let $a \in V$ and $\tilde{a} \in \psi^{-1}(a)$ be arbitrary. Using a trail $\tilde{\alpha}$ from \tilde{a} to \tilde{c} we get that

$$\begin{aligned} \psi(\pi(\tilde{G}, \tilde{a})) &= \psi(\tilde{\alpha} \cdot \pi(\tilde{G}, \tilde{c}) \cdot \tilde{\alpha}^{-1}) = \alpha \cdot \psi(\pi(\tilde{G}, \tilde{c})) \cdot \alpha^{-1} = \alpha \cdot (\text{ncl}_G(R) \cap \pi(G, c)) \cdot \alpha^{-1} \\ &= \alpha \cdot \text{ncl}_G(R) \cdot \alpha^{-1} \cap \alpha \cdot \pi(G, c) \cdot \alpha^{-1} = \text{ncl}_G(R) \cap \pi(G, a). \end{aligned}$$

It is straightforward to show that (*) is also equivalent to (ii). \square

We call the R -granular cover specified in the previous proposition the *free R -granular cover*.

An R -granular cover $\varphi: \hat{G} \xrightarrow{\text{unb}} G$ is ℓ -universal for $\ell \in \mathbb{N}$ if $(\psi: \tilde{G}, \tilde{a} \xrightarrow{\text{unb}} G, a) \leq_\ell (\varphi: \hat{G}, \hat{a} \xrightarrow{\text{unb}} G, a)$ for all R -granular covers $\psi: \tilde{G} \xrightarrow{\text{unb}} G$ and choices of $a \in V$ and $\tilde{a} \in \psi^{-1}(a)$, $\hat{a} \in \varphi^{-1}(a)$.

Proposition 1.3.12. *Let $\varphi: \hat{G} \xrightarrow{\text{unb}} G$ be an R -granular cover with $\hat{G} = (\hat{V}, \hat{E})$. Then the following are equivalent:*

(i) $\varphi^{-1}(\text{ncl}_G(R)) \cap \hat{E}^{\leq 2\ell} = \pi_{\text{cl}}(\hat{G}) \cap \hat{E}^{\leq 2\ell}$.

(ii) $\varphi: \hat{G} \rightarrow G$ is an ℓ -universal R -granular cover.

Proof. The proof is similar to the proof of Proposition 1.3.11. One can show that for R -granular covers (i) and (ii) are both equivalent to

$$\varphi(\pi(\hat{G}, \hat{a})) \cap E^{\leq 2\ell} = \text{ncl}_G(R) \cap \pi(G, a) \cap E^{\leq 2\ell} \quad \text{for all } a \in V \text{ and } \hat{a} \in \varphi^{-1}(a). \quad \square$$

An ℓ -free R -granular cover is an R -granular cover that has the property specified in the previous proposition.

Note that ordinary ℓ -acyclic covers are ℓ -free \emptyset -granular covers. It can be shown by a simple product construction with groups of high girth that ℓ -acyclic covers do exist [34]. Using the existing theory we can also provide a different take of this proof.

Lemma 1.3.13. *Every finite multidigraph has finite, ℓ -acyclic unbranched covers for any fixed degree $\ell \in \mathbb{N}$.*

Proof. We first note that there are finite P -generated groups G of arbitrary large girth ℓ , i.e. $[u]_G \neq 1$ for all $u \in \text{FG}(P)$ with $|u| \leq \ell$. There is a simple construction of these groups attributed to Biggs in [34]. This implies that there are normal subgroups N of finite index s.t. $N \cap P^{\leq \ell} = \{\varepsilon\}$ (these can also be obtained using Hall's Theorem (Theorem 1.3.5, [16]) with the fact that every subgroup of finite index of the free group contains a normal subgroup of finite index).

Let G be a finite multidigraph and d its girth. Choose any vertex c of G . By the discussing above there is a normal subgroup N of $\pi(G, c)$ s.t. $N \cap E^{\leq 2\ell+2d} = \{\varepsilon\}$.

1 Acyclicity and covers

Consider not the pointed cover $\varphi: \widehat{G}, \widehat{c} \xrightarrow{\text{unb}} G, c$ corresponding to N . Let \widehat{a} in \widehat{G} . Then there is a path \widehat{a} of length at most d to some \widehat{c}' . Then

$$\begin{aligned} \varphi(\pi(\widehat{G}, \widehat{a})) \cap E^{\leq 2\ell} &= \varphi(\widehat{a} \cdot \pi(\widehat{G}, \widehat{c}) \cdot \widehat{a}^{-1}) \cap E^{\leq 2\ell} = \alpha \cdot \varphi(\pi(\widehat{G}, \widehat{c})) \cdot \alpha^{-1} \cap \widehat{E}^{\leq 2\ell} \\ &= \alpha \cdot N \cdot \alpha^{-1} \cap E^{\leq 2\ell} \subseteq \alpha(N \cap E^{\leq 2\ell + 2d})\alpha^{-1} = \{\varepsilon\}. \end{aligned}$$

□

Finite, locally free granular covers and the Novikov–Boone Theorem

In Section 1.3.1 we used the Theorem of M. Hall to show that finite, locally free pointed granular covers generally exist. We recall the statement of M. Hall’s Theorem: in a free group F we can separate a finitely generated subgroup G and an element $x \in F \setminus G$ by a subgroup of finite index. In order to adapt this idea for finite, locally free granular covers we would need a variant of M. Hall’s Theorem for finitely generated normal subgroups. However, using an undecidability result of Novikov and Boone, we can show that such a variant is false, and ultimately that finite, locally free granular covers do not always exist.

We assume that the reader is familiar with basic computability theory in particular the terms decidability and semi-decidability.

A finite presentation of a P -generated group G is a tuple (P, R) where $R \subseteq P^*$ s.t. G is the ‘freest’ group subject to the conditions that $[r]_G = 1$ for all $r \in R$. We may assume that all elements of R are reduced since replacing an element $r \in R$ by its reduction $\text{red}(r)$ does not effectively change the condition $[r]_G = 1$. Now given a presentation (P, R) , its represented group is simply $\text{FG}(P)/\text{ncl}(R)$.

The ‘word problem’ for a presentation (P, R) is the problem to decide whether $u \in P^*$ evaluates to 1 in $\text{FG}(P)/\text{ncl}(R)$, or equivalently whether $[u]_{\text{FG}(P)} \in \text{ncl}(P)$. Thus, if we can determine the members of $\text{ncl}(P)$, we can decide the word problem for (P, R) .

The Novikov–Boone Theorem [9, 33] says that there is a finite presentation for which the word problem is undecidable. For our purposes the following formulation is more suitable.

Theorem 1.3.14 (Novikov–Boone). *There is a finite P and a finite $R \subseteq \text{FG}(P)$ s.t. $\text{ncl}(R) \subseteq \text{FG}(P)$ is undecidable.*

In general, $\text{ncl}(R)$ is semi-decidable: we can use an unbounded search to ‘check’ whether $u \in \text{ncl}(R) = \{u_1 \cdot r_1 \cdot u_1^{-1} \cdot \dots \cdot u_n r_n u_n^{-1} \mid n \in \mathbb{N}, u_i \in \text{FG}(P), r_i \in R\}$. Thus the computationally interesting part of the Novikov–Boone Theorem is that $\text{FG}(P) \setminus \text{ncl}(R)$ is not semi-decidable.

We can use the Theorem of Novikov–Boone to show that the normal variant of M. Hall’s Theorem is false.

Lemma 1.3.15. *There are finite P , finite $R \subseteq \text{FG}(P)$ and $u \in \text{FG}(P) \setminus \text{ncl}(R)$ s.t. there is no normal subgroup N of $\text{FG}(P)$ of finite index that contains R but not u .*

Proof. Let P and $R \subseteq \text{FG}(P)$ be finite. We show that, if for all $u \in \text{FG}(P) \setminus \text{ncl}(R)$ there is a normal subgroup N with

$$N \text{ has finite index, } R \subseteq N, u \notin R, \quad (*_u)$$

then $\text{FG}(P) \setminus \text{ncl}(P)$ is semi-decidable.

There is a bijection between normal subgroups $N \subseteq \text{FG}(P)$ and P -generated groups G via $N \mapsto \text{FG}(P)/N$, and this bijection translates $(*_u)$ into

$$G \text{ is finite, } [r]_G = 1 \text{ for all } r \in R, [u]_G \neq 1. \quad (**_u)$$

So in order to ‘check’ whether some $u \in \text{FG}(P)$ is not in $\text{ncl}(P)$ we can do an unbounded search for a P -generated group with property $(**_u)$.

Thus, for the presentation (P, R) of the Novikov–Boone Theorem there has to be a $u \notin \text{ncl}(R)$ s.t. there is no normal subgroup $N \subseteq \text{FG}(P)$ with property $(*_u)$. \square

With the same idea we can prove the next lemma.

Lemma 1.3.16. *Let $G = (V, E)$ be a finite multidigraph and $R \subseteq \pi_{\text{cl}}(G)$ a finite set of tours. Then $\text{ncl}_G(R)$ is decidable if for every $\ell \in \mathbb{N}$ there is a finite, locally ℓ -free granular cover of G .*

Proof. First note that $\text{ncl}_G(R)$ is semi-decidable for the same reason as in the group case.

Assuming that finite, ℓ -free granular covers of arbitrary degree $\ell \in \mathbb{N}$ exist, we show now that $\pi_{\text{cl}}(G) \setminus \text{ncl}_G(R)$ is also semi-decidable. For $u \in \pi_{\text{cl}}(G)$ ‘check’ by an unbounded search whether there is a finite, R -granular cover that unravels u . If $u \in \text{ncl}_G(R)$ then this search does not terminate as every R -granular cover preserves u . If $u \notin \text{ncl}_G(R)$ this search eventually encounters a finite, $|u|$ -locally free cover of G which unravels u . \square

Corollary 1.3.17. *There is a finite multidigraph G and finite $R \subseteq \pi_{\text{cl}}(G)$ s.t. there are no finite, ℓ -locally free R -granular covers for some $\ell \in \mathbb{N}$.*

Proof. Let (P, R) as in the Novikov–Boone Theorem. Let $G = (\{a\}, P)$ be a graph with one vertex. Then, by the previous lemma and the fact that $\text{ncl}_G(R) = \text{ncl}(R)$ is not decidable, G does not have finite, ℓ -free granular covers for all degrees $\ell \in \mathbb{N}$. \square

1.4 Acyclicity in hypergraphs

Hypergraphs are a generalisation of simple graphs. In a hypergraph an edge, or rather a hyperedge, can connect more than two vertices. Hypergraphs are commonly used in combinatorics and discrete mathematics and also appear in algebraic geometry in the form of abstract simplicial complexes.

Definition 1.4.1. A *hypergraph* consists of a set of vertices V and a collection S of finite, non-empty subsets of V , the *hyperedges*.




Figure 1.3: 3 examples of hypergraphs: the tetrahedron, the cartwheel and the guarded cartwheel. The tetrahedron has 4 vertices and 4 hyperedges each of which contains all but one vertex. The cartwheel can be understood as the tetrahedron with one hyperedge removed. It has 4 vertices one of which is the center vertex and 3 hyperedges each containing the center vertex and two additional vertices. The guarded cartwheel is like the cartwheel with one additional hyperedge that contains all 4 vertices.

We can cast a simple graph $G = (V, E)$ as a hypergraph $H = (V, S)$ simply by putting $S = \{ \{a, b\} \mid (a, b) \in E \}$.

A *guard* of a subset $X \subseteq V$ in a hypergraph $H = (V, S)$ is a hyperedge $s \in S$ that contains X . We call singleton sets and sets for which such a guard exist *guarded*. The *Gaifman graph* $G(H)$ of H is the simple graph (V, E) on the same domain as H in which two distinct vertices are connected by an edge if they are guarded.

We can lift graph theoretic notions to hypergraphs via the Gaifman graph. E.g., we say that $\alpha = a_0 \dots a_n$ is a walk in H if α is a walk in $G(H)$; similarly for trails, tours, cliques, neighbourhoods and so on. There is but one exception to this convention. We do not lift the notion of acyclicity in this manner as we provide a separate notion for hypergraph acyclicity (see the following subsection).

For $X \subseteq V$ the *hypergraph on X induced by H* is the hypergraph $H|_X = (X, \{ s \cap X \mid s \in S, s \cap X \neq \emptyset \})$. We call $H|_X$ an *induced subhypergraph* of H . A *weak subhypergraph* $H' = (V', S')$ of H is a hypergraph s.t. $V' \subseteq V$ and $S' \subseteq \{ s \cap V' \mid s \in S, s \cap V' \neq \emptyset \}$. The *localisation of H at a* is the particular weak subhypergraph $H_a := (N(a), \{ s \mid s \in S, a \in s \})$. For example the ‘cartwheel’ in Figure 1.3 is a localisation of the ‘tetrahedron’.

1.4.1 Definition of acyclicity for hypergraphs

A sensible notion of acyclicity for hypergraphs has to be a generalisation of ordinary graph acyclicity. However, this restriction does not determine a unique acyclicity notion for hypergraphs. There are multiple distinct acyclicity notions in the literature each of which is useful in its context.

Typical notions that appear in hypergraph theory are α -, β - and γ -acyclicity. The latter three originate from database theory. They are used as criteria for well-behaved database schemas but also find application in the theory of constraint satis-

fraction problems and finite model theory. For a comparison of the different acyclicity notions see [10].

We focus on α -acyclicity as this is the strongest acyclicity notion for which we know that we can generally obtain branched covers that satisfy this notion (see the discussion about the relationship of branched covers and acyclicity on page 66). From now on we call α -acyclicity *hypergraph acyclicity* and say that a hypergraph is *acyclic* if it is α -acyclic.

A common definition of hypergraph acyclicity consist of two conditions: chordality and conformality. The former constrains the Gaifman graph, the latter constrains the connection between Gaifman graph and hypergraph.

A hypergraph $H = (V, S)$ is *chordal* if its Gaifman graph is chordal. Recall that a simple graph is chordal if each cycle $a_0 \dots a_n$ of length at least 4 has a chord, i.e., there is an edge (a_i, a_j) for non-adjacent $i, j \in \mathbb{Z}_n$. A hypergraph $H = (V, S)$ is *conformal* if every finite clique $K \subseteq V$ is guarded, i.e., there is an $s \in S$ s.t. $K \subseteq s$. We also need localised versions of chordality and conformality that are defined in the obvious ways: a hypergraph is ℓ -*chordal* if each cycle of length at least 4 and at most ℓ has a chord. A hypergraph is ℓ -*conformal* if every clique of size up to ℓ is guarded.

Definition 1.4.2. A hypergraph is *acyclic* if it is conformal and chordal.

At first glance it seems reasonable to say that a hypergraph is ℓ -acyclic if its induced subhypergraphs of size at most ℓ are acyclic. This definition of ℓ -acyclicity is equivalent to ℓ -conformality and ℓ -chordality. However, by this definition the ‘cartwheel’ in Figure 1.3 would be 3-acyclic which seems quite odd. We get a more intuitive localised acyclicity notion if we require full conformality and not just ℓ -conformality.

Definition 1.4.3. A hypergraph is ℓ -*acyclic* ($\ell \geq 3$) if it is conformal and ℓ -chordal.

We can describe ℓ -acyclicity also in terms of induced subhypergraphs. A hypergraph $H = (V, S)$ is ℓ -acyclic if, and only if, all induced subhypergraphs on unions of ℓ hyperedges are acyclic, i.e., for all $s_1, \dots, s_\ell \in S$ the hypergraph $H|_{s_1 \cup \dots \cup s_\ell}$ is acyclic.

It is easy to see that acyclicity is preserved under passage to induced substructures, i.e., if $H = (V, S)$ is ℓ -acyclic then so is $H|_A$ for $A \subseteq V$. For weak subhypergraphs this is not true in general. Consider the ‘cartwheel’ and the ‘guarded cartwheel’ of Figure 1.3. The former is not 3-acyclic and is a weak subhypergraph of the latter which is acyclic. Then again, localisations do preserve acyclicity.

Lemma 1.4.4. *Localisations of ℓ -acyclic hypergraphs are ℓ -acyclic.*

Proof. Let $H = (V, S)$ be a hypergraph and H_c a localisation for $c \in V$.

Let K be a clique in H_c . Then $K \cup \{c\}$ is a clique in H and thus guarded by a hyperedge of H which is also present in H_c and thus forms a guard K in H_c .

Let α be a tour in H_c . Then this tour has a chord (a, b) in H . The vertices a, b, c form a clique in H and thus are contained in a hyperedge of H which guarantees that the chord (a, b) also exists in H_c . \square

1.4.2 Acyclicity as absence of induced cycles

Brault-Baron [10] introduces the notion of cycle-freeness, an acyclicity notion for hypergraphs that is weaker than hypergraph acyclicity. We give a connection between cycle-freeness and hypergraphs that foreshadows a similar connection between simple connectivity and hypergraph acyclicity that is given in Theorem 1.5.9 below. We also discuss triangulations that play an important role in the proof of said theorem.

Definition 1.4.5. A hypergraph is ℓ -cycle-free if it is ℓ -chordal and 3-conformal. It is cycle-free if its is chordal and 3-conformal.

Clearly ℓ -cycle-freeness is weaker than ℓ -acyclicity. The tetrahedron in Figure 1.3 is an example of a hypergraph that is cycle-free but not 3-acyclic.

We can characterise ℓ -cycle-free hypergraphs as hypergraphs in which closed walks can be triangulated.

Definition 1.4.6. A *triangulation* of a closed walk $a_0 \dots a_n$ in a hypergraph H is a non-empty set $T \subseteq \{s \subseteq \mathbb{Z}_n \mid |s| = 3\}$ s.t. for all $s \in T$

- (i) $\{a_i \mid i \in s\}$ is guarded,
- (ii) if $i, j \in s$ are non-adjacent, then there is a distinct $s' \in T$ s.t. $i, j \in s'$,
- (iii) if $i, j \in s$ are non-adjacent, then for all $s' \in T$ either $s' \subseteq \{i, i+1, \dots, j-1, j\}$ or $s' \subseteq \{j, j+1, \dots, i-1, i\}$.


Lemma 1.4.7. A hypergraph is ℓ -cycle-free if, and only if, all closed walks of length at least 3 and at most ℓ can be triangulated.

Proof. We start with the implication from right to left. Let H be a hypergraph in which closed walks of length at least 3 and at most ℓ have triangulations. Since all triangles have a triangulation, H is 3-conformal. H is also ℓ -chordal as triangulations of cycles of length at least 4 induce chords on that cycle.

Now we show the converse. We show by induction that a closed walk $\alpha = a_0 \dots a_n$ of length at least 3 and at most ℓ has a triangulation.

If $\alpha = a_0 a_1 a_2 a_0$ then $\{a_0, a_1, a_2\}$ is a clique of size 3 and thus has a guard. So $\{\{0, 1, 2\}\}$ is a triangulation of α .

For the induction step it suffices to show that there is some guarded set $\{a_i, a_j, a_k\}$ ($0 \leq i < j < k \leq n-1$).



Then we can combine the triangulations of $a_i a_{i+1} \dots a_j a_i$, $a_j a_{j+1} \dots a_k a_j$ and $a_k a_{k+1} \dots a_i a_k$ together with $\{a_i, a_j, a_k\}$ to a triangulation of α .

If α is not a cycle, then $a_i = a_j$ for some $i, j \in \mathbb{Z}_n$. So $\{a_i, a_{i+1}, a_j\}$ is guarded. If α is a cycle, then it has a chord (a_i, a_j) . Let this chord be s.t. its distance

$d = \min\{|a_i a_{i+1} \dots a_j|, |a_j a_{j+1} \dots a_i|\}$ is minimal. W.l.o.g. $d = |a_i a_{i+1} \dots a_j|$. Then $d = 2$, as we could otherwise find a shorter chord in $a_i \dots a_j a_i$. Then $\{i, i+1, j\}$ is a clique and thus is guarded. \square

An ℓ -cycle-free hypergraph is ℓ -acyclic if it is conformal. The following lemma is helpful to boost 3-conformality to full conformality.

Lemma 1.4.8. *A 3-conformal hypergraph for which all localisations are ℓ -conformal is $\ell+1$ -conformal.*

Proof. Let $H = (V, S)$ be as in the premise of the lemma and let K be a clique in H of size $\ell + 1$. Pick any $a \in K$. We observe that $K \setminus \{a\}$ is also a clique in H_a as for any two vertices $b, c \in K \setminus \{a\}$ the set $\{a, b, c\}$ forms a 3-clique in H and thus its guard induces a connection of b and c in H_a . Now $K \setminus \{a\}$ is contained in some hyperedge in H_a which has to contain a as well and thus K is guarded in H . \square

We say that H' is an iterated localisation of H if there is a sequence of hypergraphs H_0, \dots, H_n s.t. $H_0 = H$, $H_n = H'$ and H_{i+1} is a localisation of H_i . Note that H itself is its 0th iterated localisation. Iterated localisations can also be described as multi-localisations $H_A = \{N(A) \mid \{s \in S \mid A \subseteq s\}\}$ where A is a guarded vertices of H and $N(A) := \{b \mid \{b\} \cup A \text{ is guarded}\}$.

Proposition 1.4.9. *A hypergraph is ℓ -acyclic if, and only if, all its iterated localisations are ℓ -cycle-free.*

Proof. If H is ℓ -acyclic then by Corollary 1.4.4 all its iterated localisations are ℓ -acyclic. So they are ℓ -cycle-free. This settles the ‘only if’.

For the converse we need to show that for all $k \in \mathbb{N}$ the iterated localisations of H are k -conformal. By assumption all iterated localisations are 3-conformal. Using Lemma 1.4.8 we can boost any degree k of conformality established for all iterated localisations to the next degree $k + 1$. \square

1.4.3 Hypercycles

There are numerous characterisations of hypergraph acyclicity in the literature. In [8] alone, Beeri, Fagin, Maier and Yannakakis list 12 of them. In [51] one such characterisation via the absence of some type of cyclic configuration is given.⁴ Even if not new this treatment here is far more direct and we also discuss in detail the connection to ordinary cycles.

Let $H = (V, S)$ be a hypergraph:

- A *hypertriangle* is a triple of hyperedges $(s_i)_{i \in \mathbb{Z}_3} \subseteq S$ such that there is no $s \in S$ with $\bigcap_{i \in \mathbb{Z}_3} s_i \cap s_{i+1} \subseteq s$.

⁴At the time of writing this thesis this article was not known to me (it is also not mentioned in the recent overview article [10]) and I thought that alpha-cycles were a genuinely new notion. For that reason this section is given such prominence.

1 Acyclicity and covers

- A *hyperchord* of a sequence $(s_i)_{i \in \mathbb{Z}_\ell} \subseteq S$ is a set s such that $s_i \cap s_{i+1}, s_j \cap s_{j+1} \subseteq s$ for non-adjacent $i, j \in \mathbb{Z}_\ell$.
- A *hypercycle* is a sequence $(s_i)_{i \in \mathbb{Z}_n} \subseteq S$ of length at least 3 s.t. there are no three distinct $i, j, k \in \mathbb{Z}_n$ and $s \in S$ with $s_i \cap s_{i+1}, s_j \cap s_{j+1}, s_k \cap s_{k+1} \subseteq s$.
- A *hypertour* is a sequence $(s_i)_{i \in \mathbb{Z}_n} \subseteq S$ s.t. there are no $i \in \mathbb{Z}_n$ and $s \in S$ with $s_i \cap s_{i+1}, s_{i+1} \cap s_{i+2}, s_{i+2} \cap s_{i+3} \subseteq s$.

We note some easy observations about these definitions: first, the definitions of hypertriangles, hypercycles of length 3 and hypertours of length 3 are all equivalent. Second, hypertours are at least of length 3, otherwise s_0 would be a witness that $(s_i)_{i \in \mathbb{Z}_2}$ is not a hypertour, as $s_0 \cap s_1, s_1 \cap s_0, s_0 \cap s_1 \subseteq s_0$. Third, the intersection of s_i and s_{i+1} in a hypercycle or hypertour $(s_i)_{i \in \mathbb{Z}_n}$ is non-empty, otherwise the intersections $s_i \cap s_{i+1}, s_{i+1} \cap s_{i+2}, s_{i+2} \cap s_{i+3}$ would be guarded by s_{i+2} and thus $(s_i)_{i \in \mathbb{Z}_n}$ were not a hypertour (and also not a hypercycle).

We give three properties of hypergraphs that are based on forbidding these cyclic configurations. In Theorem 1.4.10 we show that these three properties all are equivalent to hypergraph ℓ -acyclicity.

- (A1 $^\ell$) H has no hypercycles of length at most ℓ .
- (A2 $^\ell$) H has no hypertours of length at most ℓ .
- (A3 $^\ell$) H has no hypertriangles, and every sequence $(s_i)_{i \in n}$ of hyperedges of length at least 4 and at most ℓ has a hyperchord.

We want to argue that the definitions of hypertriangles, hyperchords, hypercycles and hypertours are natural by showing that they perfectly correspond to their counterparts over simple graphs, i.e., hypertriangles correspond to triangles over simple graphs, chordless sequences of edges to chordless cycles and so on. So, the properties (A1 $^\ell$)–(A3 $^\ell$) for hypergraphs are natural generalisations of the properties (A1 $^\ell$)–(A3 $^\ell$) for graphs (see Section 1.1 p. 22). We discuss this correspondence in the example of hypertours.

Let $G = (V, E)$ be a simple graph and $H = (V, S)$ this simple graph cast as a hypergraph. For a non-trivial tour $a_0 \dots a_n$ in G we define its translation into a hypertour in H as $T(a_0 \dots a_n) = (\{a_{i-1}, a_i\})_{i \in \mathbb{Z}_n}$. We show now that T is indeed a bijection between tours and hypertours.

First, note that $T(a_0 \dots a_n) = (s_i)_{i \in \mathbb{Z}_n}$ is indeed a hypertour as for each $i \in \mathbb{Z}_n$

$$\bigcup_{j=i-1}^{i+1} s_j \cap s_{j+1} = \bigcup_{j=i-1}^{i+1} \{a_{j-1}, a_j\} \cap \{a_j, a_{j+1}\} = \{a_{i-1}, a_i, a_{i+1}\}.$$

So no $s \in S$ can contain three consecutive intersections.

Now we need to show that a hypertour $(s_i)_{i \in \mathbb{Z}_n}$ has a preimage under this translation. Note that $|s_i \cap s_{i+1}| = 1$: the intersection of s_i and s_{i+1} in a hypertour cannot be




Figure 1.4: Examples of the three basic obstacles why a sequence of edges $(e_i)_{i \in \mathbb{Z}_n}$ in a simple graph does not describe a tour: jumping, repetition, not being reduced. In each of the three examples e_2 is a witness that these sequences are also no hypertours.

empty, and if $|s_i \cap s_{i+1}| = 2$ then $s_i = s_{i+1}$ and thus $s_{i-1} \cap s_i, s_i \cap s_{i+1}, s_{i+1} \cap s_{i+2} \subseteq s_i$. We let a_i be the the unique element in $s_i \cap s_{i+1}$. Then clearly $T(a_0 \dots a_n) = (s_i)_{i \in \mathbb{Z}_n}$. We have to check that $a_0 \dots a_n$ is indeed a reduced walk. For that we have to rule out the following cases: a_i and a_{i+1} are not connected (jumping), $a_i = a_{i+1}$ (repetition), $a_{i-1} = a_{i+1}$ (not being reduced). Jumping does not occur as $a_i, a_{i+1} \in s_i$. Repetition does not occur, as otherwise $s_i \cap s_{i+1} = \{a_i\} = \{a_{i+1}\} = s_{i+1} \cap s_{i+2}$ and thus $s_i \cap s_{i+1}, s_{i+1} \cap s_{i+2}, s_{i+2} \cap s_{i+3} \subseteq s_{i+2}$. Also $a_0 \dots a_n$ is reduced since if $a_{i-1} = a_{i+1}$, then $s_i = \{a_{i-1}, a_i\} = \{a_i, a_{i+1}\} = s_{i+1}$. The whole argument why hypertours describe tours over simple graphs is also concisely illustrated in Figure 1.4.

Now we show that $(A1^\ell)$ – $(A3^\ell)$ indeed capture hypergraph ℓ -acyclicity.

Theorem 1.4.10. *Let H be a hypergraph and $\ell \geq 3$. Then the following are equivalent*

- (i) H is ℓ -acyclic.
- (ii) H has no hypertriangles and every sequence $(s_i)_{i \in \mathbb{Z}_n}$ of hyperedges of length at least 4 and at most ℓ has a hyperchord.
- (iii) H has no hypercycles of length at most ℓ .
- (iv) H has no hypertours of length at most ℓ .


As a preparation for the proof we provide a graph theoretic lemma that basically ensures the existence of ‘hyperchords’ in conformal and chordal hypergraphs.

Lemma 1.4.11. *Let $G = (V, E)$ be an ℓ -chordal graph, $n \in \{4, \dots, \ell\}$ and $(V_i)_{i \in \mathbb{Z}_n} \subseteq V$ be a cyclically indexed sequence of sets of vertices s.t. $V_i \cup V_{i+1}$ is a clique for all $i \in \mathbb{Z}_n$. Then there exist non-adjacent $i, j \in \mathbb{Z}_n$ s.t. $V_i \cup V_j$ is a clique.*

1 Acyclicity and covers


Proof. We proceed by contradiction and assume that for all non-adjacent i, j the set $V_i \cup V_j$ is not a clique.

The first claim is that there is an $a \in V_0$ s.t. for each $i \in \{2, \dots, n-2\}$ there is an $a_i \in V_i$ that is not adjacent to a , i.e., $(a, a_i) \notin E$. For this we prove by induction over k that there are $a \in V_0$ and $a_i \in V_i$ for $i \in \{2, \dots, k\}$ with the desired properties. For $k = 2$ this is clear, otherwise $V_0 \cup V_2$ would form a clique. For the induction step we assume that the claim holds for some $k < n - 2$. So there are $a \in V_0$ and $a_i \in V_i$ for $i \in \{2, \dots, k\}$ s.t. a and a_i are not adjacent. Furthermore, there are non-adjacent $b \in V_0$ and $a_{k+1} \in V_{k+1} \setminus V_0$, otherwise $V_0 \cup V_{k+1}$ would be a clique. If a is not adjacent to a_{k+1} or b is not adjacent to all the a_i we are done. So we assume that a is adjacent to a_{k+1} and b is adjacent to some a_j with $j \in \{2, \dots, k\}$. Let j be the largest index in $\{2, \dots, k\}$ s.t. b is adjacent to a_j . We give a sketch of the current situation where we indicate the information of being adjacent and being non-adjacent by lines and crossed-out lines. Note that for connectivity reasons a, b, a_j and a_{k+1} are pairwise distinct.



We consider now the closed walk $\alpha = baa_{k+1} \dots a_j b$. W.l.o.g. we may assume that α is a cycle as we can just pass to the cycle in α that contains a_j, b, a and a_{k+1} (note that a, b, a_j and a_{k+1} are pairwise distinct). Regarding the vertices appearing in α the vertex b is only adjacent to a and a_j and a is only adjacent to b and a_{k+1} . Thus no chord of α includes a or b . Since G is ℓ -chordal, α has a chord. This chord splits α in two cycles and one of these contains a, b, a_j and a_{k+1} . Iteratively drawing chords and passing to smaller and smaller cycles we eventually obtain the cycle $baa_{k+1}a_jb$. However, this cycle does not have a chord. So the assumption that a or b is adjacent to some of the a_i for $i \in \{2, \dots, k+1\}$ leads to a contradiction, hence a or b is not adjacent to all the a_i . This concludes the proof of the claim.

Now we can finish the proof. Let $a_0 \in V_0$ and $a_i \in V_i$ for $i \in \{2, \dots, n-2\}$ such that a_0 is not adjacent to any of the a_i . Also let $a_1 \in V_1$ and $a_{n-1} \in V_{n-1}$ be non-adjacent, distinct vertices, which exist as otherwise $V_1 \cup V_{n-1}$ would be a clique. We thus have the following configuration:




Note that a_{n-1}, a_0 and a_1 are pairwise distinct. We consider the cycle that contains a_{n-1}, a_0 and a_1 within the walk $a_0, a_1, \dots, a_{n-2}, a_{n-1}$. Similar as above, we can iteratively shorten this cycle by drawing chords and keep a_{n-1}, a_0, a_1 as vertices of the shorter cycles. However, since a_1 and a_{n-1} are not adjacent we this process has to stop before we have reached a triangle.

This provides the final contradiction and so the initial assumption that all the $V_i \cup V_j$ are not cliques is wrong. \square

We can now prove Theorem 1.4.10.

Proof of Theorem 1.4.10. (i) \Rightarrow (ii): Let $(s_i)_{i \in \mathbb{Z}_3}$ a sequence of length 3. Then $(s_0 \cap s_1) \cup (s_1 \cap s_2) \cup (s_2 \cap s_0)$ is a clique and hence by conformality of H guarded by some hyperedge. Now let $3 < n \leq \ell$. We set $V_i := s_i \cap s_{i+1}$ and apply Lemma 1.4.11. We get two non-adjacent indices $i, j \in \mathbb{Z}_n$ s.t. $V_i \cup V_j$ is a clique which by conformality is guarded by some hyperedge s which is a hyperchord since $s_i \cap s_{i+1}, s_j \cap s_{j+1} \subseteq s$.

(ii) \Rightarrow (iii): We prove by induction that every sequence $(s_i)_{i \in \mathbb{Z}_n}$ of length $3 \leq n \leq \ell$ does have the desired property. For $n = 3$ this is clear. So let $n > 3$. Then there is a hyperchord s.t. $s_i \cap s_{i+1}, s_j \cap s_{j+1} \subseteq s$.




Now we consider the sequence $(s'_i)_{i \in \mathbb{Z}_{n'}}$ that traverses cyclically the hyperedges s, s_{i+1}, \dots, s_j . By induction hypothesis there is an s' and distinct indices i', j', k' s.t. $s'_{i'} \cap s'_{i'+1}, s'_{j'} \cap s'_{j'+1}, s'_{k'} \cap s'_{k'+1} \subseteq s'$. But since $s_i \cap s_{i+1} \subseteq s_{i+1} \cap s$ and $s_j \cap s_{j+1} \subseteq s \cap s_j$ we can also find for this particular s' indices i, j, k s.t. $s_i \cap s_{i+1}, s_j \cap s_{j+1}, s_k \cap s_{k+1} \subseteq s'$.

(iii) \Rightarrow (iv): If the sequence $(s_i)_{i \in \mathbb{Z}_n}$ has length $n \leq 2$ then the condition of (iv) does hold trivially. For $n \geq 3$ one can use (iii) to find a ‘hypertriangulation’ of $(s_i)_{i \in \mathbb{Z}_n}$ which necessarily has a witness showing that $(s_i)_{i \in \mathbb{Z}_n}$ is not a hypercycle.

(iv) \Rightarrow (i): We show that the absence of hypertours of length up to ℓ implies conformality and ℓ -chordality. Let K be a clique in H . Let s_0 be a hyperedge whose

1 Acyclicity and covers

intersection with K is maximal, i.e., there is no hyperedge s' with $s_0 \cap K \subset s' \cap K$. We show that the assumption that $K \not\subseteq s_0$ leads to a contradiction. So assume that there is a $a \in K \setminus s_0$. Let s_1 be a hyperedge that contains a and has maximal intersection with $K \cap s_0$. By maximality of s_0 there is a $b \in (s_0 \cap K) \setminus s_1$ and since V is a clique there is a hyperedge s_2 containing v and w . So we have the following constellation.



Since s_0, s_1, s_2 , do not form a hypertour there is an $s \in S$ s.t. $s_0 \cap s_1, s_1 \cap s_2, s_2 \cap s_0 \subseteq s$. Thus $s_1 \cap (s_0 \cap K)$ is strictly contained in $s \cap (s_0 \cap K)$ which gives the desired contradiction.

Towards ℓ -chordality consider a cycle $a_0 \dots a_n$ of length greater than 3 and at most ℓ . Choose hyperedges $(s_i)_{i \in \mathbb{Z}_n}$ s.t. $(a_i, a_{i+1}) \subseteq s_i$. Then a witness showing that $(s_i)_{i \in \mathbb{Z}_n}$ is not a hypertour also gives a chord. \square

1.4.4 Acyclicity as decomposability

Trees can be characterised by some form of decomposability. This is also true for acyclic hypergraphs [8]. We reprove this result using the characterisation of hypergraph acyclicity as the absence of hypertours.

A *leaf* of a simple graph is a vertex that has exactly one adjacent vertex. A finite simple graph is a tree, if and only if it can be reduced to a single vertex by iteratively removing leaves.

This characterisation of trees relies on two basic facts. First, every non-trivial, connected, finite, acyclic graph has a leaf and second, removing leaves does not make an acyclic graph cyclic.

Graham [15], and Yu and Ozsoyoglu [53] independently generalised this idea to hypergraphs. The corresponding ‘reduction algorithm’ is named accordingly GYO algorithm (see [1] for applications of this algorithm).

Definition 1.4.12. A *hyperleaf* of a hypergraph $H = (V, S)$ is a hyperedge $s \in S$ for which there is a distinct hyperedge $s' \in S$ that guards the intersection of s with any distinct hyperedge $s'' \in S$, i.e., $s \cap s'' \subseteq s \cap s'$ for all $s'' \in S$ with $s \neq s''$.

It is easy to see that removing a hyperleaf s does not make an acyclic hypergraph H cyclic. If s' is a witness that s is a hyperleaf, then for all sequences of hyperedges for which s is a witness that this sequence is not a hypertour s' is also such a witness.

It is a little harder to show that every finite, acyclic hypergraph with at least two hyperedges has a leaf. The principal idea is the same as for simple graphs: assume that a simple graph does not have leaves. Then we can construct a tour by starting at some vertex and stepwise adding adjacent vertices while avoiding backtracking. As the graph is finite, this reduced walk has to intersect itself at some vertex. This then creates a tour.

Lemma 1.4.13. *Every finite, acyclic hypergraph with at least two hyperedges has a hyperleaf.*

Proof. For this proof we use the following terminology. We say that s' has *maximal intersection with s* if $s' \neq s$ and $s \cap s'$ is not properly contained in any hyperedge besides s . We say that $X \subseteq s$ is *maximal within s* if $X = s \cap s'$ is the intersection of s with some hyperedges s' that has maximal intersection with s .

We show the statement by induction on the number of hyperedges. The base case is trivial. For the induction step we assume for the sake of contradiction that $H = (V, S)$ is acyclic and does not have a hyperleaf.

We show below that this assumption leads to the existence of a sequence of hyperedges $(s_i)_{i \in \mathbb{Z}_n}$ where s_{i+1} has maximal intersection with s_i , and $s_{i-1} \cap s_i$ and $s_i \cap s_{i+1}$ are incomparable. Since H is acyclic, there are $s \in S$ and $j \in \mathbb{Z}_n$ s.t. $\bigcap_{i=j-1}^{j+1} s_i \cap s_{i+1} \subseteq s$. Then $s_{j+1} \cap s_{j+2}$ is a proper subset of s as $s_j \cap s_{j+1}$ and $s_{j+1} \cap s_{j+2}$ are incomparable and both contained in s . Also $s \neq s_{j+1}$ as $s_{j-1} \cap s_j$ and $s_j \cap s_{j+1}$ are not comparable and so $s_{j-1} \cap s_j \not\subseteq s_{j+1}$ but $s_{j-1} \cap s_j \subseteq s$. This contradicts that s_{j+2} has maximal intersection with s_{j+1} .

Now we show that such a sequence $(s_i)_{i \in \mathbb{Z}_n}$ exists, assuming that H is acyclic and without a leaf. Let s_0 be any hyperedge of H and $X_0 \subseteq s_0$ maximal within s_0 . We show that there are $s_1 \in S$, $X_1 \subseteq s_1$ s.t. $X_0 = s_0 \cap s_1$, X_1 is maximal within s_1 , and X_0 and X_1 are incomparable: let $H_{X_0} := (\{a \in V \mid \{a\} \cup X_0 \text{ is guarded}\}, \{s \in S \mid X_0 \subseteq s\})$ be the multi-localisation of H . Then H_{X_0} has fewer hyperedges than H as there has to be a hyperedge that does not contain X_0 since otherwise s_0 would be a hyperleaf. Also H_{X_0} is acyclic, since localisations of acyclic hypergraphs are acyclic. So, by induction hypothesis, H_{X_0} has a hyperleaf s_1 . As s_1 is not a hyperleaf in H , there is a hyperedge that has maximal intersection with s_1 but does not contain X_0 .

We can continue this process, generating a sequence $s_0, X_0, s_1, X_1, s_2, X_2, \dots$ s.t. X_i and X_{i+1} are incomparable and $X_i = s_i \cap s_{i+1}$ is maximal within s_i . We can now forget about the X_i and obtain a sequence s_0, s_1, \dots s.t. s_{i+1} has maximal intersection with s_i , and $s_{i-1} \cap s_i$ and $s_i \cap s_{i+1}$ are incomparable. Eventually, as H is finite, this sequence has to repeat an entry. W.l.o.g. $s_n = s_0$. Note that $n > 2$, as otherwise $s_0 \cap s_1$ and $s_1 \cap s_2$ are not incomparable.

The sequence $(s_i)_{i \in \mathbb{Z}_n}$ has nearly the desired property, but we cannot guarantee that $s_{n-1} \cap s_0$ and $s_0 \cap s_1$ are incomparable. It could be that $s_{n-1} \cap s_0 \subseteq s_0 \cap s_1$. Then $s_{n-1} \cap s_0 = s_{n-1} \cap s_1$ and so we can consider the shorter sequence s_1, \dots, s_{n-1} .

1 Acyclicity and covers

We can repeat this process and by the remark above the resulting sequences is at least of length 3. \square

Corollary 1.4.14 (Graham, Yu and Ozsoyoglu). *A finite, connected hypergraph is acyclic if, and only if, it can be reduced to a hypergraph with one or no hyperedge by iteratively removing hyperleaves.*

1.5 Simply connected hypergraphs

Simple connectivity is an acyclicity notion for topological spaces. A topological space is simply connected if every loop is homotopic to a trivial loop, i.e., every loop can be continuously transformed to a stationary one. In this section we give a combinatorial description of homotopy and simple connectivity for hypergraphs and discuss a class of hypergraphs over which simple connectivity and hypergraph acyclicity do coincide.

If we think of the hyperedges as ‘solid’ building blocks of a topological realisation of a hypergraph, then a combinatorial description of homotopy should allow transformations of walks within a hyperedge. We capture this idea by the notion of basic transformations. A *basic transformation* of a walk is a folding or unfolding of that walk. A *folding* is performed by replacing a subwalk of the form abc with ac provided that $\{a, b, c\}$ is guarded. In the special case that $a = c$ we instead replace aba with a . An *unfolding* is the reverse operation of a folding.

We define a congruence \approx on the walks in a hypergraph:

$$\alpha \approx \beta \quad :\iff \quad \alpha \text{ can be transformed into } \beta \text{ by basic transformations.}$$

A closed walk is *contractible* if it is equivalent to a trivial walk w.r.t. \approx .

Definition 1.5.1. A connected hypergraph is *simply connected* if all closed walks are contractible.

As stated, \approx is a congruence on the category of walks in a hypergraph H . It is straightforward to verify that this category modulo \approx forms a groupoid. we call this groupoid the *edge path groupoid* $E(H)$. The local groups of $E(H)$ at $a \in V$ are the so-called *edge path groups* $E(H, a)$ [40].

We want to characterise contractibility and simple connectivity by a normal closure operation (recall that $\pi_{\text{cl}}(G)$ is the set of all tours in G and for $R \subseteq \pi_{\text{cl}}(G)$, $\text{ncl}_G(R)$ is the normal closure of R in $\pi(G)$, see Definition 1.3.7 and Definition 1.3.9).

Lemma 1.5.2. *Let H be a hypergraph, G its Gaifman graph and $R = \{abca \in \pi_{\text{cl}}(G) \mid \{a, b, c\} \text{ is guarded in } H\}$. Then*

$$\alpha \text{ is contractible} \quad \iff \quad \text{red}(\alpha) \in \text{ncl}_G(R)$$

for closed walks α in H .

1.5 Simply connected hypergraphs

Proof. We show ‘ \implies ’ by induction on the length of the contraction process. If no basic transformation is required then α is a trivial walk and thus $\alpha \in \text{ncl}_G(R)$. Assume that α is the result of a folding of α' and $\text{red}(\alpha') \in \text{ncl}_G(R)$. Say $\alpha' = \alpha_1 \cdot acb \cdot \alpha_2$ and $\alpha = \alpha_1 \cdot ab \cdot \alpha_2$ then

$$\begin{aligned} \text{red}(\alpha) &= \text{red}(\alpha_1 \cdot ab \cdot \alpha_2) = \text{red}(\alpha_1 \cdot abca \cdot \alpha_1^{-1} \cdot \alpha_1 \cdot acb \cdot \alpha_2) \\ &= \text{red}(\alpha_1 \cdot abca \cdot \alpha_1^{-1}) \cdot \text{red}(\alpha') \in \text{ncl}_G(R). \end{aligned}$$

If α is the result of an unfolding of α' the argument is similar.

Note that $\text{red}(\alpha)$ is contractible, if and only if α is contractible. So for the converse, it suffices to show that every element in $\text{ncl}_G(R)$ is contractible. Obviously, all trivial tours are contractible. Also all tours $abca \in R$ are contractible as $acba \approx aca \approx a$. Being contractible is preserved under multiplication and conjugation. From this we get that $\text{ncl}_G(R)$ is a subset of the set of all contractible tours. \square

As a corollary we get a characterisation of simple connectivity.

Corollary 1.5.3. *A connected hypergraph H is simply connected if, and only if,*

$$\text{ncl}_G(R) = \pi_{\text{cl}}(G)$$

where G is the Gaifman graph of H and $R = \{abca \mid \{a, b, c\} \text{ is guarded}\}$.

We also introduce the notion of local simple connectivity.

Definition 1.5.4. A hypergraph is ℓ -*simply connected* if all closed walks of length at most 2ℓ are contractible.

Note that local simple connectivity is not strictly a local notion as the contraction-process of a short closed walk can have intermediate steps with closed walks of arbitrary length.

We can characterise local simple connectivity similar to simple connectivity.

Corollary 1.5.5. *A connected hypergraph H is ℓ -simply connected if, and only if,*

$$\text{ncl}_G(R) \cap V^{2\ell+1} = \pi_{\text{cl}}(G) \cap V^{2\ell+1}$$

where $G = (V, E)$ is the Gaifman graph of H and $R = \{abca \mid \{a, b, c\} \text{ is guarded}\}$.

The next theorem (see [40] for a proof) says that \approx is indeed a combinatorial description of homotopy. We just state the theorem here as we only need it as justification/motivation for the definition of simple connectivity for hypergraphs.

The geometric realisation $|H|$ of a hypergraph H is defined as for abstract simplicial complexes; $|H|$ is a quotient space of a collection of simplices, one for each hyperedge, and the equivalence relation is induced by the intersections of the hyperedges.

Theorem 1.5.6. *Let H be a hypergraph and $|H|$ its geometric realisation. Then*

$$E(H, a) \simeq \pi(|H|, a)$$

for all $a \in V$.

Let α be a walk in a hypergraph H . Then each reduction step of the reduction process for $\text{red}(\alpha)$ is a special case of a folding. Thus $\text{red}(\alpha) \approx \alpha$. For simple graphs G , reduction steps are the only folding operations possible and thus $\alpha \approx \beta \iff \text{red}(\alpha) = \text{red}(\beta)$ for walks α, β in G . It is no surprise then that $E(G) \rightarrow \pi(G); [\alpha]_{\approx} \mapsto \text{red}(\alpha)$ is an isomorphism of the edge path groupoid and the fundamental groupoid of G .

It is a well-known fact of algebraic topology that the word problem for finite simplicial complexes is not decidable in general, i.e., there is a simplicial complex for which it is not decidable whether a closed walk is contractible or not (see e.g. [45]). This is proved by reducing the word problem for finitely represented groups to the word problem for finite simplicial complexes, and the latter is undecidable by the Novikov–Boone Theorem. We state this result for later reference.

Theorem 1.5.7. *There is a finite hypergraph H for which the word problem*

$$\{ \alpha \in \pi_{\text{cl}}(H) \mid \alpha \text{ is contractible} \} \subseteq \pi_{\text{cl}}(H)$$

is undecidable.

1.5.1 Simple connectivity vs. hypergraph acyclicity

In this section we discuss how simple connectivity can be related to hypergraph acyclicity. Every acyclic, connected hypergraph is simply connected (see Theorem 1.5.9) but the converse is false. For example the ‘cart wheel’ in Figure 1.3 is simply connected but not acyclic. Over apex acyclic hypergraphs, however, these two notions are equivalent.

Definition 1.5.8. A hypergraph H is *apex acyclic* if all localisations of H are acyclic.

Apex acyclic hypergraphs can be characterised as hypergraphs in which every hypertour does not have a commonly shared element, i.e., $\bigcap_{i \in \mathbb{Z}_n} s_i = \emptyset$ for hypertours $(s_i)_{\mathbb{Z}_n}$.

Theorem 1.5.9. *The following are equivalent for a connected hypergraph H for $\ell \geq 3$*

- (i) H is ℓ -acyclic
- (ii) H is apex acyclic and ℓ -simply connected.

Proof. We show that (i) and (ii) are both equivalent to

(iii) H is apex acyclic and ℓ -cycle-free.

If H is ℓ -acyclic then it is clearly apex acyclic and ℓ -cycle-free. If H is ℓ -cycle-free and apex acyclic, then, in particular, all iterated localisations of H are ℓ -cycle-free and so H is ℓ -acyclic by Proposition 1.4.9.

So we are left to show that (ii) \iff (iii). In Lemma 1.5.12 below we show that a walk of length at least 3 on an apex acyclic hypergraph is contractible, if and only if it has a triangulation. So, if H is apex acyclic and ℓ -simply connected, then all closed walks of length at least 3 and at most ℓ have a triangulation which implies that H is ℓ -cycle-free by Lemma 1.4.7. This proves (ii) \implies (iii).

It is easy to see that a walk that has a triangulation is contractible. So, using Lemma 1.4.7 again we see that ℓ -cycle-freeness implies simple ℓ -connectivity. \square

An interesting corollary of Theorem 1.5.9 gives a topological characterisation of hypergraph acyclicity. We set $H_{\downarrow a}$ to be the *punctured localisation at a* . That is the hypergraph H_a with the vertex a removed, i.e., $H_{\downarrow a} = H_a|_{N(a)\setminus\{a\}}$. H' is an iterated punctured localisation of H if there is a sequence $H = H_0, \dots, H_n = H'$ where H_{i+1} is a punctured localisation of H_i .

Corollary 1.5.10. *A finite hypergraph is acyclic if, and only if, all connected components of all its iterated punctured localisations are simply connected.*

Proof. It is easy to prove that H_a is acyclic iff $H_{\downarrow a}$ is acyclic. So if H is acyclic then all its iterated punctured localisations are acyclic. Then by Theorem 1.5.9 all the connected components of its iterated punctured localisation are simply connected.

We prove the converse by induction over the number of vertices in H . The claim is trivial for a hypergraph with a single vertex. If all iterated punctured localisations of H are simply connected then the same is true for all $H_{\downarrow a}$, which by induction hypothesis are then acyclic. Thus, by Theorem 1.5.9, H is acyclic. \square

The finiteness condition in the aforementioned corollary cannot be dropped. Let H be the hypergraph on vertex set $V = \{a, b, c\} \cup \mathbb{N}$ and hyperedges $S = \{h_{x,n} \mid x \in \{a, b, c\}, n \in \mathbb{N}\}$ with $h_{x,n} = \{a, b, c\} \setminus \{x\} \cup \{i \in \mathbb{N} : i \leq n\}$. Then H is not acyclic but all of its iterated punctured localisations are simply connected.


The next lemma is a preparation for Lemma 1.5.12. We say that a triangulation T of $a_0 \dots a_n$ isolates $i \in \mathbb{Z}_n$ if $|\{s \in T \mid i \in s\}| = 1$.

Lemma 1.5.11. *Let H be an apex acyclic hypergraph and $\alpha = a_0 \dots a_n$ a closed walk that has a triangulation. If $\{a_{j-1}, a_j, a_{j+1}\}$ is guarded, then there is a triangulation of α that isolates j .*


Proof. We show that for every triangulation T of $a_0 \dots a_n$ with $|\{s \in T \mid j \in s\}| > 1$ there is another triangulation T' s.t. $|\{s \in T' \mid j \in s\}| < |\{s \in T \mid j \in s\}|$.

We can order the elements of $\{s \in T \mid j \in s\}$ into (s_1, \dots, s_m) , s.t. $j+1 \in s_1$, $j-1 \in s_m$ and $|s_i \cap s_{i+1}| = 2$, as in the sketch

1 Acyclicity and covers



Let j_0, \dots, j_{m+1} be an enumeration of the elements ‘traversed’ by these triangles, i.e., $j_0 = j$, $j_1 = j + 1$, $j_{m+1} = j - 1$ and j_i ($1 < i < m + 1$) is the single element in $s_i \cap s_{i+1} \setminus \{j\}$. We want to argue now that there is a hyperedge s^* that contains a_{j_0} and $a_{j_{k-1}}, a_{j_k}, a_{j_{k+1}}$ for a suitable $k \in \{1, \dots, m\}$. With such an s^* we could alter the triangulation of α as in the following sketch and obtain a suitable T' .



In order to show that such an s^* exists we introduce the notion of a clamp. A clamp is a hyperedge s that contains $a_{j_0}, a_{j_{k_1}}, a_{j_{k_2}}$ for some $1 \leq k_1 < k_2 \leq m + 1$. The length of a clamp is $k_2 - k_1$. Note that the guard in the assumption of the lemma is a clamp.

If we have a clamp of length 2 we obtain a suitable s^* : consider the sequence of hyperedges s, s_{k_1}, s_{k_2} . Since $H_{a_{j_0}}$ is acyclic there is an s^* containing $s \cap s_{k_1}, s_{k_1} \cap s_{k_2}, s_{k_2} \cap s$. In particular, $a_{j_0}, a_{j_{k_1}}, a_{j_{k_1+1}}$ and $a_{j_{k_1+2}} = a_{j_{k_2}}$ are all in s^* .

We finish by showing that if there is a clamp s of length greater than 2 then either we get an s^* as desired or a clamp of shorter length. For this look at the sequence $\alpha = s, s_{j_{k_1+1}}, \dots, s_{j_{k_2}}$. As $H_{a_{j_0}}$ is acyclic there is a witness s^* that this is not a hypertour. In nearly all possible cases, how this s^* lies over α it has the desired property. The only exceptions are $s_{j_{k_1+2}} \cap s_{j_{k_1+1}}, s_{j_{k_1+1}} \cap s, s \cap s_{j_{k_2}} \subseteq s^*$ or $s_{j_{k_2-1}} \cap s_{j_{k_2}}, s_{j_{k_2}} \cap s, s \cap s_{j_{k_1+1}} \subseteq s^*$, but in both cases s^* is then a shorter clamp. \square

Lemma 1.5.12. *A closed walk in an apex acyclic hypergraph is contractible if, and only if, it has a triangulation.*


Proof. That closed walks which have a triangulation are contractible is trivial.

For the converse we show that for closed walks of length at least 3


the existence of a triangulation is preserved by basic transformations. $(*)$

Let α and α' be two closed walks of length at least 3 where α' is the result of a fold/unfold of α and α has a triangulation.

First we treat the case in which α' is the result of an unfolding of α . In this case we can simply ‘add’ the guard s' of the unfolding to a triangulation of α to get a triangulation of α' , as in the following sketch:



Now assume that α' is the result of a fold of $\alpha = a_0 \dots a_n$ where $a_{j-1}a_ja_{j+1}$ gets folded to $a_{j-1}a_{j+1}$. Since α has a triangulation and $\{a_{j-1}, a_j, a_{j+1}\}$ is guarded, there is a triangulation T that isolates j (Lemma 1.5.11). This can then induce a triangulation of α' as in the following sketch:



□

1.6 Acyclic covers of hypergraphs

In this section we investigate the relationship between acyclicity and (un)branched covers in the context of hypergraphs. We focus on two results ‘there is a finite hypergraph that does not have finite, ℓ -simply connected unbranched covers for all degrees $\ell \in \mathbb{N}$ ’ (Proposition 1.6.10), and ‘every finite hypergraph has finite, ℓ -acyclic branched covers for all degrees $\ell \in \mathbb{N}$ ’ (Theorem 1.6.23). The first result is new and at its core lies the undecidability of the word problem for hypergraphs. The second result is from Otto [34]. These branched covers have been linked to ‘coset ℓ -acyclicity’ groupoids [35]. Our contribution here is that we show that coset ℓ -acyclicity is naturally connected to the hypergraph ℓ -acyclicity via the notion of hypertours.

We motivate briefly why these results are of interest. It can be shown (see the remark after Corollary 1.6.7) that every hypergraph property Q for which it is guaranteed that every hypergraph has an unbranched cover with property Q is weaker than simple connectivity. In other words, simple connectivity is the strongest

1 Acyclicity and covers

property that *can be achieved in unbranched covers*. So it is natural to ask whether we can approximate simply connected covers by finite covers.

We would like to have a similar relationship between hypergraph acyclicity and branched covers. However, we can easily find a hypergraph property that is incompatible with hypergraph acyclicity but also can be achieved in branched covers (see ‘branched covers and acyclicity’ in Section 1.6.2). So hypergraph acyclicity and branched covers are not as closely linked as simple connectivity and unbranched covers are. Still, of all the acyclicity notions appearing in the literature it is the strongest one which can be guaranteed for branched covers, at least to the best of my knowledge. So again it is natural to ask, can we approximate this by finite means?

We give the definition of unbranched/branched covers in the respective subsections. Both notions are based on hypergraph homomorphisms.

Definition 1.6.1. A *homomorphism* $\varphi: \widehat{H} \xrightarrow{\text{hom}} H$ between hypergraphs $\widehat{H} = (\widehat{V}, \widehat{E})$ and $H = (V, E)$ is a map $\varphi: \widehat{V} \rightarrow V$ s.t. φ restricted to \widehat{s} is injective and $\varphi(\widehat{s}) \in S$ for every $\widehat{s} \in \widehat{S}$.

Note that a hypergraph homomorphism between simple graphs is just a graph homomorphism. In particular, a homomorphism $\varphi: \widehat{H} \xrightarrow{\text{hom}} H$ between hypergraphs induces a homomorphism $\varphi: \widehat{G} \xrightarrow{\text{hom}} G$ between their Gaifman graphs.

1.6.1 Unbranched covers and granular covers

We reduce the theory of unbranched covers of hypergraphs to the theory of granular covers. In particular, we can translate (locally) simply connected unbranched covers of hypergraphs to (locally) free granular covers.

An unbranched cover of hypergraphs is a homomorphism $\varphi: \widehat{H} \xrightarrow{\text{hom}} H$ that is locally an isomorphism, i.e., φ induces an isomorphism $\widehat{H}_{\widehat{a}} \simeq H_a$ for all pairs $\widehat{a} \in \varphi^{-1}(a)$. We give an alternative definition.

Definition 1.6.2. An *unbranched cover* $\varphi: \widehat{H} \xrightarrow{\text{unb}} H$ is a homomorphism of connected hypergraphs $\widehat{H} = (\widehat{V}, \widehat{S})$ and $H = (V, S)$ with the *unique lifting property*: for all guarded $\widehat{A} \subseteq \widehat{V}$ and $s \in S$ with $\varphi(\widehat{A}) \subseteq s$ there is a unique $\widehat{s} \in \widehat{S}$ that contains \widehat{A} and maps to s .

Unbranched covers of hypergraphs are a generalisation of unbranched covers of simple graphs. Also, as for homomorphisms, every cover $\varphi: \widehat{H} \xrightarrow{\text{unb}} H$ of hypergraphs induces a cover $\varphi: \widehat{G} \xrightarrow{\text{unb}} G$ of their Gaifman graphs. Note that the cover $\varphi: \widehat{G} \xrightarrow{\text{unb}} G$ preserves every ‘triangle’ $abca \in \pi(G)$ for which $\{a, b, c\}$ is guarded in H . Thus $\varphi: \widehat{G} \xrightarrow{\text{unb}} G$ is an $\text{Tr}(H)$ -granular cover of G , where $\text{Tr}(H) := \{abca \in \pi_{\text{cl}}(G) \mid \{a, b, c\} \text{ is guarded in } H\}$ is the set of guarded ‘triangles’ in G (recall that $\pi_{\text{cl}}(G)$ is the set of all tours in G). We show that this connection between unbranched covers of H and $\text{Tr}(H)$ -granular covers of G is one-to-one.

We denote the class of unbranched covers of a hypergraph H by $\text{Cov}(H)$. Similarly, for a given graph G and $R \subseteq \pi_{\text{cl}}(G)$, we denote the class of R -granular covers of G by $\text{Cov}(G, R)$. We can give a bijection $F: \text{Cov}(G, \text{Tr}(H)) \rightarrow \text{Cov}(H)$ between the $\text{Tr}(H)$ -granular covers of G and the unbranched covers of H : for an $\text{Tr}(H)$ -granular cover $\varphi: \widehat{G} \xrightarrow{\text{unb}} G$ with $\widehat{G} = (\widehat{V}, \widehat{E})$ we let

$$F(\varphi: \widehat{G} \xrightarrow{\text{unb}} G) = \varphi: \widehat{H} \xrightarrow{\text{unb}} H, \text{ with} \\ \widehat{H} = (\widehat{V}, \widehat{S}) \text{ and } \widehat{S} = \{ \text{connected components of } \varphi^{-1}(s) \mid s \in S \}.$$

Lemma 1.6.3. *The map $F: \text{Cov}(G, \text{Tr}(H)) \rightarrow \text{Cov}(H)$ is a well-defined bijection.*

Proof. It is easy to see that the inverse of F is given by the function that maps every unbranched cover of H to the respective $\text{Tr}(H)$ -granular cover of the Gaifman graphs.

It remains to show that $F(\varphi: \widehat{G} \xrightarrow{\text{unb}} G) = (\varphi: \widehat{H} \xrightarrow{\text{unb}} H)$ is actually an unbranched cover.

First note that if \widehat{s} is a connected component of $\varphi^{-1}(s)$, then \widehat{s} is a clique: if $\widehat{a}, \widehat{b}, \widehat{c} \in \widehat{s}$ and $(\widehat{a}, \widehat{b}), (\widehat{b}, \widehat{c}) \in \widehat{E}$, then their images a, b, c are guarded by s and thus $abca \in \text{Tr}(H)$. The lift \widehat{a} of $abca$ to \widehat{a} is closed as $\varphi: \widehat{G} \xrightarrow{\text{unb}} G$ is a $\text{Tr}(H)$ -granular cover. So \widehat{a} has the form $\widehat{a}\widehat{b}\widehat{c}\widehat{a}$ and thus $(\widehat{a}, \widehat{c}) \in \widehat{E}$.

Using this we can show that φ is a hypergraph homomorphism. Clearly $\varphi(\widehat{s}) = s$ as for any $\widehat{a} \in \widehat{s}$ and $b \in s$ there is a preimage \widehat{b} of b s.t. $(\widehat{a}, \widehat{b}) \in \widehat{E}$. The map φ is also injective on \widehat{s} , since if $\varphi(\widehat{a}) = \varphi(\widehat{b})$ for $\widehat{a}, \widehat{b} \in \widehat{s}$ there cannot be an edge between \widehat{a} and \widehat{b} . Thus $\widehat{a} = \widehat{b}$ as they are in the same connected component.

The map φ also has the unique lifting property: for guarded $\widehat{A} \in \widehat{V}$ with $\varphi(\widehat{A}) \in s$, the connected component of \widehat{A} in $\varphi^{-1}(s)$ is the unique hyperedge that maps to s and contains \widehat{A} . \square

Simply connected covers

Lemma 1.5.2 says that the set of all contractible tours in H is given by $\text{ncl}_G(\text{Tr}(H))$ where G is the Gaifman graph. With this in mind, the next lemma tells us that a tour in the covering graph of an unbranched cover is contractible, if and only if its image is contractible. So unbranched covers ‘preserve and reflect’ contractibility.

Lemma 1.6.4. *Let $\varphi: \widehat{H} \xrightarrow{\text{unb}} H$ be an unbranched cover of hypergraphs and \widehat{G} and G their respective Gaifman graphs. Then*

$$\text{ncl}_{\widehat{G}}(\text{Tr}(\widehat{H})) = \varphi^{-1}(\text{ncl}_G(\text{Tr}(H))).$$

Proof. We let $\widehat{R} = \text{Tr}(\widehat{H})$ and $R = \text{Tr}(H)$. It is easy to show that $\widehat{R} = \varphi^{-1}(R)$. For example, if $\varphi(\widehat{a}\widehat{b}\widehat{c}\widehat{d}) \in R$, then for $\varphi(\widehat{a}\widehat{b}\widehat{c}\widehat{d}) = abca$ there is a guard $s \in S$. Thus $\widehat{a}, \widehat{b}, \widehat{c}, \widehat{d}$ are all in the same connected component of $\varphi^{-1}(s)$ and so $\widehat{a} = \widehat{d}$ and $\{\widehat{a}, \widehat{b}, \widehat{c}\}$ is guarded.

1 Acyclicity and covers

Note that $\varphi^{-1}(\text{ncl}_G(R))$ is closed under products and conjugation. Thus $\text{ncl}_{\widehat{G}}(\widehat{R}) \subseteq \varphi^{-1}(\text{ncl}_G(R))$.

For the converse let $\widehat{\alpha} \in \varphi^{-1}(\text{ncl}_G(R))$. Then $\varphi(\widehat{\alpha})$ has the form $\alpha_1 r_1 \alpha_1^{-1} \dots \alpha_n r_n \alpha_n^{-1}$ with $r_1, \dots, r_n \in R$. So $\widehat{\alpha} = \widehat{\alpha}_1 \widehat{r}_1 \widehat{\alpha}_1^{-1} \dots \widehat{\alpha}_n \widehat{r}_n \widehat{\alpha}_n^{-1}$ with $\widehat{r}_i \in \varphi^{-1}(R) = \widehat{R}$. Thus $\widehat{\alpha} \in \text{ncl}_{\widehat{G}}(\widehat{R})$. \square

We can now show that an unbranched cover is simply connected, if and only if the cover at the level of the Gaifman graphs is a free $\text{Tr}(H)$ -granular cover.

Lemma 1.6.5. *Let H be a hypergraph and G its Gaifman graph. Then $\psi: \widetilde{G} \xrightarrow{\text{unb}} G$ is a free $\text{Tr}(H)$ -granular cover if, and only if, $F(\psi: \widetilde{G} \xrightarrow{\text{unb}} G)$ is simply connected.*

Proof. Let $\psi: \widetilde{H} \xrightarrow{\text{unb}} H$ be the image of $\psi: \widetilde{G} \xrightarrow{\text{unb}} G$ under F . Then

$$\begin{aligned} \widetilde{H} \text{ is simply connected} &\stackrel{1.5.3}{\iff} \text{ncl}_{\widetilde{G}}(\text{Tr}(\widetilde{H})) = \pi_{\text{cl}}(\widetilde{G}) \\ &\stackrel{1.6.4}{\iff} \psi^{-1}(\text{ncl}_G(\text{Tr}(H))) = \pi_{\text{cl}}(\widetilde{G}) \\ &\stackrel{1.3.11}{\iff} \psi: \widetilde{G} \xrightarrow{\text{unb}} G \text{ is the free } \text{Tr}(H)\text{-granular cover. } \square \end{aligned}$$

We can use this lemma to show that simply connected covers exist and are universal.

Definition 1.6.6. A homomorphism from $\psi: \widetilde{H} \xrightarrow{\text{unb}} H$ to $\varphi: \widehat{H} \xrightarrow{\text{unb}} H$ is a homomorphism $f: \widetilde{H} \xrightarrow{\text{hom}} \widehat{H}$ s.t. $\psi = \varphi \circ f$.

Similar to how homomorphisms and unbranched covers of hypergraphs induce homomorphisms and unbranched covers of Gaifman graphs, homomorphisms of unbranched hypergraph covers induce homomorphisms of graph covers.

The following is a direct corollary of the previous lemma and Proposition 1.3.11.

Corollary 1.6.7. *Every hypergraph has a simply connected cover and it is universal.*

Using Theorem 1.5.9 we get:

Corollary 1.6.8. *Every apex acyclic hypergraph has an acyclic unbranched cover.*

Corollary 1.6.7 tells us that every unbranched cover of a simply connected hypergraph is trivial. This shows us that simple connectivity is the strongest hypergraph property that can be guaranteed for unbranched covers: if every hypergraph has an unbranched cover with property Q , then a simply connected hypergraph has property Q as it only allows for the trivial unbranched cover.

Locally simply connected covers

We show that finite approximations of simply connected covers do not generally exist.

We start by giving a localised version of Lemma 1.6.5.

Lemma 1.6.9. *Let H be a hypergraph and G its Gaifman graph. Then a $\text{Tr}(H)$ -granular cover $\varphi: \widehat{G} \xrightarrow{\text{unb}} G$ is ℓ -free if, and only if, $F(\varphi: \widehat{G} \xrightarrow{\text{unb}} G)$ is ℓ -simply connected.*

Proof. Let $F(\varphi: \widehat{G} \xrightarrow{\text{unb}} G) = \varphi: \widehat{H} \xrightarrow{\text{unb}} H$ with $\widehat{G} = (\widehat{V}, \widehat{E})$. Then

$$\begin{aligned} \widehat{H} \text{ is } \ell\text{-simply connected} &\stackrel{1.5.5}{\iff} \text{ncl}_{\widehat{G}}(\text{Tr}(\widehat{H})) \cap \widehat{V}^{\leq 2\ell+1} = \pi_{\text{cl}}(\widehat{G}) \cap \widehat{V}^{\leq 2\ell+1} \\ &\stackrel{1.6.4}{\iff} \psi^{-1}(\text{ncl}_G(\text{Tr}(H))) \cap \widehat{V}^{\leq 2\ell+1} = \pi_{\text{cl}}(\widehat{G}) \cap \widehat{V}^{\leq 2\ell+1} \\ &\stackrel{1.3.12}{\iff} \psi: \widehat{G} \xrightarrow{\text{unb}} G \text{ is } \ell\text{-free.} \quad \square \end{aligned}$$

Similarly to simply connected covers, we get the corollary that the ℓ -simply connected unbranched covers are exactly the ℓ -universal unbranched covers, i.e., their ℓ -neighbourhoods are isomorphic to the ℓ -neighbourhoods of the simply connected cover. So finite, ℓ -simply connected unbranched covers indeed approximate simply connected covers. However, they might not exist.

Proposition 1.6.10. *There is a finite hypergraph which does not have finite, ℓ -locally simply connected covers for all degrees $\ell \in \mathbb{N}$.*

Proof. By Theorem 1.5.7, there is a hypergraph H that has an undecidable word problem $\{\alpha \in \pi_{\text{cl}}(H) \mid \alpha \text{ is contractible}\}$. By Corollary 1.5.3 this set is given by $\text{ncl}_G(\text{Tr}(H))$, where G is the Gaifman graph of H . By Lemma 1.3.16, G does not have finite, ℓ -free $\text{Tr}(H)$ -granular covers for all degrees $\ell \in \mathbb{N}$. Thus, by the previous lemma, H does not have finite, ℓ -simply connected covers for all degrees $\ell \in \mathbb{N}$. \square

We have shown that for computability reasons, arbitrarily good, finite approximations of simply connected unbranched covers do not generally exist. It would be interesting to know whether there are also other reasons that forbid such approximations, i.e., are there finite hypergraphs that have decidable word problems but do not have finite, ℓ -simply connected covers for all degrees $\ell \in \mathbb{N}$? We do not know the answer to this question.

This question is particularly interesting for the class of apex acyclic hypergraphs. Note that the covering graph \widehat{H} of an unbranched cover $\varphi: \widehat{H} \xrightarrow{\text{unb}} H$ is apex acyclic, if and only if the base graph H is apex acyclic, since $\widehat{H}_{\widehat{a}} \simeq H_a$ for every pair $\widehat{a} \in \varphi^{-1}(a)$. So, by Theorem 1.5.9, it is equivalent to ask: does every finite, apex acyclic hypergraph have finite, ℓ -acyclic unbranched covers for all degrees of $\ell \in \mathbb{N}$? We know that apex acyclic hypergraphs have acyclic unbranched covers. We also know that a non-apex acyclic hypergraph does not have ℓ -acyclic unbranched covers for all degrees $\ell \in \mathbb{N}$, as an ℓ -cyclic neighbourhood in the base graph induces an ℓ -cyclic neighbourhood in all covering graphs. A positive answer would then render the class of apex acyclic hypergraph as the class of hypergraphs for which acyclic unbranched covers can be approximated by finite covers.

Note that we cannot adapt the argument of Proposition 1.6.10 to prove a negative answer for this problem. The word problem for apex acyclic hypergraphs is decidable: by Lemma 1.5.12, a walk is contractible, if and only if it has a triangulation, which is obviously decidable.


1.6.2 Branched covers and coset acyclic groupoids

In [34] Otto shows that that finite hypergraphs have finite, ℓ -acyclic branched covers of arbitrary degree $\ell \in \mathbb{N}$. In [35] there is an alternative but ultimately erroneous proof of this theorem. Nevertheless this different approach seems to be more natural and better tailored to the problem itself. The crucial ingredient for this approach are finite groupoids that are ‘coset ℓ -acyclic’ for every $\ell \in \mathbb{N}$. However, we do not know at this point whether these objects exists. The part in [35] where these groupoids are used in a reduced product with a hypergraph to produce ℓ -acyclic hypergraphs is correct. As noted before, the main problem is to provide these groupoids. However, the second step is interesting in itself, as the connection between the acyclicity of the groupoids and the hypergraphs is very natural. This connection is not very developed in [35]. Our contribution is that we show how coset ℓ -acyclicity can be motivated as a natural notion for groupoids and link it to the notion of hypertours.

Branched covers and hypergraph acyclicity

Definition 1.6.11. A branched cover $\varphi: \widehat{H} \xrightarrow{\text{unb}} H$ is a homomorphism of connected hypergraphs $\widehat{H} = (\widehat{V}, \widehat{S})$ and $H = (V, S)$ with the *lifting property*: for all guarded $\widehat{A} \subseteq \widehat{V}$ and $s \in S$ with $\varphi(\widehat{A}) \subseteq s$ there is an $\widehat{s} \in \widehat{S}$ that contains \widehat{A} and maps to s .


Below in Lemma 1.6.18 we show that every hypergraph has an acyclic branched cover. However, we also can show that every non-trivial hypergraph has a cyclic branched cover, even if the base hypergraph is acyclic (non-trivial means that the hypergraph has not just singleton hyperedges): every non-trivial hypergraph H has a 3-cyclic branched cover. Simply glue three copies H', H'', H''' of H together at vertices a, b that are guarded by some s so that its copies s', s'', s''' form a hypertriangle like this:



With this we can see that the following property Q : ‘the hypergraph is trivial or cyclic’ can be guaranteed for branched covers. Thus hypergraph acyclicity is not the strongest such property. However, Q certainly is not an acyclicity notion for hypergraphs, and all acyclicity notions that appear in the literature are either weaker than hypergraph acyclicity or cannot be guaranteed for branched covers.

This begs the question whether there is a hypergraph property that is an acyclicity notion, that can be guaranteed by branched covers, and that is stronger than or incomparable to hypergraph acyclicity? To answer this, we must first have a definition of what an acyclicity notion for hypergraphs should be. However, there is little

research into this topic. One attempt to formalize this is from Brault-Baron [10] who introduces the term ‘good acyclicity notion’. It is shown that any ‘good acyclicity notion’ is either at most as strong as γ -acyclicity and at least as strong as β -acyclicity, or at most as strong as α -acyclicity (what we call here hypergraph acyclicity) and at least as strong as cycle-freeness (for the definitions of β - and γ -acyclicity we refer to the source). So any ‘good acyclicity notion’ lies somewhere in the gray areas of the following sketch:



It is a fact that the guarded cartwheel does not have β -acyclic branched covers. Thus hypergraph acyclicity is the strongest ‘good acyclicity notion’ that can be guaranteed for branched covers. Note that simple connectivity is not a ‘good acyclicity notion’ as it is weaker than cycle-freeness (the ‘cartwheel’ in fig. 1.3 is not cycle-free but simply connected).

Constructing branched covers by glueing

The basic idea of how to provide acyclic branched covers of hypergraphs $H = (V, S)$ is to construct these by glueing together copies of the hyperedges of H . To know how to glue these copies together we need a ‘glueing schema’ (in [35] these are called ‘overlap patterns’). Basically, a glueing schema is a simple graph $\hat{I} = (\hat{S}, \hat{E})$ with a labelling $\varphi: \hat{S} \rightarrow S$. To obtain from a glueing schema a hypergraph, we take for each $\hat{s} \in \hat{S}$ a copy of $\varphi(\hat{s})$ and for $(\hat{s}, \hat{t}) \in \hat{E}$ we glue these copies together as specified by the intersection of $\varphi(\hat{s})$ and $\varphi(\hat{t})$.

We now define glueing schemas and the result of the glueing process formally.

Definition 1.6.12. The *intersection graph* $I(H)$ of a hypergraph $H = (V, S)$ is the simple graph (S, E) with $E = \{(s, s') \mid s \cap s' \neq \emptyset\}$.

A *glueing schema* for H is simply an unbranched cover $\varphi: \hat{I} \xrightarrow{\text{unb}} I$ of the intersection graph of H . Usually we do not mention the covering map φ explicitly and just write \hat{I} for the whole glueing schema, where φ is understood implicitly (this notation resonates well with the view of \hat{I} as a labelled graph).

Let $\hat{I} = (\hat{S}, \hat{E})$ be a glueing schema for $H = (V, S)$. We set $H \otimes \hat{I}$ to be the following hypergraph $\hat{H} = (\hat{V}, [\hat{S}])$:

$$(i) \quad \hat{V} := \left(\bigcup_{\hat{s} \in \hat{S}} \varphi(\hat{s}) \times \{\hat{s}\} \right) / \approx, \text{ where}$$

$$(a, \hat{s}) \approx (a, \hat{t}) \iff \text{there is a walk } \hat{s}_0 \dots \hat{s}_\ell \text{ from } \hat{s} \text{ to } \hat{t} \text{ s.t. } a \in \bigcap_{i=0}^{\ell} \varphi(\hat{a}_i).$$

1 Acyclicity and covers

- (ii) $[\widehat{S}] := \{[\widehat{s}] \mid \widehat{s} \in \widehat{S}\}$ where $[\widehat{s}] = \{[a, \widehat{s}] \mid a \in \varphi(\widehat{s})\}$ and $[a, \widehat{s}]$ stands for the equivalence class of (a, \widehat{s}) .


The following lemma is straightforward to prove.

Lemma 1.6.13. *Let \widehat{I} be a glueing schema for the hypergraph H . Then $\pi: H \otimes \widehat{I} \xrightarrow{\text{bra}} H$ is a branched cover where the covering map is the projection $\pi: \widehat{V} \rightarrow V; [a, \widehat{s}] \mapsto a$.*


A notion of acyclicity for glueing schemas

We have seen that we can obtain branched covers of hypergraphs by using glueing schemas. Now we describe degrees of acyclicity of glueing schemas which guarantee acyclic hypergraphs for the glueing construction.

We start by considering two examples of glueing schemas for the tetrahedron H . The tetrahedron is the hypergraph $H = (V, S) = (\{a_0, a_1, a_2, a_3\}, \{s_0, s_1, s_2, s_3\})$ with $s_i = \{a_j \mid j \neq i\}$ and its intersection graph is given by $I = (S, E)$ with $E = \{(s_i, s_j) \mid i \neq j\}$:




The first glueing schema is \widehat{I}_1 (the labels of the nodes indicate the map $\varphi: \widehat{I}_1 \xrightarrow{\text{unb}} I$):




The resulting hypergraph $\widehat{H}_1 = H \otimes \widehat{I}_1$ is just the tetrahedron again. For example, we have $[\widehat{s}] = [\widehat{t}]$ for the two nodes \widehat{s} and \widehat{t} with label s_0 : for each $a \in s_0$ there is a path from \widehat{s} to \widehat{t} whose labels all contain a and thus $[a, \widehat{s}] = [a, \widehat{t}]$.

Note that even though \widehat{I}_1 is a 3-acyclic graph the hypergraph \widehat{H}_1 has a 3-cycle. Even worse, for every $\ell \in \mathbb{N}$ there is a glueing schema that is ℓ -acyclic as a graph but it produces just the tetrahedron. This shows us that ordinary graph acyclicity is too weak for glueing schemas.

Now consider the following glueing schema \widehat{I}_2 for H (the outer edges ‘wrap around’ as indicated by the decorations):



Using \widehat{I}_2 we obtain a 3-acyclic hypergraph $\widehat{H}_2 = H \otimes \widehat{I}_2$ (the vertices are identified as indicated by the symbols):




Motivated by these two examples, we introduce the notion of ℓ -acyclicity for glueing schemas.

We start by introducing 2-acyclicity which holds a special status among the different levels of acyclicity as it does not correspond to an acyclicity property of the result of the glueing construction but rather controls how hyperedges in the glueing construction can intersect.

Let $\widehat{I} = (\widehat{S}, \widehat{E})$ be a glueing schema for the hypergraph $H = (V, S)$. Let \widehat{s}, \widehat{t} be vertices of \widehat{I} . We say that \widehat{s} and \widehat{t} intersect over $a \in V$ if $a \in \varphi([\widehat{s}] \cap [\widehat{t}])$, i.e., $[a, \widehat{s}] = [a, \widehat{t}]$. By definition, \widehat{s} and \widehat{t} intersect over a if there is a witnessing walk, a walk from \widehat{s} to \widehat{t} whose labels all contain a . We say that a walk $\widehat{\alpha}$ describes the intersection of \widehat{s} and \widehat{t} if $\widehat{\alpha}$ witnesses that \widehat{s} and \widehat{t} intersect over a for all $a \in \varphi([\widehat{s}] \cap [\widehat{t}])$. In \widehat{I}_1 not all pairs of vertices have walks that describe their intersection. On the other hand, in \widehat{I}_2 each pair of vertices has a walk that describes their intersection. Note that these walks do not have to be unique. We can see this on \widehat{I}_2 :

1 Acyclicity and covers



Even though walks that describe intersections do not have to be unique, their set of labels are, at least in the case of the tetrahedron. If $\hat{s}_0 \dots \hat{s}_n$ and $\hat{s}'_0 \dots \hat{s}'_m$ both describe the intersection of \hat{s} and \hat{t} then $\{\varphi(\hat{s}_i) \mid i \leq n\} = \{\varphi(\hat{s}'_i) \mid i \leq m\}$ as both walks have to ‘lose’ exactly the same vertices and every label loses exactly one particular vertex. This motivates the following definition.

Definition 1.6.14. A glueing schema $\hat{I} = (\hat{S}, \hat{E})$ for $H = (V, S)$ is *2-acyclic* if for all vertices $\hat{s}, \hat{t} \in \hat{S}$ there is a unique minimal set of labels $S(\hat{s}, \hat{t}) \subseteq S$ s.t. there is a walk $\hat{s}_0 \dots \hat{s}_n$ from \hat{s} to \hat{t} whose labels are all in $S(\hat{s}, \hat{t})$, i.e., $\varphi(\hat{s}_i) \in S(\hat{s}, \hat{t})$ for all $i = 0, \dots, n$.

2-acyclicity guarantees that each pair of vertices has walks that describe their intersections. In the case of the tetrahedron the converse is true as well, a glueing schema \hat{I} for the tetrahedron is 2-acyclic if it has for each pair of vertices walks that describe their intersections (this is also true for ‘higher dimensional simplices’, i.e., hypergraphs with n vertices that have for each vertex a hyperedge that contains all but this vertex).


2-acyclicity gives us a useful criterion for when $[\hat{s}_0] \cap [\hat{s}_1] \subseteq [\hat{s}]$. For $S' \subseteq S$ we set $\hat{I}(S')$ to be the induced subgraph on those vertices $V' \subseteq V$ whose labels are in S' , i.e., $\hat{I}(S') := \hat{I}|_{\varphi^{-1}(S')}$. The following lemma is immediate.

Lemma 1.6.15. *Let \hat{I} be a 2-acyclic glueing schema for H . Then $[\hat{s}_0] \cap [\hat{s}_1] \subseteq [\hat{s}]$ if \hat{s} is in the connected component of \hat{s}_0 and \hat{s}_1 in $\hat{I}(S(\hat{s}_0, \hat{s}_1))$.*

In the case of the tetrahedron and its higher dimensional generalisations the converse is also true.

We now motivate the notion of 3-acyclicity for glueing schemas. We want to ensure that $\hat{H} = H \otimes \hat{I}$ has no hypertriangles, i.e., that for all $\hat{s}_0, \hat{s}_1, \hat{s}_2 \in \hat{I}$ there is an $\hat{s} \in \hat{I}$ s.t. $[\hat{s}_0] \cap [\hat{s}_1], [\hat{s}_1] \cap [\hat{s}_2], [\hat{s}_2] \cap [\hat{s}_0] \subseteq [\hat{s}]$. If \hat{I} is 2-acyclic we can rewrite this condition to $\hat{s} \in C_0 \cap C_1 \cap C_3$ where the C_i are the connected components of the \hat{s}_i in $\hat{I}(S(\hat{s}_i, \hat{s}_{i+1}))$. We take this as the definition of 3-acyclicity: \hat{I} is 3-acyclic if it is 2-acyclic and for all vertices $\hat{s}_0, \hat{s}_1, \hat{s}_2$ of \hat{I} we have $C_0 \cap C_1 \cap C_2 \neq \emptyset$ where C_i is the connected component of \hat{s}_i in $\hat{I}(S(\hat{s}_i, \hat{s}_{i+1}))$.

One can check that \hat{I}_2 is 3-acyclic, i.e., for all vertices $\hat{s}_0, \hat{s}_1, \hat{s}_2$ of \hat{I}_2 this ‘non-empty intersection’ property is true. Consider for example the following three vertices:



We can generalise this idea to higher degrees of acyclicity by translating in the same way the condition that $H \otimes I$ has no hypertours of short length.

Definition 1.6.16. A glueing schema $\hat{I} = (\hat{S}, \hat{E})$ for $H = (V, S)$ is ℓ -acyclic if it is 2-acyclic and for all sequences of vertices $(\hat{s}_i)_{i \in \mathbb{Z}_n}$ of length at most ℓ there is a $j \in \mathbb{Z}_n$ s.t. $C_{j-1} \cap C_j \cap C_{j+1} \neq \emptyset$ where C_i is the connected component of \hat{s}_i in $\hat{I}(S(\hat{s}_i, \hat{s}_{i+1}))$.

The notion of ℓ -acyclic glueing schema is derived from the notion of ‘component ℓ -acyclic’ groupoids, introduced in [35]. In the next section we discuss how these two terms relate and that the existence of the latter implies the existence of the latter.

The discussion so far makes the following connection between ℓ -acyclic glueing schemas and ℓ -acyclic branched covers of hypergraphs obvious.

Lemma 1.6.17. *If \hat{I} is an ℓ -acyclic glueing schema for H , then $H \otimes \hat{I}$ is ℓ -acyclic.*

Proof. We show that $H \otimes \hat{I}$ has no hypertours of length at most ℓ . Let $([\hat{s}_i])_{i \in \mathbb{Z}_n}$ be a sequence of hyperedges in $H \otimes \hat{I}$ of length at most ℓ . Let C_i be the connected component of \hat{s}_i in $\hat{I}(S(\hat{s}_i, \hat{s}_{i+1}))$. By ℓ -acyclicity of \hat{I} there are \hat{s} and $j \in \mathbb{Z}_n$ s.t. $\hat{s} \in C_{j-1} \cap C_j \cap C_{j+1}$. Then by Lemma 1.6.15, $[\hat{s}_{j-1}] \cap [\hat{s}_j], [\hat{s}_j] \cap [\hat{s}_{j+1}], [\hat{s}_{j+1}] \cap [\hat{s}_{j+2}] \subseteq [\hat{s}]$. \square

Note that the definition of an ℓ -acyclic glueing schema for H only depends on the intersection graph of I . So, it would be more fitting to speak of an ℓ -acyclic glueing schema for I . We can rephrase the statement of the previous lemma to ‘if \hat{I} is an ℓ -acyclic glueing schema for I , then $H \otimes I$ is an ℓ -acyclic hypergraph for every hypergraph with intersection graph I ’. Actually, it can be shown that, provided that \hat{I} is 2-acyclic, the converse is also true, i.e., that \hat{I} is ℓ -acyclic for I , if and only if, $H \otimes I$ is ℓ -acyclic for all hypergraphs with intersection graph I . So the notion of ℓ -acyclic glueing schema is optimal for our purposes.

We can now show fairly easily that hypergraphs have acyclic branched covers.

Lemma 1.6.18. *Every hypergraph has an acyclic branched cover.*

1 Acyclicity and covers

Proof. Let $\varphi: \tilde{I} \xrightarrow{\text{unb}} I$ be the free cover I . We have to show that \tilde{I} is an ℓ -acyclic glueing schema for all $\ell \in \mathbb{N}$.

Since \tilde{I} is acyclic there is for each pair of vertices \tilde{s}_0 and \tilde{s}_1 a unique connecting trail. The labels appearing in this trail are $S(\tilde{s}_0, \tilde{s}_1)$. Thus \tilde{I} is 2-acyclic.

Let $(\tilde{s}_i)_{i \in \mathbb{Z}_\ell}$ be a sequence of vertices and C_i the connected component of s_i in $\tilde{I}(S(\tilde{s}_i, \tilde{s}_{i+1}))$. Let $\tilde{\alpha}_i$ be the unique walk from \tilde{s}_i to \tilde{s}_{i+1} and α_i its projection. Then $\tilde{\alpha}_0 \dots \tilde{\alpha}_\ell$ is closed walk in \tilde{I} and thus its image $\alpha_0 \dots \alpha_\ell$ reduces to a trivial walk. By Lemma 0.2.5, there is a j such that α_j can be decomposed into $\alpha_j = \alpha_- \alpha_+$ such that α_-^{-1} is a suffix of α_{j-1} and α_+^{-1} is a prefix of α_{j+1} . Let $\tilde{s} = t(\tilde{\alpha}_-)$ where $\tilde{\alpha}_-$ is the lift of α_- to \tilde{s}_j . Then $\tilde{s} \in C_{j-1}$ as the walk $\alpha_{j-1} \alpha_-$ from \tilde{s}_{j-1} to \tilde{s} only uses labels that also appear in α_{j-1} . Similarly $\tilde{s} \in C_j$ and $\tilde{s} \in C_{j+1}$. \square

From glueing schemas to groupoids

The previous considerations reduce the problem of finding finite, ℓ -acyclic covers of finite hypergraphs to the question of finding finite, ℓ -acyclic glueing schemas. This problem was seemingly solved by Otto in [35] by showing the existence of finite ‘component ℓ -acyclic’ hypergraphs. However, it turns out that the presented proof is erroneous (and there is no know fix to this point). We thus state the existence of finite ‘component acyclic’ groupoids as conjecture below.

We discuss firsts how the original notion of ‘component acyclicity’ from [35] relates to acyclic glueing schemas. A glueing schema for $I = (V, E)$ is an unbranched cover $\varphi: \hat{I} \xrightarrow{\text{unb}} I$. There is a tight connection between these covers and groupoids that have generators ‘indexed’ by I .

An I -groupoid \mathbb{G} is a groupoid that has a famliy $(g_e)_{e \in E}$ of generators s.t. $s(g_e) = s(e)$, $t(g_e) = t(e)$ and $g_e^{-1} = g_{e^{-1}}$. Note that the vertices of I constitute the objects of \mathbb{G} . We can ‘evaluate’ walks $\alpha = e_1 \dots e_n$ in I over \mathbb{G} via $[\alpha]_{\mathbb{G}} = g_{e_1} \dots g_{e_n}$. For a subset $E' \subseteq E$ that is closed under $(\cdot)^{-1}$ we let $\mathbb{G}(E')$ the subgroupoid of \mathbb{G} generated by $\{g_e \mid e \in E'\}$. This subgroupoid is also given by $\mathbb{G}(E') = \{[\alpha]_{\mathbb{G}} \mid \alpha \text{ is a walk in } I \text{ that uses only edges in } E'\}$.

The aforementioned correspondence between unbranched covers of I and I -groupoids manifests itself as follows: The Cayley graph of an I -groupoid can be seen as an unbranched cover $\varphi: \hat{I} \xrightarrow{\text{unb}} I$ which is regular, i.e., its group of automorphisms acts transitively on each $\varphi^{-1}(a)$, and conversely every regular unbranched cover can be seen as the Cayley graph of an I -groupoid (we gloss over some minor details about connected components here).

We can translate the acyclicity notions developed for glueing schemas one-to-one to hypergraphs. We do so below but with a slight change. In the definition of ℓ -acyclicity we state the acyclicity conditions in a stronger manner that requires the existence of walks in certain weak subgraphs and not as before in induced subgraphs. This corresponds to a shift in view from vertex-coloured graphs to edge-coloured graphs. Below, after Definition 1.6.20 we discuss this shift on the example of \hat{I}_2 .

Apart from this change, the translation can basically be performed by the following

substitutions: $\widehat{I} \rightsquigarrow \mathbb{G}$, $S' \subseteq S \rightsquigarrow E' \subseteq E$, $\widehat{I}(S') \rightsquigarrow \mathbb{G}(E')$ and “connected component of s in $I(S')$ ” $\rightsquigarrow g\mathbb{G}(E')$. Here $\mathbb{G}(E')$ stands for the subgroupoid generated by $\{g_e \mid e \in E'\}$.

We give now the corresponding definitions and lemmas. We omit the proofs, as they are basically the same as above.

Let $H = (V, S)$ be a hypergraph with intersection graph $I = (V, E)$. For $a \in V$ we set $E(a) = \{e \in E \mid a \in s(e) \cap t(e)\}$. We define the *reduced product* $H \otimes \mathbb{G}$ of H with an I -groupoid \mathbb{G} as the following hypergraph $\widehat{H} = (\widehat{V}, \widehat{S})$:

- (i) $\widehat{V} := (\bigcup_{g \in \mathbb{G}} t(g) \times \{g\}) / \approx$, where $(a, g) \approx (a, g') \iff g^{-1}g' \in \mathbb{G}(E(a))$
- (ii) $\widehat{S} := \{[g] : g \in \mathbb{G}\}$, where $[g] := \{[a, g] : a \in s(g)\} \subseteq \widehat{V}$ and $[a, g]$ is the equivalence class of (a, g) .

Lemma 1.6.19. *Let H be a hypergraph with intersection graph I and \mathbb{G} an I -groupoid. Then $H \otimes \mathbb{G}$ induces a branched cover of H via $\pi: \widehat{V} \rightarrow V; [a, g] \mapsto a$.*

Definition 1.6.20. An I -groupoid \mathbb{G} is

- (i) *component 2-acyclic* if for all $g_0, g_1 \in \mathbb{G}$ with the same source there is a unique minimal set $E(g_0, g_1) \subseteq E$ that is closed under $(\cdot)^{-1}$ s.t. $g_0^{-1}g_1 \in \mathbb{G}(E(g_0, g_1))$.
- (ii) *component ℓ -acyclic* if it is component 2-acyclic and if for each sequence $(g_i)_{i \in \mathbb{Z}_n} \subseteq \mathbb{G}$ of length at most ℓ , all with the same source, there is a j s.t.

$$g_{j-1}\mathbb{G}(E(g_{j-1}, g_j)) \cap g_j\mathbb{G}(E(g_j, g_{j+1})) \cap g_{j+1}\mathbb{G}(E(g_{j+1}, g_{j+2})) \neq \emptyset.$$

As mentioned before, component ℓ -acyclicity is a stronger notion than the notion of ℓ -acyclicity that we introduced for glueing schemas. We compare both notions in the case $\ell = 2$ for a (regular) glueing schema \widehat{I} of $I = (V, E)$. 2-acyclicity requires that for all $\widehat{s}, \widehat{t} \in \widehat{V}$ there is a minimal vertex set $S' \subseteq S$ such that there is a walk from \widehat{s} to \widehat{t} that only uses vertices whose labels are in S' . Component 2-acyclicity requires that for all $\widehat{s}, \widehat{t} \in \widehat{V}$ there is a minimal edge set $E' \subseteq E$ such that there is a walk from \widehat{s} to \widehat{t} that only uses edges whose labels are in E' .

For example the glueing structure \widehat{I}_2 described above is not component 2-acyclic. Considering two neighbouring vertices $\widehat{s}_0, \widehat{s}_1$, say with label s_0 and s_1 , we get that $S(\widehat{s}_0, \widehat{s}_1) = \{s_0, s_1\}$ but $E(\widehat{s}_0, \widehat{s}_1)$ does not exist, as we can find two walks from \widehat{s}_0 to \widehat{s}_1 where for the first one the edge labels are in $\{(s_0, s_1), (s_1, s_0)\}$ and for the second one the edge labels are in $S \times S \setminus \{(s_0, s_1), (s_1, s_0)\}$.

Lemma 1.6.21. *Let \mathbb{H} be a hypergraph with intersection graph I and \mathbb{G} an I -groupoid.*

- (i) *If \mathbb{G} is component 2-acyclic then*

$$g \in \mathbb{G}(E(g_0, g_1)) \implies [g_0] \cap [g_1] \subseteq [g] \text{ in the product } H \otimes I.$$

1 Acyclicity and covers

(ii) $H \otimes \mathbb{G}$ is an ℓ -acyclic hypergraph if \mathbb{G} is component ℓ -acyclic.

This lemma shows how finite, component ℓ -acyclic groupoids can be used to construct finite, ℓ -acyclic branched covers of hypergraphs. However, the status of the existence of such groupoids is not clear. Originally a proof of this statement was published by Otto in [35] but it ultimately turned out to be erroneous. Still, this statement plays a central role in this thesis and we call it Otto's conjecture.⁵

Conjecture 1.6.22 (Otto). *For every finite multidigraph I and $\ell \in \mathbb{N}$ there are finite, component ℓ -acyclic I -groupoids.*

The original statement of Otto is stated for 'coset ℓ -acyclic' I -groupoids. Below we give the definition and show that this notion is equivalent to component ℓ -acyclicity.

As a corollary of this conjecture we would get the existence of finite, ℓ -acyclic branched covers of finite hypergraphs via the following argument: let H be a hypergraph with intersection graph I . Then if Conjecture 1.6.22 were true there were a finite, component ℓ -acyclic I -groupoid. By Lemma 1.6.21 $\widehat{H} = H \otimes \mathbb{G}$ is ℓ -acyclic. By Lemma 1.6.19 this induces an ℓ -acyclic branched cover.

There is an alternative in [34] of the existence of finite, ℓ -acyclic covers that does not rely on the existence of ℓ -acyclic groupoids.

Theorem 1.6.23 (Otto [34]). *Every finite hypergraph has a finite ℓ -acyclic cover.*

We now introduce Otto's notion of coset ℓ -acyclicity and compare it to component ℓ -acyclicity.

Definition 1.6.24. An I -groupoid is coset ℓ -acyclic if for all cyclic indexed sequences $((g_i, E_i))_{i \in \mathbb{Z}_n}$ of length at most ℓ

$$g_i^{-1}g_{i+1} \in \mathbb{G}(E_i) \text{ for all } i \in \mathbb{Z}_n \implies \\ g_j \mathbb{G}(E_{i-1} \cap E_j) \cap g_{j+1} \mathbb{G}(E_j \cap E_{j+1}) \neq \emptyset \text{ for some } j \in \mathbb{Z}_n$$

Lemma 1.6.25. *Let \mathbb{G} be an I -groupoid with finite $I = (V, E)$. Then \mathbb{G} is ℓ -coset acyclic if, and only if, \mathbb{G} is ℓ -component acyclic.*

Proof. We treat $\ell = 2$ separately. We show that the following are equivalent:

- (i) \mathbb{G} is component 2-acyclic
- (ii) \mathbb{G} is coset 2-acyclic,
- (iii) $\mathbb{G}(E_0) \cap \mathbb{G}(E_1) = \mathbb{G}(E_0 \cap E_1)$ for subsets $E_0, E_1 \subseteq E$ that are closed under the involutive operation on E .

⁵There is a stronger conjecture that is also concerned with symmetries of I . In Section 2.5.3 we state this conjecture (see Conjecture 2.5.16).

(i) \implies (ii). If $g_0^{-1}g_1 \in \mathbb{G}(E_0)$ and $g_1^{-1}g_0 \in \mathbb{G}(E_1)$, then, by component 2-acyclicity of \mathbb{G} we have $E(g_0, g_1) \subseteq E_0 \cap E_1$. So $g_0^{-1}g_1 \in \mathbb{G}(E(g_0, g_1)) \subseteq \mathbb{G}(E_0 \cap E_1)$. This is equivalent to $g_0\mathbb{G}(E_1 \cap E_0) \cap g_1\mathbb{G}(E_1 \cap E_0) \neq \emptyset$.

(i) \implies (ii). That $\mathbb{G}(E_0 \cap E_1) \subseteq \mathbb{G}(E_0) \cap \mathbb{G}(E_1)$ is clear. Now let $g \in \mathbb{G}(E_0) \cap \mathbb{G}(E_1)$. Then we have $1^{-1}g \in \mathbb{G}(E_0)$ and $g^{-1}1 \in \mathbb{G}(E_1)$. So $1\mathbb{G}(E_0 \cap E_1) \cap g\mathbb{G}(E_0 \cap E_1) \neq \emptyset$. This is equivalent to $g \in \mathbb{G}(E_0 \cap E_1)$.

(ii) \implies (i). $E(g_0, g_1) = \bigcap \{ E' \subseteq E \mid g_0^{-1}g_1 \in \mathbb{G}(E') \}$ for $g_0, g_1 \in \mathbb{G}$ with the same source.

Now we can treat the general case. Using (iii) we can rewrite the statement of coset ℓ -acyclicity to: for all $((g_i, E_i))_{i \in Z_n}$

$$g_i^{-1}g_{i+1} \in \mathbb{G}(E_i) \implies g_{j-1}\mathbb{G}(E_{j-1}) \cap g_j\mathbb{G}(E_j) \cap g_{j+1}\mathbb{G}(E_{j+1})$$

as

$$\begin{aligned} & g_j\mathbb{G}(E_{j-1} \cap E_j) \cap g_{j+1}\mathbb{G}(E_j \cap E_{j+1}) \\ &= g_j(\mathbb{G}(E_{j-1}) \cap \mathbb{G}(E_j)) \cap g_{j+1}(\mathbb{G}(E_j) \cap \mathbb{G}(E_{j+1})) \\ &= g_j\mathbb{G}(E_{j-1}) \cap g_j\mathbb{G}(E_j) \cap g_{j+1}(\mathbb{G}(E_j) \cap g_{j+1}\mathbb{G}(E_{j+1})) \\ &= g_{j-1}\mathbb{G}(E_{j-1}) \cap g_j\mathbb{G}(E_j) \cap g_j(\mathbb{G}(E_j) \cap g_{j+1}\mathbb{G}(E_{j+1})) \\ &= g_{j-1}\mathbb{G}(E_{j-1}) \cap g_j(\mathbb{G}(E_j) \cap g_{j+1}\mathbb{G}(E_{j+1})) \end{aligned}$$

So if we take the condition of coset ℓ -acyclicity and plug in $E_i = E(g_i, g_{i+1})$ then we get the condition of component ℓ -acyclicity. On the other hand component ℓ -acyclicity specifies the condition of coset ℓ -acyclicity for the minimal E_i with $g_i^{-1}g_i \in \mathbb{G}(E_i)$. So the larger sets intersect as well in this manner. \square

2 Extension problems

An *extension problem* is the task of extending a partial symmetry to a global symmetry where the global symmetry may act on an extension of the underlying object of the partial symmetry. Extension problems arise naturally in various areas of mathematics, and are of considerable interest in their own right.

A famous construction of how to solve extension problems for groups are the so-called HNN extensions [22, 28] (HNN are the initials of the inventors Higman, Neumann and Neumann): the HNN extension of a group G relative to some partial automorphism f is a supergroup of G in which f can be extended to an inner automorphism. HNN extensions occupy a significant position in combinatorial and algorithmic group theory.

Witt's theorem [3, 52] is another example of a result that can be phrased as the solvability of an extension problem. It says that every partial isometry of a finite-dimensional quadric space over a field of characteristic different from 2 can be extended to a total isometry of that space.

This 'Witt-property', i.e., that partial isometries can be extended to full isometries, has a model-theoretic analogue: homogeneous structures. A structure is *homogeneous* if every finite partial automorphism can be extended to a full automorphism. Investigations of the automorphism group of homogeneous structures led to the following extension problem: given a finite structure \mathcal{A} , find a finite extension \mathcal{B} over which every partial automorphism of \mathcal{A} can be extended to an automorphism of \mathcal{B} . We say that such a \mathcal{B} *solves the extension problem for \mathcal{A}* . In the following section we discuss the history of this line of research and develop some basic terminology.

There is a related topic concerned with the solvability of extension problems for finite metric spaces [2, 11, 37, 39, 42, 50] which goes beyond the scope of this thesis. Nevertheless, we note that these results are very much in the style of the results presented here, and there has been quite some cross-pollination between these two fields.

In Section 2.3 we discuss the solvability of extension problems for naked structures, i.e., sets with partial bijections. This seemingly trivial task becomes interesting when we put constraints on the automorphism group of the solution. In [5, 30] a similar approach is taken, but there the focus is on the decidability of the solvability of an extension problem, whereas we provide general existence results. Furthermore, the kinds of constraints considered in [5, 30] are fundamentally different from the ones we consider here.

Basic definitions and notations

A *partial bijection* f of a set A is a function $f: X \rightarrow Y$ that is a bijection between two subsets $X, Y \subseteq A$. The *domain* of f is X and the *image* of f is Y ; we denote the former by $\text{dom}(f)$ and the latter by $\text{img}(f)$. We say that f is total if $X = Y = A$. In that case f is a permutation of A .

Set-theoretically, a partial bijection is a relation $f \subseteq A \times A$ s.t. for each $a \in A$ there is at most one $b \in A$ with $(a, b) \in f$, and for each $b \in A$ there is at most one $a \in A$ with $(a, b) \in f$. We have two ways of restricting a partial bijection f . First, the common restriction: for $Z \subseteq \text{dom}(f)$ we define the *restriction of f to Z* to be

$$f|_Z := f \cap (Z \times \text{img}(f)).$$

Second, a model-theoretic restriction where f is seen as a relation on A : for $B \subseteq A$ we define the *relational restriction of f to B* to be

$$f||_B := f \cap (B \times B).$$

If $a \in \text{dom}(f)$ we write af for the image of a under f , i.e., af denotes the unique element b s.t. $(a, b) \in f$. We use this convention of right application for all functions between structures. When we write an expression like $af = b$ we implicitly assume that f is defined on a . For a set $B \subseteq A$ and a tuple $\mathbf{b} = (b_1, \dots, b_n) \subseteq \text{dom}(f)$ we define

$$Bf := \{bf \mid b \in \text{dom}(f) \cap B\} \quad \text{and} \quad \mathbf{b}f := (b_1f, \dots, b_nf).$$

The *inverse* of f is defined as $f^{-1} = \{(b, a) \mid (a, b) \in f\}$. It can also be described extensionally by: $bf^{-1} = a$, if and only if $af = b$. The composition of two partial bijections f, g is given by $fg := \{(a, b) \mid \text{there is a } c \text{ s.t. } (a, c) \in f, (c, b) \in g\}$. It is easy to see that $\text{dom}(fg) = \text{dom}(g)f^{-1}$, $\text{img}(fg) = \text{img}(f)g$ and $a(fg) = (af)g$. We say that g is an *extension of f* (or f is a *restriction of g*) if $\text{dom}(f) \subseteq \text{dom}(g)$ and $xg = xf$ for $x \in \text{dom}(f)$. In set-theoretic terminology we can write ‘ g extends f ’ also as $f \subseteq g$.

A partial automorphism f of a structure $\mathcal{A} = (A, (R^A)_{R \in \sigma})$ is a partial bijection of A s.t. $f: \mathcal{A}|_{\text{dom}(f)} \xrightarrow{\text{iso}} \mathcal{A}|_{\text{img}(f)}$. A total partial automorphism is simply an automorphism. For a structure \mathcal{A} we write $\text{Sym}(\mathcal{A})$ for its automorphism group. For a sets X , $\text{Sym}(X)$ simply are the permutations of X .

2.1 An overview of classical results

We give an historical overview of the development of the theory of extension problems and introduce some basic definitions.

To describe the origin of this line of research we need some notions regarding Fraïssé limits and generic automorphisms. These notions are not required for the understanding of the results presented here but rather are used to provide some motivation. For the definitions of Fraïssé limits and generic automorphism we refer the reader to the survey article [29] by Macpherson.

We recall: a structure \mathcal{B} solves the extension problem for the structure \mathcal{A} if \mathcal{B} is a superstructure of \mathcal{A} in which every partial automorphism of \mathcal{A} has an extension to an automorphism. The question whether there are finite solutions for extension problems for finite structures arose in the search for ‘typical’ elements of the automorphism group of homogeneous structures, so called generic automorphisms. In order to prove that the random graph (the Fraïssé limit of the class of finite graphs) has generic automorphisms, Truss [48] shows that any finite simple graph with a given partial automorphism f admits a finite extension over which f extends to a full automorphism. The idea of generic elements is extended to the notion of generic sequences by Hodges, Hodkinson, Lascar and Shelah [23]. In order to show that the random graph has generic sequences, Hrushovski [25] shows the existence of finite solutions to extension problems for finite graphs.

Theorem 2.1.1 (Hrushovski). *Every finite simple graph G has finite extensions over which every partial automorphism of G can be extended to an automorphism.*

To be precise, Truss’ result proves his result for finite structures of width 2, whereas Hrushovski’s result is about simple graphs. Nevertheless, it is easy to extend Hrushovski’s result to directed graphs and subsequently to structures of width 2 by ‘combining’ solutions. Generalising Hrushovski’s Theorem to structures of greater width is not that straightforward, but can be done, as shown by Herwig [20].

Theorem 2.1.2 (Herwig). *Every finite structure \mathcal{A} has finite extensions over which every partial automorphism of \mathcal{A} can be extended to an automorphism.*

Hrushovski’s original proof is not particularly complex. Nevertheless, Herwig and Lascar [21] give an even simpler ‘kindergarten proof’ which they also generalise to obtain a simpler proof of Herwig’s Theorem.

Herwig’s Theorem gives a general positive answer to the existence of finite solutions of extension problems for finite structures. It is natural to vary this task by restricting the possible solutions to some class C . These ‘restricted’ extension problems are also motivated by the search for generic sequences, just as the initial question was.


An example of a positive result in this direction appears in [20], where it is shown that every extension problem for a triangle free simple graph has a solution which is also triangle free.

For a general discussion about variants of the theorems of Hrushovski and Herwig for restricted structural classes we have to refine our notation. We motivate the to-be-proposed notions by the example of coloured- ℓ -cycle-free graphs.¹ A coloured- ℓ -cycle-free graph is an edge-coloured graph with ℓ colours, i.e., a structure $\mathcal{A} = (A, (E_i^{\mathcal{A}})_{i \in \mathbb{Z}_\ell})$ over the signature $\sigma_\ell = \{E_i \mid i \in \mathbb{Z}_\ell\}$, s.t. there is no sequence $(a_i)_{i \in \mathbb{Z}_\ell}$ of vertices with $(a_i, a_{i+1}) \in E_i^{\mathcal{A}}$. Let CCF_ℓ be the class of all coloured- ℓ -cycle-free graphs for fixed $\ell \in \mathbb{N}$. Consider the structure $\mathcal{A} \in \text{CCF}_4$ depicted below

¹The class of coloured- ℓ -cycle-free graphs is very similar to the class of ‘cycle-free ℓ -partitioned graphs’ discussed in [21].

2 Extension problems

with domain $\{a_0, a_1, a_2, a'_0, a'_1, a'_2\}$, $(R_i^A)_{i \in \mathbb{Z}_4}$ represented by the solid arrows, and with a designated partial automorphism f indicated by the dashed arrows.



The extension problem for \mathcal{A} has no solution in CCF_4 , not even an infinite one: for any solution \mathcal{B} the image a_3 of a'_1 under the extension of f induces a coloured 4-cycle $(a_i)_{i \in \mathbb{Z}_4}$. Thus, for a sensible rephrasing of Herwig's Theorem, we have to take the general solvability of the extension problem into account. Also note that \mathcal{A} has solutions in CCF_4 when we exclude f . So by restricting the set of partial automorphisms the task of finding solutions may become possible again.

From these insights the following definitions are natural. (In the following we let P always be a finite set with some involution $(\cdot)^{-1}$. See Section 0.2 for basic notations regarding such P .)

Definition 2.1.3. A P -extension problem \mathcal{X} over the signature σ is a tuple of the form $(X, (R^X)_{R \in \sigma}, (p^X)_{p \in P})$ s.t. $(p^X)_{p \in P}$ is a collection of partial automorphisms of the underlying structure $\text{Str}(\mathcal{X}) := (X, (R^X)_{R \in \sigma})$, and $(p^{-1})^{\mathcal{X}} = (p^{\mathcal{X}})^{-1}$ for all $p \in P$.

A solution $\mathcal{S} = (S, (R^S)_{R \in \sigma}, (p^S)_{p \in P})$ of a P -extension problem \mathcal{X} has the same format as an extension problem, where in addition the underlying structure of \mathcal{S} extends the one underlying \mathcal{X} and p^S is total and extends $p^{\mathcal{X}}$ for $p \in P$.

Given a class C of structures, we write $\mathcal{X} \in C$ if $\text{Str}(\mathcal{X})$ is in C ; analogously for solutions \mathcal{S} . We also use the following convention, for $u = p_1 \dots p_n \in P^*$ we write $u^{\mathcal{X}} = p_1^{\mathcal{X}} \dots p_n^{\mathcal{X}}$ and $u^{\mathcal{S}} = p_1^{\mathcal{S}} \dots p_n^{\mathcal{S}}$.

Definition 2.1.4. A class C of structures has the *extension property for partial automorphisms* (EPPA for short) if every finite P -extension problem in C that has a solution in C also has a finite solution in C .

Equipped with these definitions we can formulate the following variant of Herwig's Theorem for the class CCF_ℓ , which is proved to be true in [21].

Theorem 2.1.5. *The class of coloured- ℓ -cycle-free graphs, CCF_ℓ , has EPPA.*

Note that Herwig's Theorem (Theorem 2.1.2) can also be phrased in terms of EPPA. In Section 2.2 below we show that every extension problem has a uniquely defined free solution. In particular, every extension problem has a solution.² So, if

²We can also use Fraïssé limits to observe the general solvability of extension problems. Given a finite extension problem over some signature σ , consider the class of all finite σ -structures. This has a Fraïssé limit and since it is homogeneous we can extend the partial automorphisms.

we presuppose the existence of free solutions as an elementary fact, the statements ‘the class of all σ -structures has EPPA’ and ‘every finite extension problem over σ has a solution’ are equivalent.

All of the results presented so far are special cases of a powerful result by Herwig and Lascar. In [21] they present the following sufficient criterion for a class of structures to have EPPA.

A structure \mathcal{B} is \mathcal{A} -free if there is no homomorphism $f: \mathcal{A} \xrightarrow{\text{hom}} \mathcal{B}$. If T is a class of structures, we say that \mathcal{B} is T -free if it is \mathcal{A} -free for all $\mathcal{A} \in T$. A class of structures C is defined *in terms of forbidden homomorphic images* if C is the class of T -free structures for some finite set T of structures.

Theorem 2.1.6 (Herwig-Lascar). *A class of structures that is defined in terms of forbidden homomorphic images has EPPA.*

It is not hard to see that we can deduce all results mentioned so far from this criterion. For example, CCF_ℓ can be defined as the class of structures not containing the homomorphic image of a ‘coloured- ℓ -cycle’. Nonetheless, there are also results that are not direct consequences of the Theorem of Herwig and Lascar. We present some of these now.

Hodkinson and Otto [24] show that the class of conformal structures has EPPA. *Conformal structures* $\mathcal{A} = (A, (R^A)_{R \in \sigma})$ are structures in which every clique K of the Gaifman graph is *guarded*, i.e., there is an $R \in \sigma$ and $\mathbf{a} \in R^A$ s.t. $K \subseteq \mathbf{a}$. Note that this result is not a consequence of the Theorem of Herwig and Lascar, as classes defined in terms of forbidden homomorphic images are closed under weak substructures, but the class of conformal structures is not. Hodkinson and Otto actually provide an even stronger result. To describe this we introduce the notion of the translation hypergraph (see Section 1.4 for definitions regarding hypergraphs, such as conformality and acyclicity).

Definition 2.1.7. The *translation hypergraph* $T(\mathcal{S})$ of a solution \mathcal{S} of an extension problem $\mathcal{X} = (S, (R^X)_{R \in \sigma}, (p^X)_{p \in P})$ is the hypergraph with domain $\bigcup_{u \in P^*} Xu^S$ and hyperedges $\{ Xu^S \mid u \in P^* \}$.

The main result of Hodkinson and Otto is that any finite solution of an extension problem can be ‘unravelling’ to a finite solution whose translation hypergraph is conformal. With Herwig’s Theorem we obtain the following result.

Theorem 2.1.8 (Hodkinson-Otto). *Every finite extension problem has a finite solution with conformal translation hypergraph.*

We want to emphasize that this theorem constitutes a shift in how to discuss extension problems. The theorems we have seen so far are all of the form “...there is some finite solution whose underlying structure has a certain property” whereas the theorem of Hodkinson and Otto has the form “...there is some finite solution s.t. the automorphisms translate X through the underlying domain in a certain way”. This approach shifts the focus from restricting solutions by classes (ultimately constraining their relations) to constraining the automorphisms. If we follow this idea

2 Extension problems

to the end and only put constraints on the automorphisms of a solution, then there is no need for any structural information and we can use plain sets as underlying structures. We do exactly this in Section 2.3 and also discuss how these two approaches (constraining relations vs. constraining automorphisms) can be connected. As a small preview we show how to obtain EPPA for conformal structures from the theorem of Hodkinson and Otto. As a preparation we introduce *tidy* solutions.

Definition 2.1.9. A solution $\mathcal{S} = (S, (R^{\mathcal{S}})_{R \in \sigma}, (p^{\mathcal{S}})_{p \in P})$ of an extension problem $\mathcal{X} = (X, (R^{\mathcal{X}})_{R \in \sigma}, (p^{\mathcal{X}})_{p \in P})$ is *tidy* if $\text{Str}(\mathcal{S})$ is the orbit of $\text{Str}(\mathcal{X})$ w.r.t. the automorphisms $(p^{\mathcal{S}})_{p \in P}$, i.e.,

- (i) for all $a \in S$ there is an $x \in X$ and $u \in P^*$ s.t. $xu^{\mathcal{S}} = a$, and
- (ii) for all $R \in \sigma$ and $\mathbf{a} \in R^{\mathcal{S}}$ there is an $\mathbf{x} \in R^{\mathcal{X}}$ and $u \in P^*$ s.t. $\mathbf{x}u^{\mathcal{S}} = \mathbf{a}$.

We can obtain a tidy solution from any given solution \mathcal{S} simply by passing to the orbit of $\text{Str}(\mathcal{X})$. We may generally think in terms of tidy solutions as all properties that we consider here are preserved when passing to the orbit of $\text{Str}(\mathcal{X})$.

The key feature of a tidy solution \mathcal{S} is that the Gaifman graph of the structure $\text{Str}(\mathcal{S})$ is a weak subgraph of the Gaifman graph of the hypergraph $T(\mathcal{S})$. Using this and finite, conformal, tidy solutions we can show EPPA for the class of conformal structures. If \mathcal{X} is a finite P -extension problem whose underlying structure $\text{Str}(\mathcal{X})$ is conformal and \mathcal{S} is a finite, tidy solution with conformal translation hypergraph then $\text{Str}(\mathcal{S})$ is also conformal: if K is a clique in the Gaifman graph of $\text{Str}(\mathcal{S})$, then, as $T(\mathcal{S})$ is conformal, there is some $u \in P^*$ s.t. $K \subseteq Xu^{\mathcal{S}}$. So $K' = K(u^{-1})^{\mathcal{S}}$ is a clique in the Gaifman graph of $\text{Str}(\mathcal{X})$. As $\text{Str}(\mathcal{X})$ is conformal, there is some guard \mathbf{x} of K' , and so $\mathbf{x}u^{\mathcal{S}}$ is a guard of K .

We end this overview by a conjecture that can implicitly be found in [35]. Originally this conjecture is stated as a theorem but it is based on Conjecture 2.5.16, for which an erroneous proof was given. Our contribution is that we assert the strength of those solutions given by this conjecture in a cleaner way.

Definition 2.1.10. A solution \mathcal{S} of an extension problem $\mathcal{X} = (X, (R^{\mathcal{X}})_{R \in \sigma}, (p^{\mathcal{X}})_{p \in P})$ is ℓ -free if for all $u_1, \dots, u_\ell \in P^*$ with $(u_1 \dots u_\ell)^{\mathcal{S}} = \text{id}$ there are $v_1, \dots, v_\ell \in P^*$ s.t. $u_i^{\mathcal{S}} = v_i^{\mathcal{S}}$, $u_i^{\mathcal{S}} \parallel_X = v_i^{\mathcal{X}}$ for all $i = 1, \dots, \ell$, and $[v_1 \dots v_\ell]_{\text{FG}(P)} = 1$.

Conjecture 2.1.11 (Free Extension Conjecture). *Every finite extension problem has finite, ℓ -free solutions.*

This theorem would constitute the strongest result about extension problems in the literature in general, with two noteworthy exceptions:

1. An even stronger conjecture concerns the symmetry of the solution. In Section 2.7 we discuss the possibility of symmetric solutions.
2. In [42] Solecki generalised the Herwig-Lascar Theorem to ‘coherent’ solutions. We briefly discuss coherent solution in the conclusion (Chapter 3) as well.

The Free Extension Conjecture is reduced to Conjecture 2.5.16. This is one of the main results of this work and much of the theory about extension problems we develop in the following is to provide enough tools to give a proof of this reduction. Ultimately, this reduction is given in Section 2.5.4.

Recall that $u^{\mathcal{S}}\|_X := u^{\mathcal{S}} \cap (X \times X)$. Thus we can translate the condition $u^{\mathcal{S}}\|_X = v^{\mathcal{X}}$ to: if $xu^{\mathcal{S}} = y$ for $x, y \in X$, then $xv^{\mathcal{X}} = y$.

It is easy to see that 1-freeness is a trivial notion. Every solution \mathcal{S} is 1-free as when $(u_1)^{\mathcal{S}} = \text{id}$ then $v_1 = \varepsilon$ is a suitable choice. It is also not hard to see that 2-freeness is equivalent to the following condition:

$$\text{for all } u \in P^* \text{ there is a } v \in P^* \text{ s.t. } u^{\mathcal{S}} = v^{\mathcal{S}} \text{ and } u^{\mathcal{S}}\|_X = v^{\mathcal{X}}.$$

In other words, 2-freeness requires that every partial automorphism obtained by restriction of some element in $\langle p^{\mathcal{S}} \rangle_{p \in P}$ is also in $\langle p^{\mathcal{X}} \rangle_{p \in P}$. Higher degrees of ℓ -freeness are equivalent to some form of ℓ -acyclicity (see Definition 2.3.11 and the following lemma in Section 2.3 below). In Lemma 2.3.13 we prove the following result:

$$\begin{aligned} &\text{If } \ell \geq 3 \text{ and } \mathcal{S} \text{ is an } \ell\text{-free solution of } \mathcal{X}, \\ &\text{then } T(\mathcal{S}) \text{ is an } \ell\text{-acyclic hypergraph.} \end{aligned} \tag{*}$$

which indicates that ℓ -freeness implies some acyclicity properties.³

From (*) we immediately get that the Free Extension Conjecture implies the theorem of Hodkinson and Otto, since acyclic hypergraphs are conformal. We can also use (*) to show the following lemma which in turn can be used to show that the Free extension Conjecture implies the Theorem of Herwig and Lascar.

Lemma 2.1.12. *Let \mathcal{S} be an ℓ -free solution of a P -extension problem \mathcal{X} . Then for all $u_1, \dots, u_\ell \in P^*$ there are $v_1, \dots, v_\ell \in P^*$ s.t. for any solution \mathcal{R} of \mathcal{X} the map*

$$\begin{aligned} f: \bigcup_{i=1}^{\ell} Xu_i^{\mathcal{S}} &\longrightarrow \bigcup_{i=1}^{\ell} Xv_i^{\mathcal{R}} \\ Xu_i^{\mathcal{S}} &\longmapsto xv_i^{\mathcal{R}} \end{aligned}$$

is well-defined.

Proof. We consider the finite subhypergraph of $T(\mathcal{S})$ induced on the union of the $Xu_1^{\mathcal{S}}, \dots, Xu_\ell^{\mathcal{S}}$. We write V for the domain of this hypergraph, which is $\bigcup Xu_i^{\mathcal{S}}$, and R for the set of hyperedges. Note that R is in general a proper superset of $\{Xu_i^{\mathcal{S}} \mid i = 1, \dots, \ell\}$. For each $r \in R$ we fix a $u_r \in P^*$ with $r = Xu_r \cap V$ (for $r = Xu_i^{\mathcal{S}}$ we let u_r be u_i). We recursively construct $(v_r)_{r \in R} \subseteq P^*$ s.t. f defined by $xu_r^{\mathcal{S}} \mapsto xv_r^{\mathcal{R}}$ is well-defined for $x \in V$. The case $|R| = 1$ is trivial. For

³As pointed out, we do not know the existence of ℓ -free solutions. However, it is likely that solutions with property * follows from the already existing theory and the fact that finite hypergraphs admit finite, ℓ -acyclic covers [34]

2 Extension problems

$R = \{r, s\}$ we let v_r be arbitrary and set $v_s = wv_r$ where w is s.t. $w^{\mathcal{S}} = (u_s u_r^{-1})^{\mathcal{S}}$ and $w^{\mathcal{X}} = (u_s u_r^{-1})^{\mathcal{S}} \parallel_X$ (using 2-acyclicity). To show that f is well-defined we need to prove that $xv_r^{\mathcal{R}} = yv_s^{\mathcal{R}}$ if $xu_r^{\mathcal{S}} = yu_s^{\mathcal{S}}$. This follows from

$$\begin{aligned} xu_r^{\mathcal{S}} = yu_s^{\mathcal{S}} &\implies x = y(u_s u_r^{-1})^{\mathcal{S}} \implies x = yw^{\mathcal{S}} \implies x = yw^{\mathcal{X}} \\ &\implies x = yw^{\mathcal{R}} \implies x = y(v_s v_r^{-1})^{\mathcal{R}} \implies xv_r^{\mathcal{R}} = yv_s^{\mathcal{R}}. \end{aligned}$$

Now we consider $|R| = n + 1$. By (*), (V, R) is an acyclic hypergraph and so it has a hyperleaf, i.e., there are distinct $r, s \in R$ s.t. $s \cap t \subseteq r$ for all $t \in R$ distinct from s . Let v_t , for $t \in R \setminus \{s\}$, be given by the recursion. We then define v_s similarly to the case $|R| = 2$, only now depending on v_r . The $(v_t)_{t \in R}$ have the desired property. \square

The map f of the aforementioned lemma induces a homomorphism

$$f: \text{Str}(\mathcal{S}, u_1, \dots, u_\ell) \xrightarrow{\text{hom}} \text{Str}(\mathcal{R}),$$

where $\text{Str}(\mathcal{S}, u_1, \dots, u_n)$ is the weak substructure of \mathcal{S} induced by the image of $\text{Str}(\mathcal{X})$ under the $u_1^{\mathcal{S}}, \dots, u_\ell^{\mathcal{S}}$, i.e., $\text{Str}(\mathcal{S}, u_1, \dots, u_n) := (\bigcup_{i=1}^{\ell} X u_i^{\mathcal{S}}, (\bigcup_{i=1}^{\ell} R^{\mathcal{X}} u_i^{\mathcal{S}})_{R \in \sigma})$.

We can now show how the Free Extension Conjecture implies the Theorem of Herwig and Lascar. Let \mathcal{S} be a finite, tidy, ℓ -free solution of the extension problem \mathcal{X} . If \mathcal{X} has a T -free solution \mathcal{R} and ℓ is greater than $\sum_{R \in \sigma} |R^T|$ for each $\mathcal{A} \in T$, then \mathcal{S} is also T -free. Otherwise, a homomorphism from $\mathcal{A} \in T$ to \mathcal{S} would induce a homomorphism from \mathcal{A} to \mathcal{R} via the homomorphism f of the previous lemma. In the proof of Lemma 2.3.10 this argument is given in more detail.

Finally, we want to mention some classes of structures for which EPPA is still an open problem which also remain open if the Free Extension Conjecture were true. One prominent example is the class of tournaments. A positive answer to this problem would imply the existence of generic sequences for the universal, homogeneous tournament (cf. [29]). Another example for which we do not know the answer are the classes of ℓ -acyclic simple graphs. We know that EPPA holds for $\ell = 3$, which are simply the triangle free graphs. However, for greater values the question is still open. Note that a naive approach for $\ell = 4$, by ‘forbidding’ the homomorphic image of a 4-cycle using the Theorem of Herwig and Lascar does not work, as edges are also homomorphic images of 4-cycles.

2.2 Free solutions

Every extension problem has a free solution which is uniquely defined up to isomorphism. In general, the free solution is infinite and so cannot be used directly to provide finite solutions to finite extension problems. However, its universal nature makes it an important object to study.

Definition 2.2.1. The free solution of an extension problem $\mathcal{X} = (X, (R^{\mathcal{X}})_{R \in \sigma}, (p^{\mathcal{X}})_{p \in P})$ is the tidy solution $\mathcal{U} = (U, (R^{\mathcal{U}})_{R \in \sigma}, (p^{\mathcal{U}})_{p \in P})$ with the *freeness property*:

$$u^{\mathcal{U}} \parallel_X = u^{\mathcal{X}}$$

for all $u \in \text{FG}(P)$.

The freeness property is equivalent to

$$xu^{\mathcal{U}} = y \iff xu^{\mathcal{X}} = y \quad \text{for all } x, y \in X \text{ and } u \in \text{FG}(P),$$

or

$$xu^{\mathcal{U}} = yv^{\mathcal{U}} \iff x[uv^{-1}]_{\text{FG}(P)}^{\mathcal{X}} = y \quad \text{for all } x, y \in X \text{ and } u, v \in P^*$$

(note that $[uv^{-1}]_{\text{FG}(P)}$ is just $\text{red}(uv^{-1})$).

It is important that the freeness property requires $u^{\mathcal{U}}|_X = u^{\mathcal{X}}$ only for $u \in \text{FG}(P)$ and not for $u \in P^*$. Otherwise the freeness property would be too strict (consider $u = pp^{-1}$, where $p^{\mathcal{X}}$ is not total).

As always we need to justify the definite article for *the* free solution, i.e., we have to show its existence (Lemma 2.2.2) and uniqueness (Lemma 2.2.4 with preceding remark).

Lemma 2.2.2. *Every extension problem has a free solution.*

Proof. We construct a free solution $\mathcal{U} = (U, (R^{\mathcal{U}})_{R \in \sigma}, (p^{\mathcal{U}})_{p \in P})$ of the extension problem $\mathcal{X} = (X, (R^{\mathcal{X}})_{R \in \sigma}, (p^{\mathcal{X}})_{p \in P})$ in multiple steps. Set U to be the reduced product $U = (X \times P^*) / \sim$ with

$$(x, u) \sim (y, v) \quad \text{if } x[uv^{-1}]_{\text{FG}(P)}^{\mathcal{X}} = y.$$

Then X can be seen as a subset of U via the injective map $x \mapsto [(x, 1)]_{\sim}$.

We define the automorphisms $p^{\mathcal{U}}$ by $[x, u]p^{\mathcal{U}} := [x, up]$. It is easy to check that the freeness property holds. In particular, we have $p^{\mathcal{X}} \subseteq p^{\mathcal{S}}$.

For $R \in \sigma$ we define $R^{\mathcal{U}}$ as the orbit of $R^{\mathcal{X}}$ under $(p^{\mathcal{U}})_{p \in P}$, i.e., $R^{\mathcal{U}} := \{\mathbf{x}u^{\mathcal{U}} : \mathbf{x} \in R^{\mathcal{X}}, u \in \text{FG}(P)\}$. Clearly the $R^{\mathcal{U}}$ are constructed in such a way that the $p^{\mathcal{U}}$ are automorphisms of $(U, (R^{\mathcal{U}})_{R \in \sigma})$, and that $\text{Str}(\mathcal{X})$ is a weak substructure of $\text{Str}(\mathcal{U})$. We are left to show that $\text{Str}(\mathcal{X})$ is a substructure of $\text{Str}(\mathcal{U})$. Let $\mathbf{x} \subseteq X$ be in $R^{\mathcal{U}}$. We argue that then also $\mathbf{x} \in R^{\mathcal{X}}$. Since $\mathbf{x} \in R^{\mathcal{U}}$, there are $u \in \text{FG}(P)$ and $\mathbf{y} \in R^{\mathcal{X}}$ with $\mathbf{x} = \mathbf{y}u^{\mathcal{U}}$. Then by the freeness condition $\mathbf{x} = \mathbf{y}u^{\mathcal{X}}$ and thus $\mathbf{x} \in R^{\mathcal{X}}$.

It is easy to check that \mathcal{U} is tidy as well. □

To show that the free solution is unique up to isomorphism we show that it is universal. For that we need a notion of homomorphisms of solutions.

Definition 2.2.3. A homomorphism $f: \mathcal{U} \xrightarrow{\text{hom}} \mathcal{S}$ between solutions of an extension problem $\mathcal{X} = (X, (p^{\mathcal{X}})_{p \in P}, (R^{\mathcal{X}})_{R \in \sigma})$ is a homomorphism $f: \text{Str}(\mathcal{U}) \xrightarrow{\text{hom}} \text{Str}(\mathcal{S})$ s.t.

- (i) $f|_X = \text{id}_X$ and
- (ii) $p^{\mathcal{U}}f = fp^{\mathcal{S}}$ for all $p \in P$.

2 Extension problems

Observe that if $f: \mathcal{U} \xrightarrow{\text{hom}} \mathcal{S}$ is a homomorphism between two solutions then $xu^{\mathcal{U}}f = xu^{\mathcal{S}}$ for all $x \in X$ and $u \in P^*$. Hence, there is at most one homomorphism from a tidy solution \mathcal{U} to a solution \mathcal{S} . Thus, two tidy solutions are isomorphic if, and only if, they are mutually homomorphic (the argument is the same as for the related fact for P -generated groups, see ‘Groups with generators’ in Section 0.2). So it suffices to show that the free solution is universal to establish its uniqueness.

Lemma 2.2.4. *Let \mathcal{U} be a free solution and \mathcal{S} some solution of the extension problem $\mathcal{X} = (X, (R^{\mathcal{X}})_{R \in \sigma}, (p^{\mathcal{X}})_{p \in P})$. Then*

$$\begin{aligned} f: U &\rightarrow S \\ c &\mapsto xu^{\mathcal{S}}, \quad \text{where } x \in X, u \in P^* \text{ s.t. } c = xu^{\mathcal{U}} \end{aligned}$$

is a homomorphism from \mathcal{U} to \mathcal{S} .

Proof. We need to show that f is well-defined. Let $x_1u_1^{\mathcal{U}} = x_2u_2^{\mathcal{U}}$ for $x_1, x_2 \in X$ and $u_1, u_2 \in \text{FG}(P)$. Then $x_1u_1^{\mathcal{S}} = x_2u_2^{\mathcal{S}}$ follows from

$$\begin{aligned} x_1u_1^{\mathcal{U}} = x_2u_2^{\mathcal{U}} &\implies x_1(u_1u_2^{-1})^{\mathcal{U}} = x_2 \implies x_1[u_1 \cdot u_2^{-1}]_{\text{FG}(P)}^{\mathcal{X}} = x_2 \\ &\implies x_1(u_1u_2^{-1})^{\mathcal{S}} = x_2 \implies x_1u_1^{\mathcal{S}} = x_2u_2^{\mathcal{S}}. \end{aligned}$$

Now we show that f is a homomorphism between $\text{Str}(\mathcal{U})$ and $\text{Str}(\mathcal{X})$. Let $\mathbf{c} = \mathbf{x}u^{\mathcal{U}}$ for some $\mathbf{x} \subseteq X$ and $u \in P^*$. Then

$$\mathbf{c} \in R^{\mathcal{U}} \implies \mathbf{x}u^{\mathcal{U}} \in R^{\mathcal{U}} \implies \mathbf{x} \in R^{\mathcal{X}} \implies \mathbf{x}u^{\mathcal{S}} \in R^{\mathcal{S}} \implies \mathbf{c}f \in R^{\mathcal{S}}.$$

Last, we show that f satisfies properties (i) and (ii) of the homomorphism requirement. Clearly f fulfils property (i). Towards (ii) let $c = xu^{\mathcal{U}}$ for some $x \in X$ and $u \in P^*$. Then

$$cp^{\mathcal{U}}f = (xu^{\mathcal{U}})p^{\mathcal{U}}f = x(up)^{\mathcal{U}}f = x(up)^{\mathcal{S}} = (xu^{\mathcal{S}})p^{\mathcal{S}} = (xu^{\mathcal{U}}f)p^{\mathcal{S}} = cfp^{\mathcal{S}}. \quad \square$$

Free solutions have a special status in most EPPA-statements. Typically, we consider classes of structures C for which an extension problem \mathcal{X} has a solution, if and only if the free solution of \mathcal{X} is in C . Examples of such classes are classes defined in terms of forbidden homomorphic images, and the class of conformal structures. Classes for which this is not true are the class of tournaments, and the class of 4-acyclic simple graphs, indicating why our current tools do not suffice to tackle these questions.

We finish this discussion about the free solution by comparing freeness to ℓ -freeness. It is not the case that freeness is the limit of ℓ -freeness, the limit is weaker. So, we do not use the term ‘local freeness’ to describe ℓ -freeness for some unspecified $\ell \in \mathbb{N}$ but rather the term *approximate freeness*. Approximate freeness also subsumes the yet to be introduced notions of parallel ℓ -freeness, sequential ℓ -freeness and cluster ℓ -freeness (introducing these notions is the goal of the next section).


We describe a class of extension problems for which freeness is the limit of ℓ -freeness. We recall the definitions:

freeness: for all $u \in \text{FG}(P)$ $u^{\mathcal{X}} = u^{\mathcal{S}} \parallel_X$.


ℓ -freeness: for all $u_1, \dots, u_\ell \in P^*$ with $(u_1 \dots u_\ell)^{\mathcal{S}} = \text{id}$ there are $v_1, \dots, v_\ell \in P^*$ s.t. $v_i^{\mathcal{S}} = u_i^{\mathcal{S}}, v_i^{\mathcal{X}} = u_i^{\mathcal{S}} \parallel_X$ for all $i = 1, \dots, \ell$, and $[v_1 \dots v_\ell]_{\text{FG}(P)} = 1$.

First we check that freeness implies ℓ -freeness for all $\ell \in \mathbb{N}$. This requires a little trick: if $(u_1 \dots u_\ell)^{\mathcal{U}} = \text{id}$, then $v_1 = [u_1]_{\text{FG}(P)}, \dots, v_\ell = [u_\ell]_{\text{FG}(P)}$ seem to be appropriate witnesses for ℓ -freeness. However, we cannot guarantee that $[v_1 \dots v_\ell]_{\text{FG}(P)} = 1$. The trick is that if we instead take $v_\ell = [u_{\ell-1}^{-1} \dots u_1^{-1}]_{\text{FG}(P)}$ then v_1, \dots, v_ℓ have the right properties.

In general, freeness is not the limit of ℓ -freeness. Consider the following extension problem $\mathcal{X} = (X, p^{\mathcal{X}})$ over the empty signature where $P = \{p\}$ and $p^{-1} = p$:



And the solution $\mathcal{S} = (S, p^{\mathcal{S}})$:



Clearly \mathcal{S} is not the free solution of \mathcal{X} but it is ℓ -free for all $\ell \in \mathbb{N}$. Since $p^{\mathcal{S}}$ is an involution we have that $(p^n)^{\mathcal{S}} = \varepsilon^{\mathcal{S}}$ if n is even and that $(p^n)^{\mathcal{S}} = p^{\mathcal{S}}$ if n is odd. Also note that $\varepsilon^{\mathcal{X}} = \varepsilon^{\mathcal{S}} \parallel_X$ and $p^{\mathcal{X}} = p^{\mathcal{S}} \parallel_X$. So if now $p^{n_1}, \dots, p^{n_\ell} \in P^*$ are given with $(p^{n_1} \dots p^{n_\ell}) = \text{id}$ we set v_i to be equal to ε if n_i is even and otherwise to be equal to p . Then $[v_1 \dots v_\ell]_{\text{FG}(P)} = 1$ as $n_1 + \dots + n_\ell$ is even and so there is an even number k of n_i that are odd and thus $v_1 \dots v_\ell = p^k$ which evaluates in $\text{FG}(P)$ to 1.

2 Extension problems

However, freeness is the limit of ℓ -freeness, if the partial automorphisms of \mathcal{X} are ‘independent’, in the sense that $p^{\mathcal{X}} \subseteq u^{\mathcal{X}} \implies u = p$ for $u \in P^*$: let \mathcal{S} be ℓ -free for all $\ell \in \mathbb{N}$ and $u = p_1 \dots p_n \in \text{FG}(P)$. By $(n+1)$ -freeness we obtain v_1, \dots, v_n, v for $(p_1 \dots p_n u^{-1})^{\mathcal{U}} = \text{id}$ s.t.

- (i) $v_i^{\mathcal{S}} = p_i^{\mathcal{S}}, v_i^{\mathcal{X}} = p_i^{\mathcal{S}} \parallel_X$
- (ii) $v^{\mathcal{S}} = (u^{-1})^{\mathcal{S}}, v^{\mathcal{S}} = (u^{-1})^{\mathcal{S}} \parallel_X$
- (iii) $[v_1 \dots v_n v]_{\text{FG}(P)} = 1$.

From (i) we get that $p_i^{\mathcal{X}} \subseteq p_i^{\mathcal{U}} \parallel_X = v_i^{\mathcal{X}}$ and so $v_i = p_i$. This together with (iii) implies that $v = u^{-1}$. Then using (ii) we get that $(u^{-1})^{\mathcal{U}} \parallel_X = (u^{-1})^{\mathcal{X}}$ and thus $u^{\mathcal{U}} \parallel_X = u^{\mathcal{X}}$.⁴

2.3 Structural properties of plain extension problems

Plain extension problems are extension problems over the empty signature, i.e., the underlying structures are naked sets. We denote a plain extension problem by a tuple $\mathcal{X} = (X, (p^{\mathcal{X}})_{p \in P})$. Plain extension problems form the easiest kind of extension problem and it is trivial to obtain solutions for them.

Lemma 2.3.1. *Every finite plain extension problem has a finite solution.*

Proof. We can extend any partial bijection f of a finite set X to a permutation of X . As $|\text{dom}(f)| = |\text{img}(f)|$, we have $|X \setminus \text{dom}(f)| = |X \setminus \text{img}(f)|$ and thus there is a bijection $g: X \setminus \text{dom}(f) \rightarrow X \setminus \text{img}(f)$ and so $f \cup g$ is a permutation of X extending f . \square

The task of finding solutions to plain extension problems becomes interesting again when we introduce constraints on the automorphisms of the solutions. Freeness and ℓ -freeness are examples of such constraints. They restrict the way in which the solutions are allowed to move the elements of X by their automorphisms. In this section we want to introduce further properties of this kind and relate them to the classical results about extension problems (Hrushovski, Herwig, Herwig-Lascar). Explicitly, the properties we introduce in this section are: parallel ℓ -freeness, sequential ℓ -freeness, and cluster ℓ -freeness. All of these are weaker than ℓ -freeness. We use the term *approximate freeness* as a generic term that encompasses all of these freeness notions.

We also want to point out that in Section 2.4.2 we translate the approximate freeness properties to the setting of inverse monoids. In this setting the abstract freeness properties are reduced to its essential core. So the reader may peek at this section (in particular Definition 2.4.10) to get an idea of these properties in advance.

⁴It can be argued that this ‘independence’ condition on the partial automorphisms of \mathcal{X} is very strong. However, we want to mention that we can always extend the domain of the partial automorphisms by new elements to ensure this property.

2.3 Structural properties of plain extension problems

In [36] solutions \mathcal{S} to plain extension problems $\mathcal{X} = (X, p^{\mathcal{X}})$ are viewed as ‘discrete manifolds’ over X . For each $u \in P^*$ we have a *chart*

$$\begin{aligned} \varphi_u^{\mathcal{S}}: Xu^{\mathcal{S}} &\rightarrow X \\ a &\mapsto a(u^{-1})^{\mathcal{S}}. \end{aligned}$$

The set of coordinate domains of this ‘manifold’ is given by the translation hypergraph $T(\mathcal{S})$, and the transition map $\tau_{u,v}^{\mathcal{S}} := (\varphi_u^{\mathcal{S}})^{-1}\varphi_v^{\mathcal{S}}$ for two charts $\varphi_u^{\mathcal{S}}$ and $\varphi_v^{\mathcal{S}}$ is given by

$$\tau_{u,v}^{\mathcal{S}} = (uv^{-1})^{\mathcal{S}}\|_X.$$

For the free solution \mathcal{U} the transition map can be simplified to $\tau_{u,v}^{\mathcal{U}} = [uv^{-1}]_{\text{FG}(P)}^{\mathcal{X}}$.

A related but slightly different interpretation of \mathcal{S} is that of an amalgamation. We can think of \mathcal{S} as being the result of gluing together copies of X . For each $u \in P^*$ we have a *patch*, a copy of X , located in \mathcal{S} , and the patches for u and v are glued together according to $\tau_{u,v}^{\mathcal{S}}$.

In the view of \mathcal{S} as a manifold, $\tau_{u,v}$ describes the partial automorphism of X induced by the charts, whereas in the view of \mathcal{S} as an amalgamation, the map $\tau_{u,v}$ describes the intersection of two patches. In the free solution \mathcal{U} these intersections are minimal as for any other solution \mathcal{S} we have $\tau_{u,v}^{\mathcal{U}} = [uv^{-1}]_{\text{FG}(P)}^{\mathcal{X}} \subseteq \tau_{u,v}^{\mathcal{S}}$.

We introduce properties $(*\ell)$ for solutions \mathcal{S} . Loosely speaking $(*\ell)$ says that amalgamations of size at most ℓ that occur in \mathcal{S} also occur in the free solution \mathcal{U} (but maybe at a different location):

$$\begin{aligned} \text{for all } \mathbf{a} \in S^\ell \text{ and } u_0, u_1 \in P^* \text{ with } \mathbf{a} \in Xu_0^{\mathcal{S}} \cap Xu_1^{\mathcal{S}} \text{ there are} & \quad (*^\ell) \\ v_0, v_1 \in P^* \text{ and } \mathbf{c} \in Xv_0^{\mathcal{U}} \cap Xv_1^{\mathcal{U}} \text{ s.t. } \mathbf{c}\varphi_{u_0}^{\mathcal{S}} = \mathbf{a}\varphi_{v_0}^{\mathcal{U}} \text{ and } \mathbf{c}\varphi_{u_1}^{\mathcal{S}} = \mathbf{a}\varphi_{v_1}^{\mathcal{U}}. & \end{aligned}$$

See Figure 2.1 for a sketch of $(*\ell)$.

We can give an alternative characterisation of $(*\ell)$ as ‘parallel ℓ -freeness’.

Definition 2.3.2. A solution \mathcal{S} of a plain P -extension problem \mathcal{X} is *parallel ℓ -free* if for each $u \in P^*$ and $\mathbf{x} \in \text{dom}(u^{\mathcal{S}}\|_X)$ of length ℓ there is a $v \in P^*$ s.t. $\mathbf{x}v^{\mathcal{X}} = \mathbf{x}u^{\mathcal{S}}$.

It is fairly easy to see that a solution \mathcal{S} has $(*\ell)$ if, and only if, it is parallel ℓ -free. We provide the formal argument (it might be helpful to consider Figure 2.1 in parallel): let \mathcal{S} be parallel ℓ -free and $\mathbf{a} \in Xu_0^{\mathcal{S}} \cap Xu_1^{\mathcal{S}}$ of length ℓ . Set $\mathbf{x} = \mathbf{a}\varphi_{u_0}^{\mathcal{S}}$ and $\mathbf{y} = \mathbf{a}\varphi_{u_1}^{\mathcal{S}}$. Then, as $\mathbf{x}(u_0u_1^{-1})^{\mathcal{S}}\|_X = \mathbf{y}$, there is a $v \in P^*$ s.t. $\mathbf{x}v^{\mathcal{X}} = \mathbf{y}$. So $\mathbf{y} \in Xv^{\mathcal{U}} \cap X\varepsilon^{\mathcal{U}}$ and $\mathbf{y}\varphi_v^{\mathcal{U}} = \mathbf{x}$ and $\mathbf{y}\varphi_\varepsilon^{\mathcal{U}} = \mathbf{y}$. Now let \mathcal{S} have property $(*\ell)$ and $\mathbf{x} \in \text{dom}(u^{\mathcal{S}}\|_X)$ of length ℓ . Then $\mathbf{x} \in X\varepsilon^{\mathcal{S}} \cap X(u^{-1})^{\mathcal{S}}$ and thus we get $v_0, v_1 \in P^*$ and $\mathbf{c} \in Xv_0^{\mathcal{U}} \cap Xv_1^{\mathcal{U}}$ s.t. $\mathbf{c}\varphi_{v_0}^{\mathcal{U}} = \mathbf{x}$ and $\mathbf{c}\varphi_{v_1}^{\mathcal{U}} = \mathbf{x}u^{\mathcal{S}}$. Hence, $\mathbf{x}u^{\mathcal{S}} = \mathbf{x}\tau_{v_0,v_1}^{\mathcal{U}} = \mathbf{x}[v_0v_1]_{\text{FG}(P)}^{\mathcal{X}}$.

Note that parallel ℓ -freeness implies parallel ℓ' -freeness for $\ell' \leq \ell$, as we can repeat components. Also note that, for fixed \mathcal{X} , the hierarchy of parallel ℓ -freeness collapses at least at level $\ell = |X|$ to the following property: for each $u \in P^*$ there is a $v \in P^*$ s.t. $v^{\mathcal{X}} = u^{\mathcal{S}}\|_X$. This condition is similar to 2-freeness but weaker as it is not guaranteed that $v^{\mathcal{S}} = u^{\mathcal{S}}$ (cf. the discussion after Definition 2.1.10).




Figure 2.1: A parallel ℓ -free solution \mathcal{S} . A tuple \mathbf{a} in the intersection $Xu_0^{\mathcal{S}} \cap Xu_1^{\mathcal{S}}$ has some corresponding tuple $\mathbf{c} \in Xv_0^{\mathcal{U}} \cap Xv_1^{\mathcal{U}}$ such that both induce the same partial bijection of X .

In the following we show that parallel ℓ -freeness describes the combinatorial properties of the automorphisms of the solution given by Herwig's Theorem.

A σ -*expansion* of a plain extension problem $\mathcal{X} = (X, (p^{\mathcal{X}}))$ is an extension problem of the form $\mathcal{X}' = (X, (R^{\mathcal{X}'})_{R \in \sigma}, (p^{\mathcal{X}})_{p \in P})$ (in particular the $p^{\mathcal{X}}$ are partial automorphisms of $(X, (R^{\mathcal{X}'})_{R \in \sigma})$). If \mathcal{S} is a solution of \mathcal{X} , we can 'lift' the relations $R^{\mathcal{X}'}$ to \mathcal{S} :

Definition 2.3.3. Let $\mathcal{S} = (S, (p^{\mathcal{S}})_{p \in P})$ be a solution of a plain extension problem \mathcal{X} and \mathcal{X}' a σ -expansion of \mathcal{X} . We define $\mathcal{S}' = \mathcal{S}(\mathcal{X}')$ as the triple $(S, (R^{\mathcal{S}'})_{R \in \sigma}, (p^{\mathcal{S}})_{p \in P})$ where $R^{\mathcal{S}'} := \{ \mathbf{x}u^{\mathcal{S}} \mid u \in P^*, \mathbf{x} \in R^{\mathcal{X}'} \}$.

By construction, the $p^{\mathcal{S}}$ are automorphisms of $\text{Str}(\mathcal{S}')$ and for each $\mathbf{a} \in R^{\mathcal{S}'}$ there are $\mathbf{x} \in R^{\mathcal{X}'}$ and $u \in P^*$ s.t. $\mathbf{a} = \mathbf{x}u^{\mathcal{S}}$. If \mathcal{U} is the free solution of \mathcal{X} , then $\mathcal{U}(\mathcal{X}')$ is the free solution of \mathcal{X}' . However, generally $\mathcal{S}(\mathcal{X}')$ is not a solution of \mathcal{X}' ; $\text{Str}(\mathcal{X}')$ is in general only a weak substructure of $\text{Str}(\mathcal{S}(\mathcal{X}'))$.

We show that for parallel ℓ -free solutions \mathcal{S} , $\mathcal{S}(\mathcal{X}')$ is a solution of \mathcal{X}' . (Recall that the width of a signature is the maximal arity of its relational symbols.)

Lemma 2.3.4. *Let \mathcal{S} be a solution of a plain extension problem \mathcal{X} and $\ell \in \mathbb{N}$. Then \mathcal{S} is parallel ℓ -acyclic if, and only if, $\mathcal{S}(\mathcal{X}')$ is a solution of every expansion \mathcal{X}' of \mathcal{X} over signatures of width at most ℓ .*

2.3 Structural properties of plain extension problems

Proof. Let \mathcal{S} be parallel ℓ -acyclic. Given a signature σ of width at most ℓ and a σ -expansion $\mathcal{X}' = (X, (R^{\mathcal{X}'})_{R \in \sigma}, (p^{\mathcal{X}'})_{p \in P})$ of \mathcal{X} , we show that $\mathcal{S}' = \mathcal{S}(\mathcal{X}')$ is a solution of \mathcal{X}' .

For that we have to show that $\text{Str}(\mathcal{S}')|_X = \text{Str}(\mathcal{X}')$. If $\mathbf{x} \subseteq X$ with $\mathbf{x} \in R^{\mathcal{S}'}$, then there are $\mathbf{y} \in R^{\mathcal{X}'}$ and $u \in P^*$ s.t. $\mathbf{x} = \mathbf{y}u^{\mathcal{S}'}$. By parallel ℓ -freeness there is a $v \in P^*$ s.t. $\mathbf{x} = \mathbf{y}v^{\mathcal{X}'}$, whence $\mathbf{x} \in R^{\mathcal{X}'}$.

For the converse we set $\sigma := \{R_{\mathbf{x}} \mid \mathbf{x} \in X^\ell\}$ and expand \mathcal{X} by the relations $R_{\mathbf{x}}^{\mathcal{X}'}$ where $R_{\mathbf{x}}^{\mathcal{X}'} := \{\mathbf{x}u^{\mathcal{X}'} \mid u \in P^*, \mathbf{x} \in \text{dom}(u^{\mathcal{X}'})\}$. By assumption $\mathcal{S}' = \mathcal{S}(\mathcal{X}')$ is a solution of \mathcal{X}' . If $\mathbf{x} \in \text{dom}(u^{\mathcal{S}'}|_X)$, then $\mathbf{x}u^{\mathcal{S}'} \in R_{\mathbf{x}}^{\mathcal{S}'}$ and so $\mathbf{x}u^{\mathcal{S}'} \in R_{\mathbf{x}}^{\mathcal{X}'}$. Thus there is a $v \in P^*$ with $\mathbf{x}v^{\mathcal{X}'} = \mathbf{x}u^{\mathcal{S}'}$. \square

So, if we prove that every finite plain extension problem has finite, parallel ℓ -free solutions of arbitrary high degree $\ell \in \mathbb{N}$, we automatically get a proof of Herwig's Theorem. Conversely, the proof of the previous lemma tells us how we can apply Herwig's Theorem to obtain finite parallel ℓ -acyclic solutions. In this sense parallel ℓ -freeness captures the strength of the solutions given by the Theorem of Herwig. Hrushovski's Theorem and parallel 2-freeness are also related in this manner.

We want to find a similar result for the Theorem of Herwig and Lascar. As a preparatory step we introduce a freeness property that captures the strength of the solution given by Theorem 2.1.5, which is the statement that CCF_ℓ has EPPA.

We say that a solution \mathcal{S} has the property $(**^\ell)$ if short 'cycles' that are introduced by intersections of patches are also reflected in the free solution \mathcal{U} . Formally $(**^\ell)$ is defined as:

$$\begin{aligned} &\text{for all } (a_i)_{i \in \mathbb{Z}_\ell} \subseteq S \text{ and } (u_i)_{i \in \mathbb{Z}_\ell} \subseteq P^* \text{ with } a_i \in Xu_{i-1}^{\mathcal{S}} \cap Xu_i^{\mathcal{S}} \text{ there are} \\ &(v_i)_{i \in \mathbb{Z}_\ell} \subseteq P^* \text{ and } c_i \in Xv_{i-1}^{\mathcal{U}} \cap Xv_i^{\mathcal{U}} \text{ s.t.} \\ &c_i \varphi_{v_{i-1}}^{\mathcal{U}} = a_i \varphi_{u_{i-1}}^{\mathcal{S}} \text{ and } c_i \varphi_{v_i}^{\mathcal{U}} = a_i \varphi_{u_i}^{\mathcal{S}} \end{aligned} \quad (**^\ell)$$

Figure 2.2 shows a sketch of $(**^3)$.

Similarly to $(**^\ell)$, we can characterise $(**^\ell)$ as some sort of freeness.

Definition 2.3.5. A solution \mathcal{S} of a P -extension problem \mathcal{X} is *sequential ℓ -free* if for all $(u_i)_{i \in \mathbb{Z}_\ell} \subseteq P^*$ and $(x_i)_{i \in \mathbb{Z}_\ell} \subseteq X$ with $x_i \in \text{dom}(u_i^{\mathcal{S}}|_X)$ and $(u_0 \dots u_{\ell-1})^{\mathcal{S}} = \text{id}$ there are $(v_i)_{i \in \mathbb{Z}_\ell} \subseteq P^*$ s.t. $x_i v_i^{\mathcal{X}} = x_i u_i^{\mathcal{S}}$ and $[v_0 \dots v_{\ell-1}]_{\text{FG}(P)} = 1$.

It is not hard to see that sequential ℓ -freeness is equivalent to $(**^\ell)$. We take this fact for granted.

Sequential ℓ' -freeness implies sequential ℓ -freeness for $\ell \leq \ell'$: given $(u_i)_{i \in \mathbb{Z}_\ell} \subseteq P^*$ and $(x_i)_{i \in \mathbb{Z}_\ell} \subseteq X$ s.t. $x_i \in \text{dom}(u_i^{\mathcal{S}}|_X)$ and $(u_0 \dots u_{\ell-1})^{\mathcal{S}} = \text{id}$, let $(u_i)_{i \in \mathbb{Z}_{\ell'}}$ and $(x'_i)_{i \in \mathbb{Z}_{\ell'}}$ be defined as

$$u'_i := \begin{cases} u_i, & \text{if } i \in \mathbb{Z}_\ell \\ \varepsilon, & \text{if } i \in \mathbb{Z}_{\ell'} \setminus \mathbb{Z}_\ell \end{cases} \quad x'_i := \begin{cases} x_i, & \text{if } i \in \mathbb{Z}_\ell \\ y, & \text{if } i \in \mathbb{Z}_{\ell'} \setminus \mathbb{Z}_\ell \end{cases}$$

2 Extension problems




Figure 2.2: A sequential 3-free solution \mathcal{S} . A closed walk $(a_i)_{i \in \mathbb{Z}_3}$ in $T(\mathcal{S})$ has a corresponding closed walk $(c_i)_{i \in \mathbb{Z}_3}$ in $T(\mathcal{U})$.

where $y = x_{\ell-1}u_{\ell-1}$. Then, if $(v'_i)_{i \in \mathbb{Z}_{\ell'}}$ is a witness in the sense of sequential ℓ' -freeness, $(v_i)_{i \in \mathbb{Z}_{\ell'}}$ with $v_i = v'_i$ ($i = 0, \dots, \ell - 2$) and $v_{\ell-1} = v'_{\ell-1} \dots v'_{\ell'-1}$ is a witness in the sense of sequential ℓ -freeness.

Interestingly, parallel 2-freeness and sequential 2-freeness are equivalent. We show that a parallel 2-free solution \mathcal{S} is also sequential 2-free: if $(u_0u_1)^{\mathcal{S}} = \text{id}$ with $x_0 \in \text{dom}(u_0^{\mathcal{S}}|_X)$ and $x_1 \in \text{dom}(u_1^{\mathcal{S}}|_X)$, then $(x_0, x'_1)u_0^{\mathcal{S}} = (x'_0, x_1)$ for $x'_i = x_iu_i^{\mathcal{S}}$. By parallel 2-freeness, there is a $v \in P^*$ s.t. $(x_0, x'_1)v^{\mathcal{X}} = (x'_0, x_1)$. Then $v_0 = v$ and $v_1 = v^{-1}$ are suitable witnesses, as $x_0v_0^{\mathcal{X}} = x_0v^{\mathcal{X}} = x'_0$, $x_1v_1^{\mathcal{X}} = x_1(v^{-1})^{\mathcal{X}} = x'_1$ and $[v_0v_1]_{\text{FG}(P)} = [vv^{-1}]_{\text{FG}(P)} = 1$. The converse can be shown similarly.

We show now that sequential ℓ -freeness captures the structure of the partial automorphisms of the solutions given by Theorem 2.1.5. Recall that CCF_{ℓ} is the class of σ_{ℓ} -structures \mathcal{A} over $\sigma_{\ell} = \{E_i \mid i \in \mathbb{Z}_{\ell}\}$ that have no coloured- ℓ -cycles, i.e., there are no sequences $(a_i)_{i \in \mathbb{Z}_{\ell}}$ s.t. $(a_i, a_{i+1}) \in E_i^{\mathcal{A}}$.

Lemma 2.3.6. *Let \mathcal{S} be a solution of a plain extension problem \mathcal{X} and $\ell \in \mathbb{N}$. Then the following are equivalent:*

- (i) \mathcal{S} is sequential ℓ -free,
- (ii) $\mathcal{S}(\mathcal{X}')$ is a solution in CCF_{ℓ} if the σ_{ℓ} -expansion \mathcal{X}' of \mathcal{X} has a solution in CCF_{ℓ} .

Proof. Let \mathcal{U} be the free solution of \mathcal{X} . We set $\mathcal{S}' = \mathcal{S}(\mathcal{X}')$ and $\mathcal{U}' = \mathcal{U}(\mathcal{X}')$. Note that \mathcal{U}' is the free solution of \mathcal{X}' and that \mathcal{X}' has a solution in CCF_{ℓ} , if and only if $\mathcal{U}' \in \text{CCF}_{\ell}$.

2.3 Structural properties of plain extension problems

(i) \implies (ii). Let \mathcal{X}' be a σ_ℓ -expansion of \mathcal{X} . Since sequential 2-freeness is equivalent to parallel 2-freeness we can apply Lemma 2.3.4. So \mathcal{S}' is a solution of \mathcal{X}' (but we do not know whether $\mathcal{S}' \in \text{CCF}_\ell$ yet).


We need to show that $\mathcal{S}' \notin \text{CCF}_\ell \implies \mathcal{U}' \notin \text{CCF}_\ell$. Let $(a_i)_{i \in \mathbb{Z}_\ell} \subseteq S$ be a coloured- ℓ -cycle. Since $(a_i, a_{i+1}) \in R_i^{\mathcal{S}'}$ there are $(x_i, y_{i+1}) \in R_i^{\mathcal{X}'}$ and $u_i \in P^*$ s.t. $(x_i, y_{i+1})u_i^{\mathcal{S}'} = (a_i, a_{i+1})$. So $a_i = x_i u_i^{\mathcal{S}'} = y_i u_{i-1}^{\mathcal{S}'}$ and thus $a_i \in X u_{i-1}^{\mathcal{S}'} \cap X u_i^{\mathcal{S}'}$. Then, by $(**^\ell)$, there are $v_i \in P^*$ and $c_i \in U$ s.t. $c_i = x_i v_i^{\mathcal{U}'} = y_i v_{i-1}^{\mathcal{U}'}$. So $(x_i, y_{i+1})v_i^{\mathcal{S}'} = (c_i, c_{i+1})$ and thus $(c_i)_{i \in \mathbb{Z}_\ell}$ is a coloured- ℓ -cycle in \mathcal{U}' .

(ii) \implies (i). Let $(u_i)_{i \in \mathbb{Z}_\ell} \subseteq P^*$ and $(x_i)_{i \in \mathbb{Z}_\ell} \subseteq X$ be given s.t. $x \in \text{dom}(u_i^{\mathcal{S}'}|_X)$ and $(u_0 \dots u_{\ell-1})^{\mathcal{S}'} = \text{id}$.


Set $y_i = x_i u_i^{\mathcal{S}'}$. Let \mathcal{X}' be the σ_ℓ -expansion of \mathcal{X} by the following relations $E_i^{\mathcal{X}'} := \{(x_i, y_{i+1})u^{\mathcal{X}'} \mid x_i, y_{i+1} \in \text{dom}(u^{\mathcal{S}'})\}$. Then $\mathcal{S}' \notin \text{CCF}_\ell$: let $u'_0 = \varepsilon$ and $u'_{i+1} = u'_i u_{i+1}^{-1}$. Let $(a_i)_{i \in \mathbb{Z}_\ell} \subseteq S$ s.t. $a_i (u'_i)^{\mathcal{S}'} = x_i$. Then $(a_i, a_{i+1})(u'_i)^{\mathcal{S}'} = (x_i, y_{i+1})$ and thus $(a_i)_{i \in \mathbb{Z}_\ell}$ is a coloured- ℓ -cycle. Since $\mathcal{S}' \notin \text{CCF}_\ell$ also $\mathcal{U}' \notin \text{CCF}_\ell$. Let $(c_i)_{i \in \mathbb{Z}_\ell} \subseteq U$ be a coloured- ℓ -cycle in \mathcal{U}' . Then there are $(v'_i)_{i \in \mathbb{Z}_\ell} \subseteq P^*$ s.t. $(c_i, c_{i+1})(v'_i)^{\mathcal{S}'} = (x_i, y_{i+1})$. Put $v_i = \text{red}((v'_i)^{-1} v'_{i-1})$. Then $[v_0 \dots v_{\ell-1}]_{\text{FG}(P)} = 0$ and also $x_i v_i^{\mathcal{X}'} = x_i u_i^{\mathcal{S}'}$ since $x_i v_i^{\mathcal{U}'} = y_i = x_i u_i^{\mathcal{S}'}$. \square

Now we introduce ‘cluster ℓ -freeness’ a property that generalises both parallel ℓ -freeness and sequential ℓ -freeness, and that captures the strength of the Theorem of Herwig and Lascar.

We can express sequential freeness as the following relabelling property for cycle graphs: A solution \mathcal{S} of \mathcal{X} is sequential ℓ -free if, and only if, for every cycle graph $I = (V, E)$ of size ℓ that has an edge-labelling as in the following sketch




that satisfy $x_i \in X$, $u_i \in P^*$, $x_i u_i^{\mathcal{S}} \in X$ and $(u_0 \dots u_{\ell-1})^{\mathcal{S}} = \text{id}$ there is a relabelling




s.t. $v_i \in P^*$, $x_i v_i^{\mathcal{X}} = x_i u_i^{\mathcal{S}}$ and $[v_0 \dots v_{\ell-1}]_{\text{FG}(P)} = 1$.

Similarly we can express parallel ℓ -freeness as a relabelling property. A solution \mathcal{S} of \mathcal{X} is parallel ℓ -free if, and only if, for labelled graphs as in the following sketch



2 Extension problems

that satisfy $x_i \in X$, $u_i \in P^*$, $x_i u_i^{\mathcal{S}} \in X$ and $(u_i u_j^{-1})^{\mathcal{S}} = \text{id}$ for all i, j there is a relabelling



s.t. $v_i \in P^*$, $x_i v_i^{\mathcal{X}} = x_i u_i^{\mathcal{S}}$ and $[v_i v_j^{-1}]_{\text{FG}(P)} = 1$ (so basically $v_i = v_j$).

Cluster freeness generalises this relabelling idea to arbitrary graphs.

Definition 2.3.7. An \mathcal{X} -cluster of a P -extension problem is a multidigraph $I = (V, E)$ with edge-labellings $\eta: E \rightarrow X$ and $\mu: E \rightarrow P^*$ s.t. $\mu(e^{-1}) = \mu(e)^{-1}$.

A cluster (I, η, μ) is *compatible* with a solution \mathcal{S} of \mathcal{X} if

$$\eta(e)\mu(e)^{\mathcal{S}} = \eta(e^{-1}) \quad \text{for } e \in E, \quad \text{and } \mu(u)^{\mathcal{S}} = \text{id} \quad \text{for closed walks } u \text{ in } I.$$

A *relabelling* of an \mathcal{X} -cluster (I, η, μ) is an \mathcal{X} -cluster (I, η, μ') for $\mu': E \rightarrow P^*$. An \mathcal{X} -cluster is *free* if $\eta(e)\mu(e)^{\mathcal{X}} = \eta(e^{-1})$ for $e \in E$, and $[\mu(u)]_{\text{FG}(P)} = 1$ for closed walks u in G .

Note that if (I, η, μ) is compatible with \mathcal{S} then $\eta(e)\mu(e)^{\mathcal{S}} \in X$ as $\eta(e^{-1}) \in X$.

Definition 2.3.8. A solution \mathcal{S} of a P -extension problem \mathcal{X} is *cluster ℓ -free* if every \mathcal{X} -cluster (I, η, μ) with at most 2ℓ edges that is compatible with \mathcal{S} has a relabelling (I, η, μ') that is free. ⁵

As discussed above, we obtain parallel ℓ -freeness as a special case of cluster ℓ -freeness when restricting the clusters to graphs with two vertices and ‘parallel’ edges between them. Similarly, we obtain sequential ℓ -freeness by restricting the form of the possible clusters to cycles.

We introduce a property $(***)^{\ell}$ that gives an alternative description of cluster ℓ -freeness (though, a change of the parameter ℓ is required). A solution \mathcal{S} has the property $(***)^{\ell}$ if:

$$\begin{aligned} &\text{for all } (u_i)_{i \in \mathbb{Z}_{\ell}} \subseteq P^* \text{ there are } (v_i)_{i \in \mathbb{Z}_{\ell}} \subseteq P^* \text{ s.t.} \\ &f: \bigcup X u_i^{\mathcal{S}} \rightarrow \bigcup X v_i^{\mathcal{U}}; x u_i^{\mathcal{S}} \mapsto x v_i^{\mathcal{U}} \text{ is well-defined,} \end{aligned} \quad (***)^{\ell}$$

where \mathcal{U} is the free solution.

Lemma 2.3.9. *Let \mathcal{X} be fixed. Then cluster $|X|^{\ell^2}$ -freeness implies $(***)^{\ell}$, and $(***)^{\ell}$ implies cluster ℓ -freeness.*

⁵We choose 2ℓ here since in the definition of multidigraph every edge e comes automatically with its reverse e^{-1} .

2.3 Structural properties of plain extension problems

Proof. We show the first implication. Given $(u_i)_{i \in \mathbb{Z}_\ell} \subseteq P^*$, we construct an \mathcal{X} -cluster (I, η, μ) that is compatible with \mathcal{S} that has the $(u_i)_{i \in \mathbb{Z}_\ell} \subseteq P^*$ as vertices. Two vertices u_i and u_j are connected by an edge e labelled with $\eta(e) = x$ and $\mu(e) = u_i u_j^{-1}$ if $x \in \text{dom}((u_i u_j^{-1})^{\mathcal{S}} \parallel_X)$. This cluster has at most $|X| \ell^2$ many edges and so it has a free relabelling (I, η, μ') . We set $v_0 = \varepsilon$ and $v_i = [\mu'(u)]_{\text{FG}(P)}$, where u is a connecting walk from u_0 to u_i . The v_i are well-defined as

$$(I, \eta, \mu')$$

is free and we have $xu_i^{\mathcal{S}} = yu_j^{\mathcal{S}} \implies x(u_i u_j^{-1})^{\mathcal{S}} = y \implies x(v_i v_j)^{\mathcal{X}} = y \implies xu_i^{\mathcal{U}} = yu_j^{\mathcal{U}}$.

The converse direction is done by a similar argument. Suppose we are given an \mathcal{X} -cluster (I, η, μ) that is compatible with \mathcal{S} , label the vertices $\lambda: V \rightarrow P^*$ of $I = (V, E)$ s.t. $\mu(e)^{\mathcal{S}} = (\lambda(s(e))\lambda(t(e))^{-1})^{\mathcal{S}}$. Then if we apply $(***)^\ell$ to the set $\{\lambda(v) \mid v \in V\}$ we can use the witnesses for a free relabelling (I, η, μ') . \square

Property $(***)^\ell$ is a consequence of ℓ -freeness as shown in Lemma 2.1.12. Hence, ℓ -freeness implies cluster ℓ -freeness. Also note that in the argument that the Free Extension Conjecture implies the Theorem of Herwig and Lascar (the discussion after Lemma 2.1.12) we actually only used that ℓ -free solutions have the property $(***)^\ell$. This basically provides the direction $(i) \implies (ii)$ of the following lemma.

Lemma 2.3.10. *Let \mathcal{S} be a solution of a plain extension problem \mathcal{X} and $\ell \in \mathbb{N}$. Then the following are equivalent*

- (i) \mathcal{S} has the property $(***)^\ell$,
- (ii) if a σ -expansion \mathcal{X}' of \mathcal{X} has an \mathcal{A} -free solution and $\sum_{R \in \sigma} |R^{\mathcal{A}}| \leq \ell$ then $\mathcal{S}(\mathcal{X}')$ is also \mathcal{A} -free.

Proof. Let \mathcal{U} be the free solution of \mathcal{X} and put $\mathcal{U}' = \mathcal{U}(\mathcal{X}')$ and $\mathcal{S}' = \mathcal{S}(\mathcal{X}')$. Note that \mathcal{X}' has an \mathcal{A} -free solution, if and only if \mathcal{U}' is \mathcal{A} -free.

$(i) \implies (ii)$. We show that if \mathcal{S}' is not \mathcal{A} -free, then \mathcal{U}' is not \mathcal{A} -free. Let $f: \mathcal{A} \xrightarrow{\text{hom}} \text{Str}(\mathcal{S}')$ be a homomorphism. For each $R \in \Sigma$ and $\mathbf{a} \in R^{\mathcal{A}}$ we have $f(\mathbf{a}) \in R^{\mathcal{S}'}$. Thus there are $u_{R, \mathbf{a}} \in P^*$ and $\mathbf{x} \in R^{\mathcal{X}'}$ s.t. $\mathbf{x}u_{R, \mathbf{a}}^{\mathcal{S}} = f(\mathbf{a})$. Let $v_{R, \mathbf{a}}$ be according to the property $(***)^\ell$ s.t. the map $g: \bigcup X u_{R, \mathbf{a}}^{\mathcal{S}} \rightarrow \bigcup X v_{R, \mathbf{a}}^{\mathcal{U}'}; xu_{R, \mathbf{a}}^{\mathcal{S}} \mapsto xv_{R, \mathbf{a}}^{\mathcal{U}'}$ is well-defined.

Then $g \circ f: \mathcal{A} \xrightarrow{\text{hom}} \text{Str}(\mathcal{U}')$: if $\mathcal{A} \in R^{\mathcal{A}}$ then $g \circ f(\mathcal{A}) = \mathbf{x}_{R, \mathcal{A}} v_{R, \mathcal{A}}^{\mathcal{U}'} \in R^{\mathcal{U}'}$.

$(ii) \implies (i)$. Suppose we are given $(u_i)_{i \in \mathbb{Z}_\ell} \subseteq P^*$. Let \mathbf{x} be an enumeration of X and set $R^{\mathcal{X}'} = \{\mathbf{x}u^{\mathcal{X}'} \mid u \in P^* \text{ s.t. } \mathbf{x} \in \text{dom}(u^{\mathcal{X}'})\}$. Set $\mathcal{A} = (\bigcup_{i \in \mathbb{Z}_\ell} X u_i^{\mathcal{S}}, R^{\mathcal{A}} = \{\mathbf{x}u_i^{\mathcal{S}} \mid i \in \mathbb{Z}_\ell\})$. Then $\mathcal{S}' = \mathcal{S}(\mathcal{X}')$ is not \mathcal{A} -free. So the free solution \mathcal{U}' is not \mathcal{A} -free. Let $f: \mathcal{A} \xrightarrow{\text{hom}} \text{Str}(\mathcal{U}')$ be a homomorphism. Then for $f(xu_i^{\mathcal{S}})$ there is an $v_i \in P^*$ s.t. $\mathbf{x}v_i^{\mathcal{U}'} = f(xu_i^{\mathcal{S}})$. So, the map $\bigcup X u_i^{\mathcal{S}} \rightarrow \bigcup X v_i^{\mathcal{U}'}; xu_i^{\mathcal{S}} \mapsto xv_i^{\mathcal{U}'}$ is well-defined. \square

We finish this section by discussing ℓ -freeness some more. We recall: a solution \mathcal{S} is ℓ -free if

$$\begin{aligned} &\text{for all } (u_i)_{i \in \mathbb{Z}_\ell} \subseteq P^* \text{ with } (u_0 \dots u_{\ell-1})^{\mathcal{S}} = \text{id} \text{ there are } (v_i)_{i \in \mathbb{Z}_\ell} \subseteq P^* \text{ s.t.} \\ &v_i^{\mathcal{S}} = u_i^{\mathcal{S}}, v_i^{\mathcal{X}} = u_i^{\mathcal{S}} \parallel_X \text{ for } i \in \mathbb{Z}_\ell, \text{ and } [v_0 \dots v_{\ell-1}]_{\text{FG}(P)} = 1. \end{aligned}$$

2 Extension problems




Figure 2.3: Sketch for ℓ -acyclicity: the u_i form a ‘cycle’ and v_- and v_+ act as a ‘chord’ over u_j .

We introduce an acyclicity property that is equivalent to ℓ -freeness.

Definition 2.3.11. A solution \mathcal{S} of a P -extension problem \mathcal{X} is ℓ -acyclic if for all $(u_i)_{i \in \mathbb{Z}_\ell} \subseteq P^*$ with $(u_0 \dots u_{\ell-1})^{\mathcal{S}} = \text{id}$ there are $j \in \mathbb{Z}_\ell$ and $v_-, v_+ \in P^*$ s.t.

$$\begin{aligned} u_j^{\mathcal{S}} &= (v_- v_+)^{\mathcal{S}}, \text{img}(u_{j-1}^{\mathcal{S}} \|_X), \text{dom}(u_j^{\mathcal{S}} \|_X) \subseteq \text{dom}(v_-^{\mathcal{X}}) \\ &\text{and } \text{img}(u_j^{\mathcal{S}} \|_X), \text{dom}(u_{j+1}^{\mathcal{S}} \|_X) \subseteq \text{img}(v_+^{\mathcal{X}}). \end{aligned}$$

Figure 2.3 gives a sketch of ℓ -acyclicity. Note that for v_- and v_+ we have that $u_j^{\mathcal{S}} \|_X = (v_- v_+)^{\mathcal{X}}$: if $x \in \text{dom}(u_j^{\mathcal{S}} \|_X)$ then $x \in \text{dom}(v_-^{\mathcal{X}})$ and $x u_j^{\mathcal{S}} \in \text{img}(v_+^{\mathcal{X}})$. Also $x v_-^{\mathcal{X}} = x v_j^{\mathcal{S}} (v_+^{-1})^{\mathcal{X}}$ since $v_-^{\mathcal{S}} = u_j^{\mathcal{S}} (v_+^{-1})^{\mathcal{S}}$. Thus $x u_j^{\mathcal{S}} = x v_-^{\mathcal{X}} v_+^{\mathcal{X}}$.

Lemma 2.3.12. A solution is ℓ -free if, and only if, it is ℓ -acyclic.

Proof. From left to right. For $(u_i)_{i \in \mathbb{Z}_\ell} \subseteq P^*$ with $(u_0 \dots u_{\ell-1})^{\mathcal{S}} = \text{id}$ let $(v_i)_{i \in \mathbb{Z}_\ell} \in \text{FG}(P)$ be witnesses according to ℓ -freeness. Then we can decompose one of the v_j into v_-, v_+ according to Lemma 0.2.5 so that $v_{j-1} = w_- v_-^{-1}$ and $v_{j+1} = v_+^{-1} w_+$. Then

$$\begin{aligned} \text{img}(u_{j-1}^{\mathcal{S}} \|_X) &= \text{img}(v_{j-1}^{\mathcal{X}}) = \text{img}((w_- v_-^{-1})^{\mathcal{X}}) = \text{dom}((v_- w_-^{-1})^{\mathcal{X}}) \subseteq \text{dom}(v_-^{\mathcal{X}}) \\ \text{dom}(u_j^{\mathcal{S}} \|_X) &= \text{dom}(v_j^{\mathcal{X}}) = \text{dom}((v_- v_+)^{\mathcal{X}}) \subseteq \text{dom}(v_+^{\mathcal{X}}), \end{aligned}$$

and similarly $\text{img}(u_j^{\mathcal{S}} \|_X), \text{dom}(u_{j+1}^{\mathcal{S}} \|_X) \subseteq \text{img}(v_+^{\mathcal{X}})$.

From right to left. We prove the implication by induction on ℓ . We leave the case $\ell = 2$ to the reader. For the induction step let $(u_i)_{i \in \mathbb{Z}_{\ell+1}}$ with $(u_0 \dots u_\ell)^{\mathcal{S}} = \text{id}$. By ℓ -acyclicity there are $j \in \mathbb{Z}_{\ell+1}$ and $v_-, v_+ \in P^*$ s.t. $u_j^{\mathcal{S}} = (v_- v_+)^{\mathcal{S}}$ and

$$\begin{aligned} \text{img}(u_{j-1}^{\mathcal{S}} \|_X), \text{dom}(u_j^{\mathcal{S}} \|_X) &\subseteq \text{dom}(v_-^{\mathcal{X}}), \\ \text{img}(u_j^{\mathcal{S}} \|_X), \text{dom}(u_{j+1}^{\mathcal{S}} \|_X) &\subseteq \text{img}(v_+^{\mathcal{X}}). \end{aligned}$$

2.3 Structural properties of plain extension problems

We can apply the induction hypothesis on $u_0, \dots, u_{j-2}, (u_{j-1}v_-), (v_+u_{j+1}), u_{j+2}, \dots, u_\ell$ and obtain $v_0, \dots, v_{j-1}, v_{j+1}, \dots, v_\ell$ s.t. $[v_0 \dots v_{j-1} v_{j+1} \dots v_\ell]_{\text{FG}} = 1$ and in the following diagram we have $y^{\mathcal{S}} = x^{\mathcal{S}}, y^{\mathcal{X}} = x^{\mathcal{S}} \parallel_X$ for all columns (x, y) :

$$\begin{array}{cccccccc} u_0 & \dots & u_{j-2} & u_{j-1}v_- & (v_+u_{j+1}) & u_{j+2} & \dots & u_{\ell-1} \\ v_0 & \dots & v_{j-2} & v_{j-1} & v_{j+1} & v_{j+2} & \dots & v_\ell. \end{array}$$

Then $[v_0 \dots v_{j-1} (v_{j-1}v_-^{-1})(v_-v_+)(v_+^{-1}v_{j+1}) \dots v_\ell]_{\text{FG}(P)} = 1$ and we claim that we also have $y^{\mathcal{S}} = x^{\mathcal{S}}, y^{\mathcal{X}} = x^{\mathcal{S}} \parallel_X$ for all columns (x, y) in the following diagram:

$$\begin{array}{cccccccc} u_0 & \dots & u_{j-2} & u_{j-1} & u_j & u_{j+1} & u_{j+2} & \dots & u_{\ell-1} \\ v_0 & \dots & v_{j-2} & (v_{j-1}v_-^{-1}) & v_-v_+ & (v_+^{-1}v_{j+1}) & v_{j+2} & \dots & v_\ell. \end{array}$$

This is obvious for all pairs except $(u_{j-1}, v_{j-1}v_-^{-1}), (u_j, v_-v_+), (u_{j+1}, v_+^{-1}v_{j+1})$.

By the origin of v_- and v_+ we have $(v_-v_+)^{\mathcal{S}} = u_j^{\mathcal{S}}$, and in the remark after the definition of ℓ -acyclicity we showed $(v_-v_+)^{\mathcal{X}} = u_j^{\mathcal{S}} \parallel_X = (v_-v_+)^{\mathcal{X}}$.

We have $(v_{j-1}v_-^{-1})^{\mathcal{S}} = ((u_{j-1}v_-)v_-^{-1})^{\mathcal{S}} = u_{j-1}^{\mathcal{S}}$. For $(v_{j-1}v_-^{-1})^{\mathcal{X}} = u_{j-1}^{\mathcal{S}} \parallel_X$ we observe that

$$(v_{j-1}v_-^{-1})^{\mathcal{X}} = v_{j-1}^{\mathcal{X}}(v_-^{-1})^{\mathcal{X}} = (u_{j-1}v_-)^{\mathcal{S}} \parallel_X v_-^{\mathcal{X}} = u_{j-1}^{\mathcal{S}} \parallel_X,$$

where the second equality uses that $v_{j-1}^{\mathcal{X}} = (u_{j-1}v_-)^{\mathcal{S}} \parallel_X$ and the last equality uses that $\text{img}(u_{j-1})^{\mathcal{S}} \parallel_X \subseteq \text{dom}(v_-)^{\mathcal{X}}$. Similarly one shows that $(v_+^{-1}v_{j+1})^{\mathcal{X}} = u_{j+1}^{\mathcal{S}} \parallel_X$. \square

We can now prove that ℓ -acyclicity (and hence ℓ -freeness) imposes acyclicity restrictions on the translation hypergraph.

Lemma 2.3.13. *If \mathcal{S} is an ℓ -acyclic solution of \mathcal{X} for $\ell \geq 3$, then $T(\mathcal{S})$ is an ℓ -acyclic hypergraph.*

Proof. Recall that a hypergraph is ℓ -acyclic, if and only if it does not have short hypertours (Theorem 1.4.10). Let $(Xw_i^{\mathcal{S}})_{i \in \mathbb{Z}_\ell}$ be a candidate for a short hypertour. Then $(u_0 \dots u_{\ell-1})^{\mathcal{S}} = \text{id}$ for $u_i := w_i w_{i+1}^{-1}$. By ℓ -acyclicity we obtain a $j \in \mathbb{Z}_\ell$ and $v_-, v_+ \in P^*$ s.t.

$$\begin{aligned} u_j^{\mathcal{S}} &= v_-^{\mathcal{S}} v_+^{\mathcal{S}}, & \text{img}(u_{j-1}^{\mathcal{S}} \parallel_X), \text{dom}(u_j^{\mathcal{S}} \parallel_X) &\subseteq \text{dom}(v_-^{\mathcal{S}} \parallel_X), \text{ and} \\ \text{img}(u_j^{\mathcal{S}} \parallel_X), \text{dom}(u_{j+1}^{\mathcal{S}} \parallel_X) &\subseteq \text{img}(v_+^{\mathcal{S}} \parallel_X). \end{aligned}$$

We claim that for $w = v_-^{-1}w_j$ we have

$$Xw_{j-1} \cap Xw_j, Xw_j \cap Xw_{j+1}, Xw_{j+1} \cap Xw_{j+2} \subseteq Xw.$$

- $a \in Xw_{j-1}^{\mathcal{S}} \cap Xw_j^{\mathcal{S}}$: there are $x, y \in X$ s.t. $xw_{j-1}^{\mathcal{S}} = a = yw_j^{\mathcal{S}}$. Thus $y \in \text{img}(u_{j-1}^{\mathcal{S}} \parallel_X)$ and hence there is a $z \in X$ s.t. $yv_-^{\mathcal{S}} = z$. So we get $a = yw_j^{\mathcal{S}} = z(v_-^{-1}w_j)^{\mathcal{S}} \in Xw^{\mathcal{S}}$.

2 Extension problems

- $a \in Xw_j^S \cap Xw_{j+1}^S$: there are $x, y \in X$ s.t. $xw_j^S = a = yw_{j+1}^S$. Thus $x \in \text{dom}(u_j^S \parallel_X)$ and hence there is a $z \in X$ s.t. $xv_-^S = z$. So we get $a = xw_j^S = z(v_-^{-1}w_j)^S \in Xw^S$.
- $a \in Xw_{j+1}^S \cap Xw_{j+2}^S$: there are $x, y \in X$ s.t. $xw_{j+1}^S = a = yw_{j+2}^S$. Thus $x \in \text{dom}(u_{j+1}^S \parallel_X)$ and hence there is a $z \in X$ s.t. $zv_+^S = x$. So $a = xw_{j+1}^S = z(v_+w_{j-1})^S = z(v_-^{-1}w_j)^S \in Xw^S$ (note that $(v_+w_{j+1})^S = (v_-^{-1}w_j)^S$).

So Xw^S is a witness that $(Xw_i^S)_{i \in \mathbb{Z}_\ell}$ is not a short hypertour. \square

2.4 Inverse monoids as ‘abstract extension problems’

In the previous section we discuss various freeness properties and show how they are connected to classical results about extension problems. We can summarize these results in the following overview. On the left we have the classical results and the Free Extension Conjecture ordered w.r.t. their strength (on the top the weakest, on the bottom the strongest). On the right we have the corresponding structural properties that characterise the behaviour of the automorphisms of the solutions given by these theorems.

Classical results	Structural properties
Hrushovski	parallel 2-freeness / sequential 2-freeness
Herwig CCF $_\ell$ has EPPA	parallel ℓ -free. sequential ℓ -free.
Herwig–Lascar	cluster ℓ -freeness $\approx (***)^\ell$
Free Extension Conjecture	ℓ -freeness

The corresponding lemmas showing these connections are:

- Lemma 2.3.4: connects Herwig’s Theorem and parallel ℓ -freeness and Hrushovski’s Theorem and parallel 2-freeness.
- Lemma 2.3.6: connects ‘CCF $_\ell$ has EPPA’ (Theorem 2.1.5) and sequential ℓ -freeness.
- Lemma 2.3.10: connects the Theorem of Herwig and Lascar and $(***)^\ell$, which corresponds to cluster ℓ -freeness (Lemma 2.3.9).

- By definition the Free Extension Conjecture (Conjecture 2.1.11) guarantees the existence of finite ℓ -free solutions.

Our goal is it to prove the existence of finite, approximately free solutions. By ‘approximate freeness’ we mean the various freeness properties of the overview. In case of the Free Extension Conjecture we want to reduce it to a group-theoretic result.

We proceed in two steps:

In the first step, which we present in this section, we translate the problem of finding approximately free solutions to an algebraic setting where the role of plain extension problems and solutions of plain extension problems is taken by inverse monoids and groups. This translation is such that approximately free solutions to these ‘abstract extension problems’ can be used to construct approximately free solutions of plain extension problems.

In the second step, which is done in Section 2.5 we then show the existence of finite, approximately free solutions in this algebraic setting.

2.4.1 Inverse monoids

Plain extension problems are completely characterised by how their partial bijections transform under composition. The so-called Wagner-Preston Representation Theorem says that the theory of partial bijections is described by the theory of inverse monoids. We show that inverse monoids give an abstract perspective on plain extension problems.

We give a short introduction to inverse monoids which encompasses all notions and results that are needed for our purposes. For proofs and further details we refer to the monograph by Lawson [26].

Definition 2.4.1. An *inverse monoid* M is a monoid s.t. for every $f \in M$ there is a unique $g \in M$ satisfying $f = fgf$ and $g = gfg$. This g is called *the inverse of f* and is denoted by f^{-1} .

For a set X , we write $I(X)$ for the *symmetric inverse monoid on X* , i.e., the inverse monoid of all partial bijections of X . We check that the symmetric inverse monoid fulfils the axiom of an inverse monoid: in $I(X)$ we have $fgf = f \iff f^{-1} \subseteq g$ (here f^{-1} is the set-theoretical inverse of f) and thus

$$fgf = f, gfg = g \iff f^{-1} \subseteq g, g^{-1} \subseteq f \iff f^{-1} \subseteq g, g \subseteq f^{-1} \iff f^{-1} = g.$$

The Wagner-Preston Representation Theorem tells us that these are basically all examples of inverse monoids.

Theorem 2.4.2 (Wagner-Preston). *Every inverse monoid is an inverse submonoid of a symmetric inverse monoid. This symmetric inverse monoid can be chosen to be finite in case the inverse monoid is finite.*

2 Extension problems

Using this theorem we can see that every inverse monoid M is equipped with a natural partial order \leq . M is an inverse submonoid of some symmetric inverse monoid whose \subseteq -relation induces a partial order \leq on M . So $m \leq n$ expresses that m is a ‘restriction’ of n . This partial order \leq can also be characterised as follows:

$$m \leq n \iff m = ne \text{ for some idempotent } e \in M.$$

Idempotents play an important role in the theory of inverse monoids. We write $E(M)$ for the set of idempotents of M . We note that $m \leq 1 \iff m \in E(M)$. In particular, in $I(X)$ the idempotents are the restrictions of the identity.

For every $m \in M$, mm^{-1} and $m^{-1}m$ are idempotents. We set $\text{dom}(m) = mm^{-1}$ and $\text{img}(m) = m^{-1}m$. This notation is compatible with the set-theoretic definition of dom and img in the sense that the inverse monoid theoretic “ $\text{dom}(f)$ and $\text{img}(f)$ ” are the characteristic functions of the set theoretic “ $\text{dom}(f)$ and $\text{img}(f)$ ”. Using the Wagner-Preston Representation Theorem we can see that if $\text{dom}(m) \leq \text{dom}(n)$, then $nn^{-1}m = m$, and if $\text{dom}(m) \leq \text{img}(n)$, then $n^{-1}nm = m$.

Inverse monoids with generators

A P -generated inverse monoid M is an inverse monoid with a family $(m_p)_{p \in P}$ of generators s.t. $m_p^{-1} = m_{p^{-1}}$ (as always, P is a set with an associated involution $(\cdot)^{-1}$, see Section 0.2). We can associate to every plain P -extension problem a P -generated inverse monoid.

Definition 2.4.3. For a plain P -extension problem $\mathcal{X} = (X, (p^x)_{p \in P})$, $I(\mathcal{X})$ is the P -generated inverse submonoid of $I(X)$ with the generator family $(p^x)_{p \in P}$.

By the Wagner-Preston Representation Theorem we can switch between P -generated inverse monoids and plain P -extension problems. Thus we may think of P -generated inverse monoids as ‘abstract extension problems’. If we talk about P -generated inverse monoids without specifying the underlying set P , we call them *abstract extension problems*.

As for P -generated groups, we require a homomorphism between P -generated inverse monoids N and M to map the generators according to their indices $p \in P$. So, there is at most one homomorphism from N to M and if it exists it is given by

$$[u]_N \mapsto [u]_M.$$

We write $M \leq N$ if this homomorphism exists, i.e., the mapping above is well-defined. We give the formal definition of $M \leq N$ as follows.

Definition 2.4.4. For two P -generated inverse monoids we write $M \leq N$ if

$$[u]_N = [v]_N \implies [u]_M = [v]_M$$

for all $u, v \in P^*$.

To establish $M \leq N$ we can also check a seemingly weaker condition.

Lemma 2.4.5. *Two P -generated inverse monoids satisfy $M \leq N$ if, and only if,*

$$[u]_N \leq [v]_N \implies [u]_M \leq [v]_M$$

for all $u, v \in P^*$.

Proof. We use the following result for inverse monoids: $m \leq n$, if and only if $m = mm^{-1}n$.

Let $M \leq N$. If $[u]_N \leq [v]_N$, then $[u]_N = [uu^{-1}v]_N$. Hence $[u]_M = [uu^{-1}v]_M$ and so $[u]_M \leq [v]_M$.

Conversely, let $[u]_N \leq [v]_N \implies [u]_M \leq [v]_M$ for all $u, v \in P^*$. If $[u]_N = [v]_N$, then $[u]_N \leq [v]_N$ and $[u]_N \geq [v]_N$. Hence $[u]_M \leq [v]_M$ and $[u]_M \geq [v]_M$. And so $[u]_M = [v]_M$. \square

Solutions of inverse monoids

Basically, a solution of a plain extension problem consists of a set with a collection of permutations. By Cayley’s Theorem the algebraic theory of permutations is described by group theory. So it is natural that we consider groups as the right format of solutions of abstract extension problems.

Definition 2.4.6. A *solution* of a P -generated inverse monoid is a P -generated group s.t. for all $u \in P^*$

$$[u]_G = 1 \implies [u]_M \leq 1.$$

Lemma 2.4.7. *Every P -generated inverse monoid M has a solution. This solution can be chosen to be finite in case M is finite.*

Proof. Let M be a P -generated inverse monoid. By the Wagner–Preston Theorem we can assume that $M = I(\mathcal{X})$ for some plain P -extension problem \mathcal{X} . \mathcal{X} has a solution \mathcal{S} (cf. Lemma 2.3.1). Then the P -generated group given by the family $(p^{\mathcal{S}})_{p \in P}$ is a solution of M . \square

The product construction

Given a plain P -extension problem \mathcal{X} , we can use solutions of $I(\mathcal{X})$ to obtain solutions of \mathcal{X} via the following product construction.

Let G be a P -generated group. We define $\mathcal{X} \otimes G$ as the tuple $(S, (p^{\mathcal{S}})_{p \in P})$ where the set S and the family $(p^{\mathcal{S}})_{p \in P}$ of permutations are given as follows:

1. $S := (X \times P^*)/\sim$ where

$$(x, u) \sim (y, w) :\iff \exists v \in P^* \text{ s.t. } xv^{\mathcal{X}} = y \text{ and } [v]_G = [uw^{-1}]_G$$

2. $[x, u]p^{\mathcal{S}} := [x, up]$ where $[x, u]$ is the equivalence class of (x, v) .

2 Extension problems

In this product p^S ‘extends’ p^X in the sense that $xp^X = y \implies [x, \varepsilon]p^S = [xp^X, \varepsilon]$. So if $\iota: X \rightarrow S; x \mapsto [x, \varepsilon]$ is an embedding of X in S , then $\mathcal{X} \otimes G$ is a solution of \mathcal{X} .

Lemma 2.4.8. $\mathcal{X} \otimes G$ is a solution of \mathcal{X} via the map $\iota: X \rightarrow S; x \mapsto [x, \varepsilon]$, if and only if G is a solution of $I(\mathcal{X})$.

Proof. We have to prove that ι is injective, if and only if G is a solution of $I(\mathcal{X})$.

Assume that G is a solution of $I(\mathcal{X})$. If $x\iota = [x, \varepsilon] = [y, \varepsilon] = y\iota$ for $x, y \in X$, then there is a $u \in P^*$ s.t. $xu^X = y$ and $[u]_G = 1$. Since G is a solution of $I(\mathcal{X})$, $u^X \subseteq \text{id}$ and thus $x = y$.

Assume that ι is injective. Given $u \in P^*$ with $[u]_G = 1$, then $x\iota = [x, \varepsilon] = [xu^X, \varepsilon] = (xu^X)\iota$ for $x \in \text{dom}(u^X)$. So, by injectivity of ι , $x = xu^X$. This shows us that $u^X \subseteq \text{id}$. \square

Solutions constructed by the product construction are always tidy. They also satisfy the following property.

Lemma 2.4.9. Let \mathcal{X} be a P -extension problem and $\mathcal{S} = \mathcal{X} \otimes G$ a solution given by the product construction. Then for each $w \in P^*$ and $x \in \text{dom}(w^S \parallel_X)$ there is a $u \in P^*$ s.t. $xw^S = xu^X$ and $[u]_G = [w]_G$.

Proof. Assume $xw^S = y$ for $x, y \in X$. This means that $[x, \varepsilon]w^S = [y, \varepsilon]$. Thus $[x, w] = [y, \varepsilon]$. So, by definition, there is a $u \in P^*$ with $xu^X = y$ and $[u]_G = [w\varepsilon^{-1}]_G = [w]_G$. \square

Note that $[u]_G = [w]_G$ implies $u^S = w^S$. So the u in the previous lemma has the property that $u^S = w^S$ and $xu^X = xw^S$. We can see this as a very weak form of 2-freeness: instead of $u^X = w^S \parallel_X$ we just have $xu^X = xw^S$ for some previously chosen $x \in \text{dom}(w^S \parallel_X)$.⁶

2.4.2 Freeness properties

We translate the various freeness properties that are introduced in Section 2.3 to the setting of abstract extension problems.

For the analogue of cluster ℓ -acyclicity we need the following definitions. A P -cluster is a multidigraph $I = (V, E)$ with an edge labelling $\mu: E \rightarrow P^*$ that is compatible with the involution $(\cdot)^{-1}$ on E , i.e., $\mu(e^{-1}) = \mu(e)^{-1}$. A P -cluster is compatible with a P -generated group G if $[\mu(\alpha)]_G = 1$ for all closed walks α . We call a P -cluster a *free* if it is compatible with $\text{FG}(P)$.

Definition 2.4.10. A solution G of a P -generated inverse monoid M is

- (i) *parallel ℓ -free* if for all $w_1, \dots, w_\ell \in P^*$ with $[w_1]_G = \dots = [w_\ell]_G$, there is a $v \in P^*$ s.t. $[w_1]_M, \dots, [w_\ell]_M \leq [v]_M$.

⁶In [21] solutions where such a u exists are called ‘special’.

2.4 Inverse monoids as ‘abstract extension problems’

- (ii) *sequential ℓ -free* if for all $w_1, \dots, w_\ell \in P^*$ with $[w_1 \dots w_\ell]_G = 1$, there are $v_1, \dots, v_\ell \in P^*$ s.t. $[w_i]_M \leq [v_i]_M$ and $[v_1 \dots v_\ell]_{\text{FG}(P)} = 1$.
- (iii) *cluster ℓ -free* if every P -cluster (I, μ) that is compatible with G and contains at most 2ℓ edges, has a relabelling (I, μ') to a free cluster s.t. $[\mu(e)]_M \leq [\mu'(e)]_M$.
- (iv) *ℓ -free* if for all $w_1, \dots, w_\ell \in P^*$ with $[w_1 \dots w_\ell]_G = 1$ there are $v_1, \dots, v_\ell \in P^*$ s.t. $[w_i]_G = [v_i]_G$, $[w_i]_M \leq [v_i]_M$ and $[v_1 \dots v_\ell]_{\text{FG}(P)} = 1$.

We want to prove that these freeness properties guarantee solutions with the respective free behaviour when used in the product construction. We need the following lemma as a preparation.

Lemma 2.4.11. *Let M be a P -generated inverse monoid. Then G is a 2-free solution of M , if and only if for all $w_1, w_2 \in P^*$ with $[w_1]_G = [w_2]_G$ there is a $w' \in P^*$ s.t. $[w_i]_M \leq [w']_M$ and $[w']_G = [w_1]_G = [w_2]_G$.*

Proof. The proof is straightforward. Simply replace in the statement of 2-freeness the premise $[w_1 w_2]_G = 1$ by the equivalent condition $[w_1^{-1}]_G = [w_2]_G$ and in the conclusion $[v_1 v_2]_{\text{FG}(P)} = 1$ by $v_1^{-1} = v_2$ (w.l.o.g. v_1 and v_2 are reduced). \square

Lemma 2.4.12. *Let \mathcal{X} be a plain P -extension problem and G a solution of $I(\mathcal{X})$ and $\ell \geq 2$. Then*

- (i) $\mathcal{X} \otimes G$ is parallel ℓ -free if G is parallel ℓ -free,
- (ii) $\mathcal{X} \otimes G$ is sequential ℓ -free if G is sequential ℓ -free,
- (iii) $\mathcal{X} \otimes G$ is cluster ℓ -free if G is cluster ℓ -free,
- (iv) $\mathcal{X} \otimes G$ is ℓ -free if G is ℓ -free.
- (v) $\mathcal{X} \otimes \text{FG}(P)$ is free.

Proof. In the following we will use Lemma 2.4.9 multiple times without explicitly mentioning it.

- (i) Let $\mathbf{x} = (x_1, \dots, x_\ell) \in \text{dom}(u^S \parallel_X)$. Then for each x_i there is a w_i s.t. $x_i w_i^{\mathcal{X}} = x_i u^S$ and $[w_i]_G = [u_i]_G$. Hence $[w_1]_G = \dots = [w_\ell]_G$. By parallel ℓ -freeness of G , there is a $v \in P^*$ s.t. $w_1^{\mathcal{X}}, \dots, w_\ell^{\mathcal{X}} \subseteq v^{\mathcal{X}}$. In particular, $x_i v^{\mathcal{X}} = x_i w_i^{\mathcal{X}} = x_i u^S$. So $\mathbf{x} v^{\mathcal{X}} = \mathbf{x} u^S$.
- (ii) Let $(u_i)_{i \in \mathbb{Z}_\ell}$ and $(x_i)_{i \in \mathbb{Z}_\ell}$ with $x_i \in \text{dom}(u_i^S \parallel_X)$ and $(u_0 \dots u_{\ell-1})^S = \text{id}$. W.l.o.g. we can assume that $[u_0 \dots u_{\ell-1}]_G = 1$; simply replace $u_{\ell-1}$ by $u_{\ell-2}^{-1} \dots u_0^{-1}$. For each u_i let w_i be such that $x_i w_i^{\mathcal{X}} = x_i u_i^S$ and $[w_i]_G = [u_i]_G$. Then, by sequential ℓ -freeness of G , there are $(v_i)_{i \in \mathbb{Z}_\ell}$ s.t. $w_i^{\mathcal{X}} \subseteq v_i^{\mathcal{X}}$ and $[v_0 \dots v_{\ell-1}]_{\text{FG}(P)} = 1$. In particular, $x_i v_i^{\mathcal{X}} = x_i w_i^{\mathcal{X}} = x_i u_i^S$.

2 Extension problems

- (iii) Let (I, η, μ) be an \mathcal{X} -cluster compatible with \mathcal{S} where $I = (V, E)$ has at most 2ℓ edges. W.l.o.g. we can assume that (I, μ) is a P -cluster compatible with G : let $T = (V, E')$ be a spanning tree of I . If we change the value of μ for $e \in E \setminus E'$ to $\mu(e_\ell^{-1} \dots e_2^{-1})$ where $e_1 \dots e_\ell$ is the unique cycle starting with e in $(V, E' \cup \{e\})$, then this is a relabelling of (I, η, μ) where (I, μ) is compatible with G .

Similar to (i) and (ii), we can boost the compatibility condition $\eta(e)\mu(e)^S = \eta(e^{-1})$ to $\eta(e)\mu(e)^{\mathcal{X}} = \eta(e^{-1})$. By cluster ℓ -acyclicity of G , there is a free cluster (I, μ') s.t. $\mu(e)^{\mathcal{X}} \leq \mu'(e)^{\mathcal{X}}$ which makes (I, η, μ') a free relabelling.

- (iv) First we prove that if G is 2-free, then for every $u \in P^*$ there is a $w \in P^*$ s.t. $w^{\mathcal{X}} = u^S \parallel_X$ and $[w]_G = [u]_G$: for each $x \in \text{dom}(u^S \parallel_X)$ let w_x be such that $xw_x^{\mathcal{X}} = xu^S$ and $[w_x]_G = [u]_G$. Now, a repeated application of the previous lemma provides a w' with $w_x^{\mathcal{X}} \subseteq w'^{\mathcal{X}}$ and $[w']_G = [u]_G$. So, $xu^S = xw_x^{\mathcal{X}} = w'^{\mathcal{X}}$ for all $x \in \text{dom}(u^S \parallel_X)$.

Now we prove the statement of (iv). Let $(u_1 \dots u_\ell)^S = \text{id}$. As before, w.l.o.g. $[u_1 \dots u_\ell]_{\text{FG}(P)} = 1$. Now let $w_i \in P^*$ such that $w_i^{\mathcal{X}} = u_i^S \parallel_X$ and $[w_i]_G = [u_i]_G$. Then $[w_1 \dots w_\ell]_G = 1$ and by ℓ -freeness there are $v_i \in P^*$ s.t. $[v_i]_G = [w_i]_G$, $w_i^{\mathcal{X}} \subseteq v_i^{\mathcal{X}}$ and $[v_1 \dots v_\ell]_{\text{FG}(P)} = 1$. As $[u_i]_G = [w_i]_G = [v_i]_G$ we have $u_i^S = v_i^S$. Then $v_i^{\mathcal{X}} = u_i^S \parallel_X$ as $v_i^{\mathcal{X}} \subseteq u_i^S \parallel_X = w_i^{\mathcal{X}} \subseteq v_i^{\mathcal{X}}$.

- (v) Let $\mathcal{U} = \mathcal{X} \otimes \text{FG}(P)$. Let $w \in \text{FG}(P)$ and $x \in \text{dom}(w^{\mathcal{U}} \parallel_X)$. Then there is a $u \in P^*$ (w.l.o.g. $u \in \text{FG}(P)$) s.t. $xw^{\mathcal{U}} = xu^{\mathcal{X}}$ and $[u]_{\text{FG}(P)} = [w]_{\text{FG}(P)}$. Hence $u = w$ and so $x \in \text{dom}(w^{\mathcal{X}})$. \square

2.5 Finite solutions with free behaviour

In the previous section we translated approximatively free solutions of plain extension problems into approximatively free solutions of P -generated inverse monoids (abstract extension problems). In this section we present two methods of how to obtain finite, approximatively free solutions.

The first method is based on a construction used by Auinger and Steinberg [5]. Using their methods we can show the existence of finite, parallel ℓ -free solutions and finite, sequential ℓ -free solutions, thus providing alternative proofs of Herwig's Theorem (Theorem 2.1.2) and Theorem 2.1.5. It seems reasonable that this method can also be extended to produce finite, cluster ℓ -free solutions. However, this has not been achieved.

The second method we present here is based on Otto's Conjecture (Conjecture 1.6.22). We explain how Otto's Conjecture implies the Free extension Conjecture.

2.5.1 Cayley graphs and Margolis–Meakin expansions

Cayley graphs and Margolis–Meakin expansions are common, integral tools for manipulating inverse monoids and groups. Cayley graphs are known to play an important role in combinatorial group theory. In the field of inverse monoids Margolis–Meakin expansions are common tools as well (see [27] for a deeper discussion of this concept).

Cayley graphs

Definition 2.5.1. A P -coloured graph is a multidigraph $I = (V, E)$ with an edge-colouring $\pi: E \rightarrow P$ s.t. $\pi(e^{-1}) = \pi(e)^{-1}$ and every vertex $a \in V$ is incident to exactly one edge with colour $p \in P$, i.e., for all $a \in V$ and $p \in P$ there is a unique edge $e \in E$ with $s(e) = a$ and $\pi(e) = p$.

P -coloured graphs I can also be understood as unbranched covers $\pi: I \xrightarrow{\text{unb}} (\{\bullet\}, P)$ where the base graph has just one vertex \bullet and for each $p \in P$ a loop. In particular, the colouring π projects from walks in I to words in P .

We can associate P -generated groups to P -coloured graphs. The coloured edges of $I = (V, E)$ can be understood as graphs of permutations of V . Every $a \in V$ has exactly one p -neighbour, i.e., a vertex $b \in V$ such that there is an edge of colour p that starts at a and ends at b . So every $p \in P$ induces a permutation of V via

$$a \mapsto \text{the unique } p\text{-neighbour of } a.$$

We define $\text{sym}(I)$ as the P -generated subgroup of $\text{Sym}(V)$ generated by these permutations.

Conversely, we can associate to every P -generated group a P -coloured graph, its Cayley graph.

Definition 2.5.2. The *Cayley graph* of a P -generated group G is the P -coloured graph $\Gamma = (G, E)$ where $E = \{(g, p) \mid g \in G, p \in P\}$ and

$$s((g, p)) = g, \quad t((g, p)) = gp, \quad (g, p)^{-1} = (gp, p^{-1}), \quad \pi((g, p)) = p.$$

Walks in the Cayley graph Γ of G describe elements of G . If γ is a walk in Γ , then $[\pi(\gamma)]_G = s(\gamma)^{-1}t(\gamma)$. In particular, $[u]_G = t(\alpha)$ if α is the lift of $u \in P^*$ to 1.

It is a well-known fact that $\text{sym}(\Gamma) \simeq G$ and also that the elements of $\text{sym}(\Gamma)$ are automorphisms of Γ . We can understand these automorphisms as a left-action of G on Γ . The action of $g \in G$ is given by $g.g' := gg'$ on the vertices of Γ and $g.(g', p) := (gg', p)$ on the edges of Γ .

Naturally, this action extends to walks. Some simple observations about this action are: $\pi(g.\alpha) = \pi(\alpha)$ and $g.\alpha$ is the lift of u to gg' if α is the lift of $u \in P^*$ to $g' \in G$.

For a walk $\alpha \in \Gamma$ we write $E(\alpha)$ for the set of edges that are traversed by α (in either direction), i.e., $E(\alpha) := \{e \in E \mid e \in \alpha \text{ or } e^{-1} \in \alpha\}$, and for $u \in P^*$ we let

2 Extension problems

$E(u) := E(\alpha)$ where α is the lift of u to 1 in Γ . We observe that, if β is the lift of u to $g \in G$, then $E(\beta) = E(g.\alpha) = g.E(\alpha) = g.E(u)$. In particular, we can see that $E(u^{-1}) = [u]_G^{-1}.E(u)$ as the lift of u to $[u]_G^{-1}$ is just the reverse of the lift of u^{-1} to 1.

The Margolis-Meakin expansion

We can determine $[u]_G$ for a word $u \in P^*$ and a P -generated group G by tracing the labels of u when starting at 1 in the Cayley graph Γ of G . The Margolis-Meakin expansion of G is a P -generated inverse monoid that does basically this but also keeps track of the edges that are used by this walk.

Let G be a P -generated group and $\Gamma = (G, E)$ its Cayley graph. The powerset $\mathcal{P}(E)$ forms an inverse monoid when equipped with the union operation. The action of G on E extends naturally to an action of G on this inverse monoid $\mathcal{P}(E)$. We use this action to define a semidirect product $\mathcal{P}(E) \rtimes G$: the product of $(E', g), (E'', h) \in \mathcal{P}(E) \rtimes G$ is defined by

$$(E', g)(E'', h) := (E' \cup g.E'', gh).$$

It is straightforward to verify that $\mathcal{P}(E) \rtimes G$ is an inverse monoid and that the inverse of (E', g) is given by $(g^{-1}.E', g^{-1})$. We can also give simple characterisations of \leq , dom and img on $\mathcal{P}(E) \rtimes G$: the natural partial order is given by

$$(E', g) \leq (E'', h) \iff E'' \subseteq E' \text{ and } g = h,$$

and domain and image of (E', g) are

$$\text{dom}((E', g)) = (E', 1), \quad \text{img}((E', g)) = (g^{-1}.E', 1).$$

Definition 2.5.3. The *Margolis-Meakin expansion* $M(G)$ of a P -generated group G is the P -generated inverse submonoid of $\mathcal{P}(\Gamma) \rtimes G$ generated by the family $((E(p), [p]_G))_{p \in P}$ where $\Gamma = (G, E)$ is the Cayley graph of G .

A simple inductive proof shows that the evaluation $[\cdot]_{M(G)}: P^* \rightarrow M(G)$ is given by

$$[u]_{M(G)} = (E(u), [u]_G).$$

Using the characterisations of \leq , dom , img from above we get

$$\begin{aligned} [u]_{M(G)} \leq [v]_{M(G)} &\iff [u]_G = [v]_G \text{ and } E(v) \subseteq E(u), \\ \text{dom}([u]_{M(G)}) &= ((E(u), 1), \text{ and } \text{img}([u]_{M(G)}) = ((E(u^{-1}), 1). \end{aligned}$$

In particular we get that $\text{dom}([u]_{M(G)}) \leq \text{dom}([v]_{M(G)}) \iff E(v) \subseteq E(u)$, $\text{dom}([u]_{M(G)}) \leq \text{img}([v]_{M(G)}) \iff E(v^{-1}) \subseteq E(u)$ etc.

If $[u]_G = 1$, then $[u]_{M(G)} = (E(u), 1) \leq (\emptyset, 1) = 1_{M(G)}$. So, G is a solution of $M(G)$. In fact, $M(G)$ is universal among the P -generated inverse monoids for which G is a solution. To see that we need the following lemma.


Lemma 2.5.4. *Let G be a solution of the P -generated inverse monoid M . If α and γ are walks in the Cayley graph Γ of G s.t. $s(\gamma) = s(\alpha)$, $t(\gamma) = t(\alpha)$ and $E(\gamma) \subseteq E(\alpha)$, then $[\pi(\alpha)]_M \leq [\pi(\gamma)]_M$.*

Proof. We note that $[\pi(\beta)]_M$ is an idempotent, if β is a closed walk in Γ : $s(\beta) = t(\beta)$, then $[\pi(\beta)]_G = s(\beta)^{-1}t(\beta) = 1$ and thus $[\pi(\beta)]_M \leq 1$.

We prove the lemma by induction over the length of γ . If $|\gamma| = 0$, then γ and α are closed walks and thus $[\pi(\alpha)]_M \leq 1 = [\pi(\gamma)]_M$.

For the induction step let e be the last edge in γ , i.e., $\gamma = \gamma'e$. Then e or e^{-1} also appears somewhere in α . We distinguish these two cases.

1. $e \in \alpha$: α can be decomposed into $\alpha = \alpha_1 e \alpha_2$ as in the following sketch




We see that $s(\alpha_2) = t(\alpha_2)$ and thus $[\pi(\alpha_2)]_M$ is an idempotent. Note that we can apply the induction hypothesis to $\alpha_1 e \alpha_2 \alpha_2^{-1} e^{-1}$ and γ' , i.e., $[\pi(\alpha_1 e \alpha_2 \alpha_2^{-1} e^{-1})]_M \leq [\pi(\gamma')]_M$. So we get

$$\begin{aligned} [\pi(\alpha_1 e \alpha_2)]_M &= [\pi(\alpha_1)]_M [\pi(e \alpha_2)]_M = [\pi(\alpha_1)]_M [\pi(e \alpha_2)]_M [\pi(e \alpha_2)]_M^{-1} [\pi(e \alpha_2)]_M \\ &= [\pi(\alpha_1 e \alpha_2 \alpha_2^{-1} e^{-1})]_M [\pi(e \alpha_2)]_M \leq [\pi(\gamma')]_M [\pi(e \alpha_2)]_M \\ &= [\pi(\gamma' e)]_M [\pi(\alpha_2)]_M \leq [\pi(\gamma' e)]_M. \end{aligned}$$

The last inequality is true since $[\pi(\alpha_2)]_M$ is an idempotent.

2. $e^{-1} \in \alpha$: α can be decomposed into $\alpha = \alpha_1 e^{-1} \alpha_2$ as in the following sketch:



We see that $s(e^{-1} \alpha_2) = t(e^{-1} \alpha_2)$ and thus $[\pi(e^{-1} \alpha_2)]_M$ is an idempotent. So $[\pi(e^{-1} \alpha_2)]_M = [\pi(e^{-1} \alpha_2)]_M [\pi(e^{-1} \alpha_2)]_M$ and $[\pi(e^{-1} \alpha_2)]_M = [\pi(\alpha_2^{-1} e)]_M$. We can apply the induction hypothesis to $\alpha_1 e^{-1} \alpha_2 \alpha_2^{-1}$ and γ' , i.e., $[\pi(\alpha_1 e^{-1} \alpha_2 \alpha_2^{-1})]_M \leq [\pi(\gamma')]_M$. So we get

$$\begin{aligned} [\pi(\alpha_1 e^{-1} \alpha_2)]_M &= [\pi(\alpha_1)]_M [\pi(e^{-1} \alpha_2)]_M = [\pi(\alpha_1)]_M [\pi(e^{-1} \alpha_2)]_M [\pi(e^{-1} \alpha_2)]_M \\ &= [\pi(\alpha_1)]_M [\pi(e^{-1} \alpha_2)]_M [\pi(\alpha_2^{-1} e)]_M = [\pi(\alpha_1 e^{-1} \alpha_2 \alpha_2^{-1})]_M [\pi(e)]_M \\ &\leq [\pi(\gamma')]_M [\pi(e)]_M = [\pi(\gamma' e)]_M. \end{aligned}$$

In both cases we have $[\pi(\alpha)]_M \leq [\pi(\gamma' e)]_M = [\pi(\gamma)]_M$. □

2 Extension problems

We recall the definition of $M \leq N$ for P -generated inverse monoids. $M \leq N$ if $[u]_N = [v]_N \implies [u]_M = [v]_M$ for all $u, v \in P^*$. We also have that $M \leq N$, if and only if there is a homomorphism $f: N \rightarrow M$ (which is uniquely defined). So the following corollary says that $M(G)$ is universal among the P -generated inverse monoids for which G is a solution.

Corollary 2.5.5. *A P -generated group G is a solution of a P -generated inverse monoid M if, and only if, $M \leq M(G)$.*

Proof. Let G be a solution of M . If $[u]_{M(G)} = [v]_{M(G)}$, then $[u]_G = [v]_G$. So the lifts α and β of u and v to 1 have the same target. Also $E(\alpha) = E(u) = E(v) = E(\beta)$. Thus by the previous lemma $[u]_M = [\pi(\alpha)]_M = [\pi(\beta)]_M = [v]_M$.

Now let $M \leq M(G)$. If $[u]_G = 1$, then $[u]_{M(G)} = (E(u), [u]_G) = (E(u), 1) \leq 1 = [\varepsilon]_{M(G)}$ and so $[u]_M \leq [\varepsilon]_M = 1$. \square

The various freeness properties of solutions introduced in Definition 2.4.10 are invariant under homomorphic images of P -generated inverse monoids. For example, if G is a parallel ℓ -free solution of N and $M \leq N$, then G is also a parallel ℓ -free solution of M . In particular, since every finite P -generated inverse monoid M has a finite solution G (Lemma 2.4.7), it suffices to show that all Margolis-Meakin expansions of finite groups have finite approximately free solutions to show the existence of finite approximately free solutions for finite inverse monoids in general.

2.5.2 The method of Auinger and Steinberg

In [6] Auinger and Steinberg give a constructive proof of the ‘Ribes-Zaleski Theorem’, a theorem extending the Theorem of M. Hall (Theorem 1.3.5). Their construction is based on objects that, in Group Theory lingo, can be described as ‘universal efficient p -elementary co-extensions’. We give concrete definitions of these ‘co-extensions’ and use the proof ideas of Auinger and Steinberg to show the existence of finite, parallel ℓ -free solutions and finite sequential ℓ -free solutions.

Let I be some index set and $n \in \mathbb{N}$ greater than 1. We denote the elements of the abelian group $\prod_{i \in I} \mathbb{Z}_n$ as formal sums $\sum_{i \in I} x_i i$. We define the *free abelian P -generated n -group* $\text{Ab}_n(P)$ as the subgroup of $\prod_{p \in P} \mathbb{Z}_n$ generated with $((1 \cdot p + (-1) \cdot p^{-1}))_{p \in P}$ as family of generators (note that the generator for p is indeed the inverse of the generator for p^{-1}).

For a P -generated group with Cayley graph $\Gamma = (G, E)$ we can define an action on $\text{Ab}_n(E)$ via $g \cdot (\sum_{e \in E} x_e e) := \sum_{e \in E} x_e g \cdot e = \sum_{e \in E} x_{g^{-1} \cdot e} e$. So we can define a semi-direct product $\text{Ab}_n(E) \rtimes G$ with binary operation given by

$$\left(\sum_{e \in E} x_e e, g \right) \left(\sum_{e \in E} z_e e, f \right) = \left(\sum_{e \in E} (x_e + z_{g^{-1} \cdot e}) e, gf \right).$$

Definition 2.5.6. For a P -generated group G and $n > 1$ we define G^{Ab_n} as the P -generated group given as a subgroup of $\text{Ab}_n(E) \times G$ generated by the family

$$([p]_{\text{Ab}_n(P)}, [p]_G)_{p \in P}$$

where $\Gamma = (G, E)$ is the Cayley graph of G .

The evaluation over G^{Ab_n} can be concisely described as follows: we write $\#_e(\alpha)$ for the number of signed traversals of an edge e by a walk α , i.e., $\#_e(\alpha) := |\alpha|_e - |\alpha|_{e^{-1}}$. A simple inductive argument shows that

$$[u]_{G^{\text{Ab}_n}} := \left(\sum_{e \in E} (\#_e(\alpha) \bmod p) e, [u]_G \right)$$


where α is the lift of u to 1.

The group G^{Ab_n} approximates the Margolis-Meakin expansion in the sense that in $\sum_{e \in E} (\#_e(\alpha) \bmod p) e$ the entries with non-zero entry are in $E(u)$ and the reverse is only true for those edges that get not used a multiple of n times. This information suffices to ensure that G is a parallel 2-free solution of $M(G)$.

Lemma 2.5.7. *Let G be a P -generated group and $n > 1$. Then G^{Ab_n} is a parallel 2-free solution of $M(G)$.*

Proof. Suppose we are given $w_1, w_2 \in P^*$ with $[w_1]_{G^{\text{Ab}_p}} = [w_2]_{G^{\text{Ab}_p}}$. Let $\Gamma = (G, E)$ be the Cayley graph of G , $g = [w_1]_G = [w_2]_G$ and α_1, α_2 the lifts of w_1 and w_2 . We show that there is a walk γ from 1 to g s.t. $E(\gamma) \subseteq E(\alpha_1) \cap E(\alpha_2)$. Then $v = \pi(\gamma)$ is as desired: $[w_1]_{M(G)}, [w_2]_{M(G)} \leq [v]_{M(G)}$ since $E(v) = E(\gamma) \subseteq E(\alpha_1) \cap E(\alpha_2) = E(w_1) \cap E(w_2)$ and $[v]_G = s(\gamma)^{-1}t(\gamma) = 1g = [w_1]_G$.

Let Γ_1 be the weak subgraph of Γ that contains exactly the edges of $E(\alpha_1)$ and Γ_2 the weak subgraph that contains exactly the edges of $E(\alpha_2)$. Let Δ be the connected component of 1 in $\Gamma_1 \cap \Gamma_2$. If $g \in \Delta$, then we get a γ as desired. We show that $g \notin \Delta$ leads to a contradiction. Let $E' \subseteq E(\alpha_1)$ be the set of those edges whose source is in Δ but their target is not:



Then, since $g \notin \Delta$, the walk α_1 has to cross this ‘barrier’ E' effectively once, i.e., $\sum_{e \in E'} \#_e(\alpha_1) = 1$. Thus there is an $e' \in E'$ s.t. $\#_{e'}(\alpha_1) \bmod p \neq 0$. However, $\#_{e'}(\alpha_2) \bmod p = 0$ which contradicts $[w_1]_{G^{\text{Ab}_p}} = [w_2]_{G^{\text{Ab}_p}}$. \square

2 Extension problems

Corollary 2.5.8. *Every finite P -generated inverse monoid has finite, parallel ℓ -free solutions.*

Proof. Let M be a finite P -generated inverse monoid. By Lemma 2.4.7 M has a finite solution G . Set $G_1 = G$ and $G_{i+1} = G_i^{\text{Ab}_n}$. Then G_ℓ is finite and a simple inductive argument shows that G_ℓ is also an ℓ -free solution of $M(G)$: if $[w_1]_{G_\ell} = \cdots = [w_\ell]_{G_\ell}$, then, by the induction hypothesis, there is a $v' \in P^*$ s.t. $[w_1]_{M(G_2)}, \dots, [w_{\ell-1}]_{M(G_2)} \leq [v']_{M(G_2)}$. Note that $[w_1]_{M(G)}, \dots, [w_{\ell-1}]_{M(G)} \leq [v']_{M(G)}$ and $[v']_{M(G_2)} = [w_\ell]_{M(G_2)}$. By the previous lemma there is a $v \in P^*$ s.t. $[v']_{M(G)}, [w_\ell]_{M(G)} \leq [v]_{M(G)}$ and so $[w_1]_{M(G)}, \dots, [w_\ell]_{M(G)} \leq [v]_{M(G)}$. Then also $[w_1]_M, \dots, [w_\ell]_M \leq [v]_M$ since $M \leq M(G)$. Thus, G_ℓ is a finite, parallel ℓ -free solution of M . \square


We show now the existence of finite, sequential ℓ -free solutions. Note that a parallel 2-free solution is also sequential 2-free, which can be proved like the corresponding fact for plain extension problems (see the discussion after Definition 2.3.5).

Lemma 2.5.9. *Let G be a P -generated group. Then G_ℓ with $G_1 = G$ and $G_{i+1} = G_i^{\text{Ab}_p}$ is a sequential ℓ -acyclic solution of $M(G)$.*

Proof. We prove the lemma by induction. As we remarked before, the case $\ell = 2$ is treated in Lemma 2.5.7. So let $\ell > 2$. In order to prove that G_ℓ is sequential ℓ -acyclic we have to show that for given $w_1, \dots, w_\ell \in P^*$ with $[w_1 \dots w_\ell]_{G_\ell} = 1$ there are $v_1, \dots, v_\ell \in P^*$ with $[v_1 \dots v_\ell]_{\text{FG}(P)} = 1$ and $[w_i]_{M(G)} \leq [v_i]_{M(G)}$.

We set $\Gamma = (G_{\ell-1}, E)$ as the Cayley graph of $G_{\ell-1}$ and define vertices g_1, \dots, g_ℓ of Γ and walks $\alpha_1, \dots, \alpha_\ell$ in Γ as follows: $g_1 = 1$ and α_1 is the lift of w_1 to g_1 , for $i > 1$ we set $g_i = t(\alpha_{i-1})$ and α_i as the lift of w_i to g_i . Note that $g_i = [w_1, \dots, w_{i-1}]_{G_{\ell-1}}$ and $E(\alpha_i) = g_i \cdot E(w_i)$.

We show now that there are $j \in \{2, \dots, \ell - 1\}$, a vertex $g \in G_{\ell-1}$ of Γ and walks γ_-, γ_+ in Γ s.t. γ_- goes from g_j to g , γ_+ from g to g_1 , $E(\gamma_-) \subseteq E(\alpha_j)$, and $E(\gamma_+) \subseteq E(\alpha_\ell) \cap E(\alpha_1)$: let Γ_j be the weak subgraph of Γ that includes the edges $E(\alpha_j)$ and Δ the connected component of g_1 in $\Gamma_\ell \cap \Gamma_1$. If $g_2 \in \Delta$ then we can take $j = 2$, γ_- as the trivial walk at g_2 and γ_+ as the walk in Δ from g_2 to g_1 . Otherwise, by the argument given in the proof of Lemma 2.5.7, there is an edge $e \in E$ s.t. $s(e)$ is in Δ , $t(e)$ is not in Γ_ℓ nor Γ_2 , and $\#_e(\alpha_1) \bmod p \neq 0$. So e is in some Γ_j . Thus there are walks γ_-, γ_+ with the desired property from and to $s(e)$.



Now let j, g, γ_- and γ_+ be as above. We set $v_- = \pi(\gamma_-)$ and $v_+ = \pi(\gamma_+)$. Then

$$\begin{aligned} \text{dom}([w_j]_{M(G_{\ell-1})}) &\leq \text{dom}([v_-]_{M(G_{\ell-1})}), \\ \text{img}([w_\ell]_{M(G_{\ell-1})}), \text{dom}([w_j]_{M(G_{\ell-1})}) &\leq \text{img}([v_+]_{M(G_{\ell-1})}). \end{aligned} \quad (*)$$

For example, we have $g_j E(v_-) = E(\gamma_-) \subseteq E(\alpha_j) = g_j(E(w_j))$ and thus $E(v_-) \subseteq E(w_j)$, implying the first inequality. Note that these inequalities also hold in $M(G)$ since $M(G) \leq M(G_{\ell-1})$.

Now consider the following sequence of elements in P^* :

$$(v_+ w_1) \quad w_2 \quad \dots \quad w_{j-1} \quad v_- \quad (v_-^{-1} w_j) \quad w_{j+1} \quad \dots \quad w_{\ell-1} \quad (w_\ell v_+^{-1}).$$

Since $[(v_+ w_1) w_2 \dots w_{j-1} v_-]_{G_{\ell-1}} = 1$ and $[(v_-^{-1} w_j) w_{j+1} \dots w_{\ell-1} (w_\ell v_+^{-1})]_{G_{\ell-1}} = 1$ and both are of length at most $\ell-1$, we get by induction hypothesis $v_1, \dots, v_{j-1}, v, v_j, \dots, v_\ell$ such that $[v_1 \dots v_{j-1} v]_{\text{FG}(P)} = 1$, $[v_j \dots v_\ell]_{\text{FG}(P)} = 1$ and in the following table

$$\begin{array}{cccccccc} (v_+ w_1) & w_2 & \dots & w_{j-1} & v_- & (v_-^{-1} w_j) & w_{j+1} & \dots & w_{\ell-1} & (w_\ell v_+^{-1}) \\ v_1 & v_2 & \dots & v_{j-1} & v & v_j & v_{j+1} & \dots & v_{\ell-1} & v_\ell \end{array}$$

$[x]_{M(G)} \leq [y]_{M(G)}$ for a column pair (x, y) . Then also in the following table

$$\begin{array}{cccccccc} w_1 & w_2 & \dots & w_{j-1} & w_j & w_{j+1} & \dots & w_{\ell-1} & w_\ell \\ (v_+^{-1} v_1) & v_2 & \dots & v_{j-1} & (v v_j) & v_{j+1} & \dots & v_{\ell-1} & (v_\ell v_+) \end{array}$$

$[x]_{M(G)} \leq [y]_{M(G)}$ for a column pair (x, y) . For all pairs except $(w_1, v_+^{-1} v_1)$, $(w_j, v v_j)$ and $(w_\ell, v_\ell v_+)$ this is clear. We show, for example, that $[w_1]_{M(G)} \leq [v_+^{-1} v_1]_{M(G)}$. Since $\text{dom}([w_1]_{M(G)}) \leq \text{img}([v_+]_{M(G)})$ we have that $[v_+^{-1} v_+ w_1]_{M(G)} = [w_1]_{M(G)}$. Thus $[w_1]_{M(G)} = [v_+^{-1}]_{M(G)} [v_+ w_1]_{M(G)} \leq [v_+^{-1}]_{M(G)} [v_1]_{M(G)} = [v_+^{-1} v_1]_{M(G)}$. The other two inequalities are shown similarly. Note that also

$$[(v_+^{-1} v_1) v_2 \dots v_{j-1} (v v_j) v_{j+1} \dots v_{\ell-1} (v_\ell v_+)]_{\text{FG}(P)} = 1.$$

So $v_+^{-1} v_1, v_2, \dots, v_{j-1}, v v_j, v_{j+1}, \dots, v_{\ell-1}, v_\ell v_+$ are as required. \square

Corollary 2.5.10. *Every finite P -generated inverse monoid has finite, sequential ℓ -acyclic solutions.*

2.5.3 The Method of Otto

In Section 1.6 we discussed how Otto's Conjecture (Conjecture 1.6.22) can be used for the existence of finite acyclic branched covers. Here we show how a symmetrised version of Otto's Conjecture (Conjecture 2.5.16) can be used for the existence of finite ℓ -free solutions of P -generated inverse monoids.

Freeness and acyclicity of solutions

Otto's Conjecture is about acyclic groupoids. We want to find a corresponding notion of ℓ -acyclic solution for extension problems. In Lemma 2.3.12 we established a connection between ℓ -free solutions and ℓ -acyclic solutions for plain extension problems. We can adapt these results to inverse monoids easily.

Definition 2.5.11. A solution G of a P -generated inverse monoid is ℓ -acyclic if for all $(w_i)_{i \in \mathbb{Z}_\ell} \subseteq P^*$ with $[w_0 \dots w_{\ell-1}]_G = 1$ there are $j \in \mathbb{Z}_\ell$ and $v_-, v_+ \in P^*$ s.t.

$$\begin{aligned} \text{img}([w_{j-1}]_M), \text{dom}([w_j]_M) &\leq \text{dom}([v_-]_M), \\ \text{img}([w_j]_M), \text{dom}([w_{j+1}]_M) &\leq \text{img}([v_+]_M). \end{aligned}$$

Lemma 2.5.12. *A solution of a P -generated inverse monoid is ℓ -free if, and only if, it is ℓ -acyclic.*

Permutations of the Cayley graphs of Groupoids

We consider I -groupoids \mathbb{H} in which I is also a P -coloured graph (for example the Cayley graph of a P -generated group). In this case we can define permutations on the vertices of the 'Cayley graph' of \mathbb{H} , similar to how we defined $\text{sym}(G)$ for a P -generated group.

Let $I = (V, E)$ be a multidigraph and \mathbb{H} an I -groupoid with generator family $(h_e)_{e \in E}$. Recall that we can evaluate walks in I over \mathbb{H} via $[e_1 \dots e_n]_{\mathbb{H}} = h_{e_1} \dots h_{e_n}$. This evaluation preserves source and target, i.e., $s([\alpha]_{\mathbb{H}}) = s(\alpha)$ and $t([\alpha]_{\mathbb{H}}) = t(\alpha)$.

If I is the Cayley graph of a P -generated group, then the generators h_e of \mathbb{H} carry a P -colouring given by the colour of e . We can use this to define a permutation group on \mathbb{H} .

Definition 2.5.13. For an I -groupoid \mathbb{H} where $I = (G, E)$ with $E = \{(g, p) \mid g \in G, P \in P\}$ is the Cayley graph of a P -generated group G we let $\text{sym}(\mathbb{H})$ be the P -generated group given as a subgroup of $\text{Sym}(\mathbb{H})$ by the generator family $(f_p)_{p \in P}$ where

$$\begin{aligned} f_p: \mathbb{H} &\rightarrow \mathbb{H} \\ h &\mapsto h[(t(h), p)]_{\mathbb{H}} \end{aligned}$$

Inductively we can show that $[u]_{\text{sym}(P)}$ permutes the elements $h \in \mathbb{H}$ as follows:

$$h[u]_{\text{sym}(\mathbb{H})} = h[\alpha]_{\mathbb{H}}$$

where α is the lift of u to $t(h)$ in I .

Symmetric groupoids

An I -groupoid \mathbb{H} is symmetric if symmetries of I can be extended to symmetries of \mathbb{H} . For symmetric I -groupoids with highly symmetric I we can characterise $\text{sym}(\mathbb{H})$ as equivalence classes of walks over I .

A *symmetry* φ of I is an automorphism of I , i.e., a homomorphism $\varphi: I \xrightarrow{\text{hom}} I$ for which φ^{-1} is also a homomorphism. We say that \mathbb{H} is *compatible with a symmetry* φ if the function

$$\mathbb{H} \rightarrow \mathbb{H}; [\alpha]_{\mathbb{H}} \mapsto [\varphi(\alpha)]_{\mathbb{H}}$$

is well defined. If this function is defined, then it is an automorphism of \mathbb{H} .

Definition 2.5.14. An I -groupoid \mathbb{H} is symmetric if \mathbb{H} is compatible with every symmetry of I .

If I is the Cayley graph of some P -generated group G , the action of $g \in G$ on I is a symmetry of I . If \mathbb{H} is compatible with this action we obtain an action of G on \mathbb{H} by

$$g \cdot [\alpha]_{\mathbb{H}} := [g \cdot \alpha]_{\mathbb{H}}.$$

Lemma 2.5.15. Let I be the Cayley graph of a P -generated group G and \mathbb{H} a symmetric I -groupoid. Then the following are equivalent, for $u, v \in P^*$,

- (i) $[u]_{\text{sym}(\mathbb{H})} = [v]_{\text{sym}(\mathbb{H})}$
- (ii) $h[u]_{\text{sym}(\mathbb{H})} = h[v]_{\text{sym}(\mathbb{H})}$ for some $h \in \mathbb{H}$
- (iii) $[\alpha_u]_{\mathbb{H}} = [\alpha_v]_{\mathbb{H}}$ where α_u and α_v are the lifts of u, v to some $g \in G$.

Proof. We set $F = \text{sym}(\mathbb{H})$.

(i) \implies (ii). This is obvious.

(ii) \implies (iii). Let α_u and α_v be the lifts of u and v to $g = t(h)$. Then $h[\alpha_u]_{\mathbb{H}} = h[u]_F = h[v]_F = [\alpha_v]_{\mathbb{H}}$. So $[\alpha_u]_{\mathbb{H}} = [\alpha_v]_{\mathbb{H}}$.

(iii) \implies (i). Given $h \in \mathbb{H}$, let $g' = t(h)$ and α'_u, α'_v the lifts of u and v to g' . Then $\alpha'_u = (g'g^{-1}) \cdot \alpha_u$ and $\alpha'_v = (g'g^{-1}) \cdot \alpha_v$. Thus $[\alpha'_u]_{\mathbb{H}} = [\alpha'_v]_{\mathbb{H}}$ as $[\alpha'_u]_{\mathbb{H}} = [g'g^{-1} \cdot \alpha_u]_{\mathbb{H}} = g'g^{-1} \cdot [\alpha_u]_{\mathbb{H}} = \dots = [\alpha'_v]_{\mathbb{H}}$. So $h'[u]_F = h'[\alpha'_u]_{\mathbb{H}} = h'[\alpha'_v]_{\mathbb{H}} = h'[v]_F$. \square

The conjecture of Otto and approximatively free solutions

In [35] the following conjecture is proposed (originally stated as a theorem but the proof is erroneous).

Conjecture 2.5.16. *For every finite $I = (V, E)$ and $\ell \in \mathbb{N}$ there are finite, symmetric, component ℓ -acyclic I -groupoids.*

We recall some definitions. An I -groupoid \mathbb{H} for $I = (V, E)$ is component ℓ -acyclic if⁷

- (i) for $h, h' \in \mathbb{H}$ with the same source, there is a unique minimal set $E(h, h') \subseteq E$ s.t.

$$h^{-1}h' \in \mathbb{H}(E(h, h')) = \{ [\alpha]_M \mid \alpha \text{ is a walk in } I \text{ and } E(\alpha) \subseteq E(h, h') \}.$$

- (ii) for each sequence $(h_i)_{i \in \mathbb{Z}_n}$ of length at most ℓ , there is a j s.t.

$$h_j \mathbb{H}(E(h_{j-1}, h_j) \cap E(h_j, h_{j+1})) \cap h_{j+1} \mathbb{H}(E(h_j, h_{j+1}) \cap E(h_{j+1}, h_{j+2})) \neq \emptyset.$$

Note that the meaning of $E(\bullet)$ is now quite overloaded. This expression is defined for words $u \in P^*$, walks α in I , and for pairs of elements $h, h' \in \mathbb{H}$:

- (i) $E(\alpha) = \{ e \in E \mid e \in \alpha \text{ or } e^{-1} \in \alpha \}$
- (ii) $E(u) = E(\beta)$ where β is the lift of u to 1.
- (iii) $E(h, h')$ is defined as above in case \mathbb{H} is component 2-acyclic.

We also give some basic observations:

- If $h' = h[\alpha]_{\mathbb{H}}$ for $h, h' \in \mathbb{H}$ and walk α in I , then $E(h, h') \subseteq E(\alpha)$
- $E(\gamma) = g.E(u)$ if γ is the lift of $u \in P^*$ to g in I .

Lemma 2.5.17. *If Conjecture 2.5.16 is true, then there are finite, ℓ -acyclic solutions of $M(G)$ for all finite P -generated groups G and $\ell \in \mathbb{N}$.*

Proof. Let I be the Cayley graph of G , \mathbb{H} a finite, symmetric, component ℓ -acyclic I -groupoid, and $F = \text{sym}(\mathbb{H})$. Clearly F is finite. We show now that F is an ℓ -acyclic solution of $M(G)$.


Let $(w_i)_{i \in \mathbb{Z}_\ell} \subseteq P^*$ be given s.t. $[w_0 \dots w_{\ell-1}]_F = 1$. Let $(g_i)_{i \in \mathbb{Z}_\ell}$ and $(\alpha_i)_{i \in \mathbb{Z}_\ell}$ be defined as follows: for $i = 0$ set $g_0 = 1$ and α_0 to the lift of w_0 to g_0 in I ; for $i > 0$ set $g_i = t(\alpha_{i-1})$ and α_i to the lift of w_i to g_i . Furthermore, we set $h_0 = \text{id}_1$ (id_1 is the neutral element of \mathbb{H} at the object $1 \in G$) and $h_{i+1} = h_i[\alpha_i]_{\mathbb{H}}$. Note that $h_{i+1} = h_i[w_i]_F$ and $E(h_i, h_{i+1}) \subseteq E(\alpha_i)$.

⁷This is a slightly altered but equivalent formulation of component acyclicity (Definition 1.6.20).

Since \mathbb{H} is component ℓ -acyclic there are $j \in \mathbb{Z}_\ell$ and $h \in \mathbb{H}$ s.t.

$$h \in h_j \mathbb{H}(E(h_{j-1}, h_j) \cap E(h_j, h_{j+1})) \cap h_{j+1} \mathbb{H}(E(h_j, h_{j+1}) \cap E(h_{j+1}, h_{j+2})).$$

So there are walks γ_- and γ_+ in I s.t. $h_j[\gamma_-]_{\mathbb{H}} = h$, $h[\gamma_+]_{\mathbb{H}} = h_{j+1}$ and $E(\gamma_-) \subseteq E(h_{j-1}, h_j) \cap E(h_j, h_{j+1})$, $E(\gamma_+) \subseteq E(h_j, h_{j+1}) \cap E(h_{j+1}, h_{j+2})$. We can give a sketch on how these walks lie in I :



Then j , $v_- = \pi(\gamma_-)$ and $v_+ = \pi(\gamma_+)$ are witnesses according to the definition of ℓ -acyclicity (Definition 2.5.11), i.e., $[w_j]_F = [v_- v_+]_F$ and

$$\begin{aligned} \text{img}([w_{j-1}]_{M(G)}), \text{dom}([w_j]_{M(G)}) &\leq \text{dom}([v_-]_{M(G)}), \\ \text{img}([w_j]_{M(G)}), \text{dom}([w_{j+1}]_{M(G)}) &\leq \text{img}([v_+]_{M(G)}). \end{aligned}$$

First we note that $h_j^{-1}[v_- v_+]_F = h_j[\gamma_- \gamma_+]_{\mathbb{H}} = h_{j+1} = h_j[w_j]_F$. Thus by Lemma 2.5.15 $[v_- v_+]_F = [w_j]_F$. We show $\text{img}([w_{j-1}]_{M(G)}) \leq \text{dom}([v_-]_{M(G)})$, the other inequalities are shown similarly. We have $E(\gamma_-) \subseteq E(h_{j-1}, h_j) \subseteq E(\alpha_{j-1})$. And since $E(\gamma_-) = g_j.E(v_-)$ and $E(\alpha_{j-1}) = E(\alpha_{j-1}^{-1}) = g_j.E(w_{j-1}^{-1})$, we get $E(w_{j-1}^{-1}) \subseteq E(v_-)$. Thus $\text{img}([w_{j-1}]_{M(G)}) \leq \text{dom}([v_-]_{M(G)})$. \square

Corollary 2.5.18. *If Conjecture 2.5.16 is true, then every finite, P -generated inverse monoid has a finite ℓ -free solution.*

Proof. Let M be a finite P -generated inverse monoid. By Lemma 2.4.7 M has a finite solution G . Let F be a finite, ℓ -free solution of $M(G)$. By Corollary 2.5.5, $M \leq M(G)$. So F is also an ℓ -free solution of M . \square

2.5.4 The Free Extension Conjecture for extension problems

We generalise the definitions of ‘ $I(\mathcal{X})$ ’ and ‘ $\mathcal{X} \otimes G$ ’ to P -extension problems \mathcal{X} over σ in the obvious ways:

- $I(\mathcal{X})$ is the inverse submonoid of $\text{Sym}(\text{Str}(\mathcal{X}))$ generated by the family $(p^{\mathcal{X}})_{p \in P}$.
- for a P -generated group we set $\mathcal{X} \otimes P$ as the tuple $(S, (R^S)_{R \in \sigma}, (p^S)_{p \in P})$ where:

2 Extension problems

★ $S := (X \times P^*)/\sim$ where

$$(x, u) \sim (y, w) : \iff \exists v \in P^* \text{ s.t. } xv^{\mathcal{X}} = y \text{ and } [v]_G = [uw^{-1}]_G,$$

★ $[x, u]p^{\mathcal{S}} := [x, up]$ where $[x, u]$ is the equivalence class of (x, u) ,

★ $R^{\mathcal{S}} := \{ [\mathbf{x}, u] \mid \mathbf{x} \in E^{\mathcal{X}}, u \in P^* \}$ where $[(x_1, \dots, x_n), u] = ([x_1, u], \dots, [x_n, u])$.

Note that $\mathcal{X}' \otimes G = (\mathcal{X} \otimes G)(\mathcal{X}')$ if \mathcal{X}' is an expansion of a plain extension problem (cf. Definition 2.3.3). So, if G is parallel ℓ -acyclic for sufficiently high ℓ , then $\mathcal{X}' \otimes G$ is a solution of \mathcal{X} (Lemma 2.3.4).

Note that a 2-free solution is parallel ℓ -free for all $\ell \in \mathbb{N}$ (Lemma 2.4.11).

Reduction of Conjecture 2.1.11 to Conjecture 2.5.16. Assume Conjecture 2.5.16 is true. Let \mathcal{X} be a finite P -extension problem. By Corollary 2.5.18 there is a finite ℓ -free solution G of $I(\mathcal{X})$. Then $\mathcal{X} \otimes G$ is a finite solution of \mathcal{X} and by Lemma 2.4.12 it is also ℓ -free. \square

2.6 Applications

2.6.1 EPPA for tournaments and groups of odd order

The question whether the extension problem for a finite tournament has finite solutions is an important open problem [29]. We give an equivalent problem for inverse monoids.

A tournament is a ‘complete directed graph’. Formally we define a tournament as a structure $\mathcal{A} = (A, R^{\mathcal{A}})$ s.t. for distinct $a, b \in A$, either $(a, b) \in R^{\mathcal{A}}$ or $(b, a) \in R^{\mathcal{A}}$. We set TRN to be the class of all tournaments. Below we show that every extension problem in TRN has a solution in TRN.⁸

In [32] Moon shows that every finite group of odd order is the automorphism group of some finite tournament. This shows that tournaments and groups of odd order are tightly connected. We can easily convince ourselves of the converse of Moon’s Theorem. Any non-trivial involution of a tournament $\mathcal{A} = (A, R^{\mathcal{A}})$ would swap at least two vertices $a, b \in A$ but then either $(a, b), (b, a) \notin R^{\mathcal{A}}$ or $(a, b), (b, a) \in R^{\mathcal{A}}$. So, the automorphism group of a tournament does not have non-trivial involutions and thus its order is odd.

We can get a similar statement for the inverse monoid of partial automorphisms of a tournament. The idempotents of an inverse monoid are involutions, we call these the trivial involutions. So, a non-trivial involution of an inverse monoid M is an element $m \in M$ s.t. $m \not\leq 1$ and $m = m^{-1}$. Note that $m = m^{-1}$ and $mm \leq 1$ are not equivalent conditions, the latter is a weaker condition. We say that an inverse monoid is *involution-free* if it has no non-trivial involutions. By the same argument

⁸This can also be observed using Fraïssé-limits. The class of all finite tournaments has a Fraïssé-limit which can be used to solve finite extension problems in TRN.

as for the automorphism group of a tournament \mathcal{A} , we get that $I(\mathcal{A})$, the inverse monoid of partial automorphisms of \mathcal{A} , is involution-free.

We say that $\mathcal{A} = (A, R^{\mathcal{A}})$ is a pre-tournament if for distinct $a, b \in A$, not $(a, b) \in R^{\mathcal{A}}$ and $(b, a) \in R^{\mathcal{A}}$. Clearly, pre-tournaments can be completed to tournaments. We can also preserve an involution-free inverse submonoid of the partial automorphisms in this completion process.

Lemma 2.6.1. *Let $\mathcal{A} = (A, R^{\mathcal{A}})$ be a pre-tournament and $M \subseteq I(\mathcal{A})$ an involution-free inverse submonoid. Then we can augment $R^{\mathcal{A}} \subseteq R^{\mathcal{A}'}$ s.t. $\mathcal{A}' = (A, R^{\mathcal{A}'})$ is a tournament and $M \subseteq I(\mathcal{A}')$.*

Proof. If \mathcal{A} is a tournament, there is nothing to prove.

In the other case there are $a, b \in A$ s.t. $(a, b), (b, a) \notin R^{\mathcal{A}}$. Put $R^{\mathcal{A}'} = R^{\mathcal{A}} \cup \{(a, b)f \mid a, b \in \text{dom}(f), f \in M\}$. Then $\mathcal{A}' = (A, R^{\mathcal{A}'})$ is also a pre-tournament and $M \subseteq I(\mathcal{A}')$. We can repeat this process until \mathcal{A}' is a tournament. For infinite \mathcal{A} the argument is via Zorn's Lemma. \square

Lemma 2.6.2. *A finite extension problem $I(\mathcal{X}) \in \text{TRN}$ has a finite solution in TRN if, and only if, $I(\mathcal{X})$ has a finite solution of odd order.*

Proof. Let \mathcal{S} be a solution of \mathcal{X} . Then the group generated by the family $(p^{\mathcal{S}})_{p \in P}$ is a solution of $I(\mathcal{X})$ and it is involution-free.

Conversely, let G be a solution of $I(\mathcal{X})$ of odd order. Let $G' = G^{\text{Ab}_3}$. We show that G' has odd order. Let $u \in P^*$ with $[u^2]_{G'} = 1$. The value of u in G' has the form $[u]_{G'} = (\sum x_e e, [u]_G)$ where the summation is over the edges of the Cayley graph of G and the x_e are elements of \mathbb{Z}_3 . Since $[u^2]_{G'} = 1$ we also have $[u^2]_G = 1$ and thus $[u]_G = 1$. Then we have

$$(0, 1) = [u]_{G^{\text{Ab}_3}} [u]_{G^{\text{Ab}_3}} = \left(\sum_{e \in E} x_e e, 1 \right) \left(\sum_{e \in E} x_e e, 1 \right) = \left(\sum_{e \in E} 2x_e e, 1 \right).$$

So all the x_e are 0 and thus $[u]_{G'} = 1$.

We set $\mathcal{S} = \mathcal{X} \otimes G'$. By Lemma 2.3.4, \mathcal{S} is a solution of \mathcal{X} and since it is tidy it is a pre-tournament. The group generated by the $p^{\mathcal{S}}$ has no non-trivial involutions: let $u \in P^*$ with $(uu)^{\mathcal{S}} = \text{id}$. Then the order n of $[u]_{G'}$ is odd. Thus $(u^n)^{\mathcal{S}} = \text{id}$ and so $u = u(u^n)^{\mathcal{S}} = \text{id}$.

We can apply Lemma 2.6.1 in order to obtain a completion \mathcal{S}' of \mathcal{S} that is in TRN. \square

A corollary to the proof is that every involution free solution of $I(\mathcal{X})$ can be used to obtain a solution of \mathcal{X} in TRN. Since $\text{FG}(P)$ is involution-free we get the following.

Corollary 2.6.3. *Every extension problem in TRN has a solution in TRN.*

A direct corollary of Lemma 2.6.2 is the following.

Corollary 2.6.4. *The following two statements are equivalent.*

2 Extension problems

- (i) Every P -extension problem in TRN has a finite solution in TRN,
- (ii) Every finite, involution-free P -generated inverse monoid has a solution of odd order.

Proof. (i) \implies (ii): Given a finite P -generated inverse monoid M , we get by the Wagner-Preston Theorem (Theorem 2.4.2) a finite plain extension problem \mathcal{X} with $M = I(\mathcal{X})$. By Lemma 2.6.1 we can expand \mathcal{X} to an extension problem $\mathcal{X}' \in \text{TRN}$. The automorphism group of a solution of \mathcal{X}' gives a solution of M .

(ii) \implies (i) is given by Lemma 2.6.2. □

2.6.2 F-inverse Covers of inverse monoids

An important theorem in the area of inverse monoids is *McAlister's Covering Theorem* [26, 31]:

Theorem 2.6.5 (McAlister). *Every inverse monoid has E -unitary covers.*

The proof of McAlister's Covering Theorem can be adapted to finite inverse monoids as well. The finite version of McAlister's Covering Theorem is as follows:

Theorem 2.6.6. *Every finite inverse monoid has finite, E -unitary covers.*

Actually, in [26] the classical version and the finite version of McAlister's Covering Theorem are proved as a single statement.

A strengthening of McAlister's Covering Theorem is Lawson's Covering Theorem [26]. It shows the existence of F -inverse covers instead of 'just' E -unitary covers.

Theorem 2.6.7 (Lawson). *Every inverse monoid has F -inverse covers.*

We show, that we can reduce the finite version of Lawson's Covering Theorem to Conjecture 2.5.16.

Lemma 2.6.8. *If Conjecture 2.5.16 is true, then every finite inverse monoid has finite F -inverse covers.*

Lemma 2.6.8 was first conjectured by Henckell and Rhodes [19] as a possible route to an affirmative answer for the so-called pointlike conjecture, an important conjecture in the theory of monoids (cf. [18]). The 'pointlike conjecture' was proved to be true by Ash [4] but the validity of the finite version of Lawson's Covering Theorem remains an open problem. Some conditional results on the conjecture of Henckell and Rhodes are given in [7, 46, 47].

We give now the definitions of the notions used in the formulations of the Theorems of McAlister and Lawson. We then discuss a proof of McAlister's Covering Theorem that can be altered to show that the finite version of Lawson's covering theorem can be reduced to Otto's Conjecture (Conjecture 2.5.16).

Covers, E-unitary inverse monoids and F-inverse monoids

We explain the notions used in the theorems of McAlister and Lawson.

A homomorphism $\varphi: N \xrightarrow{\text{hom}} M$ between inverse monoids preserves idempotents: if e is an idempotent, then $\varphi(e)$ is an idempotent since $\varphi(e)\varphi(e) = \varphi(ee) = \varphi(e)$. We say a homomorphism is *idempotent separating* if its restriction $\varphi|_E(N): E(N) \rightarrow E(M)$ is injective.

Definition 2.6.9. An inverse monoid N is a cover of the inverse monoid M if there is a surjective, idempotent separating $\varphi: N \xrightarrow{\text{hom}} M$.

The *compatibility relation* \sim on an inverse monoid M is defined as

$$m \sim n \quad :\iff \quad mn^{-1}, m^{-1}n \in E(M).$$

Two partial bijections $f, g \in I(X)$ of the symmetric inverse monoid are compatible, if and only if $f \cup g$ is a partial bijection. In particular, $f \sim g$ is a necessary condition for f and g having a common upper bound. In general, the compatibility relation is not transitive.

Definition 2.6.10. An inverse monoid is *E-unitary* if \sim is transitive.

E-unitary inverse monoids play an important role in Inverse Semigroup Theory and there are numerous characterisations of them. A good characterisation showing that an inverse monoid is *E-unitary* is the following: An inverse monoid M is *E-unitary* if, and only if,

$$e \leq m \text{ for some } e \in E(M) \quad \iff \quad m \in E(M)$$

for all $m \in M$.

F-inverse monoids are a strengthening of *E-unitary* inverse monoids.

Definition 2.6.11. An inverse monoid M is *F-inverse* if for every $m \in M$ there is a unique maximal $n \in M$ greater than m w.r.t. the natural partial order.

We give an alternative definition of *F-inverse* monoids. An inverse monoid M is *F-inverse* if, and only if,

$$z \leq m, n \text{ for some } z \quad \implies \quad m, n \leq z' \text{ for some } z'$$

for all $m, n \in M$.

With this characterisation we can easily see that every *F-inverse* monoid is also *E-unitary*: if $e \leq m$ with $m \in M$ and $e \in E(M)$, then $e \leq m, 1$. So $m \leq 1$, since 1 is maximal.

Constructing E-unitary covers

Examples of E -unitary inverse monoids are Margolis-Meakin expansions of P -generated groups. The compatibility relation is quite simple over $M(G)$.

$$[u]_{M(G)} \sim [v]_{M(G)} \iff [uv^{-1}]_{M(G)}, [u^{-1}v]_{M(G)} \leq 1 \iff [u]_G = [v]_G.$$

In particular, it is transitive over $M(G)$.

$M(G)$ looks like a promising candidate for an E -unitary cover. If M is a P -generated inverse monoid and G a solution of M , $M \leq G(M)$. So we have a homomorphism

$$\varphi: M(G) \rightarrow M; [u]_{M(G)} \mapsto [u]_M.$$

However, in general φ is not idempotent preserving.

We can obtain an E -unitary cover by passing to a quotient of $M(G)$. Let \sim_φ be the kernel of φ , i.e., the congruence $m \sim_\varphi n \iff \varphi(m) = \varphi(n)$. Since $M(G)$ is E -unitary, \sim is a congruence on $M(G)$ as well. So we can define the congruence $\sim_M := (\sim_\varphi \cap \sim)$ on $M(G)$.

We can characterise the elements of $M(G)/\sim_M$ as follows:

$$\begin{aligned} [u]_{M(G)/\sim_M} = [v]_{M(G)/\sim_M} &\iff [u]_{M(G)} \sim_M [v]_{M(G)} \\ &\iff [u]_{M(G)} \sim_\varphi [v]_{M(G)} \text{ and } [u]_{M(G)} \sim [v]_{M(G)} \\ &\iff [u]_M = [v]_M \text{ and } [u]_G = [v]_G \end{aligned}$$

This actually suggests the following characterisation of $M(G)/\sim_M$: for a P -generated group G' , we define the product $M \times_P G'$ as the P -generated inverse submonoid of $M \times G'$ with the generating family $(([p]_M, [p]_{G'}))_{p \in P}$. By the discussion above we see that $M(G)/\sim_M \simeq M \times_P G$ if G is a solution of M .

The natural partial order on $M \times_P G$ is given by

$$(m_1, g_1) \leq (m_2, g_2) \iff m_1 \leq m_2 \text{ and } g_1 = g_2,$$

and the idempotents are given by

$$(m, g) \in E(M \times_P G) \iff m \in E(M) \text{ and } g = 1.$$

Lemma 2.6.12. *Let M be a P -generated inverse monoid and G a P -generated group. Then the projection $\pi: M \times_P G \rightarrow M; (m, g) \mapsto m$ is an idempotent separating, surjective homomorphism.*

Proof. Clearly π is a homomorphism and surjective. π is injective on $E(M \times_P G)$ since idempotents in $M \times_P G$ are purely characterised by their first component. \square

In the case that G is a solution of M , i.e., when $M \times_P G \simeq M(G)/\sim_M$ we get that $M \times_P G$ is E -unitary.

Lemma 2.6.13. *$M \times_P G$ is an E -unitary cover of M if G is a solution of the P -generated inverse monoid M .*

Proof. We are only left to show that $N = M \times_P G$ is E -unitary. Let $[u]_N \leq [w]_N$ and $[u]_N \in E(N)$. Then $[u]_G = [w]_G$ and $[u]_G = 1$. Hence $[w]_G = 1$. Since G is compatible with M , this gives $[w]_M \leq 1$, i.e., $[w]_M = e \in E(M)$. So $[w]_N = ([w]_M, [w]_G) = (e, 1) \in E(N)$. \square

We now have all the tools to prove McAlister's Covering Theorem and its finite version.

Proof of Theorems 2.6.5 and 2.6.6. Let M be a finite inverse monoid. Choose a suitable P and family of generators $(m_p)_{p \in P}$ for M such that M is a P -generated inverse monoid. In case that M is finite, we can choose P to be finite as well.

Now let G be $\text{FG}(P)$ or, in case that M is finite, a finite solution of M (guaranteed to exist by Lemma 2.4.7). Then by Lemma 2.6.13, $N = M \times_P G$ is an E -unitary cover of M and if M is finite, N is finite as well. \square

Constructing F -inverse covers (provided Conjecture 2.5.16 is true)

We can use the product construction introduced above and Corollary 2.5.18 to obtain F -inverse covers.

Lemma 2.6.14. *$M \times_P G$ is an F -inverse cover of M if G is a 2-free solution of the P -generated inverse monoid M .*

Proof. Let $N = M \times_P G$. If $[w]_N \leq [u_1]_N, [u_2]_N$, then $[u_1]_G = [w]_G = [u_2]_G$. Since G is strongly compatible with M , we obtain a $v \in P^*$ s.t. $[v]_G = [u_1]_G = [u_2]_G$ and $[u_1]_M, [u_2]_M \leq [v]_M$. So $[u_1]_N = ([u_1]_M, [u_1]_G) \leq ([v]_M, [v]_G) = [v]_N$ and $[u_2]_N = ([u_2]_M, [u_2]_G) \leq ([v]_M, [v]_G) = [v]_N$. \square

We can now prove Theorem 2.6.7 and Lemma 2.6.8.

Proof of Theorem 2.6.7 and Lemma 2.6.8. Let M be a finite inverse monoid. Choose a suitable P and family of generators $(m_p)_{p \in P}$ for M such that M is a P -generated inverse monoid. In case that M is finite, we can choose P to be finite as well.

Now let G be $\text{FG}(P)$ or, in case that M is finite, a finite 2-free solution of M (here we assume that Conjecture 2.5.16 is true which then by Corollary 2.5.18 gives the existence of such solutions). Then by Lemma 2.6.14, $N = M \times_P G$ is an F -inverse cover and if M is finite, N is finite as well. \square

2.7 Symmetric solutions

We provide a version of the Free Extension Conjecture that respects symmetries of the extension problem.

Symmetries of extension problems

Definition 2.7.1. A symmetry of an extension problem $\mathcal{X} = (X, (p^\mathcal{X})_{p \in P})$ is a tuple (η, φ) where η is a permutation of X and φ a permutation of P that is compatible with the involution $(\cdot)^{-1}$, i.e., $\varphi(p)^{-1} = \varphi(p^{-1})$, s.t.

$$\eta p^\mathcal{X} \eta^{-1} = \varphi(p)^\mathcal{X} \quad \text{for all } p \in P.$$

A symmetry of a solution \mathcal{S} of \mathcal{X} has the same format.

A symmetry (η, φ) of \mathcal{X} extends naturally to $u = p_1 \dots p_n \in P^*$ via $\varphi(u) = \varphi(p_1) \dots \varphi(p_n)$. By induction we get that the following diagram commutes:

$$\begin{array}{ccc} X & \xrightarrow{u^\mathcal{X}} & X \\ \eta \uparrow & & \uparrow \eta \\ X & \xrightarrow{\varphi(u)^\mathcal{X}} & X \end{array}$$

For example $\eta u^\mathcal{X} = \varphi(u)^\mathcal{X} \eta$. In particular, the two compositions have the same domain. So we have $x \in \text{dom}(\eta u^\mathcal{X}) \iff x \in \text{dom}(\varphi(u)^\mathcal{X} \eta)$ and since η is total, we get that $x \eta \in \text{dom}(u^\mathcal{X}) \iff x \in \text{dom}(\varphi(u)^\mathcal{X})$.

Definition 2.7.2. A symmetry (η, φ) of an extension problem \mathcal{X} is *compatible* with a solution \mathcal{S} if there is an extension η' of η s.t. (η, φ) is a symmetry of \mathcal{S} . A solution is *symmetric* if it is compatible with all symmetries of \mathcal{X} .

Obtaining symmetric solutions for plain extension problems is easy.

Lemma 2.7.3. *Every finite plain extension problem has finite, symmetric solutions.*

Proof. Let $\mathcal{X} = (X, (p^\mathcal{X})_{p \in P})$ be a finite extension problem. We construct the solution \mathcal{S} on the domain X . The idea of how to define $p^\mathcal{S}$ is by ‘closing orbits of $p^\mathcal{X}$ ’.

Let $x \in X$ then $x p^\mathcal{S} := x p^\mathcal{X}$ if $x \in \text{dom}(p^\mathcal{X})$. In the other case, let $n \in \mathbb{N}$ be the largest power s.t. $x \in \text{img}((p^n)^\mathcal{X})$ and define $x p^\mathcal{S} := x (p^{-n})^\mathcal{X}$.

Let (η, φ) be a symmetry of \mathcal{X} . We have to show that $x \eta p^\mathcal{S} = x \varphi(p)^\mathcal{S} \eta$ for all $x \in X$ and $p \in P$. If $x \eta \in \text{dom}(p^\mathcal{X})$, then the equation holds since (η, φ) is a symmetry of \mathcal{X} . If $x \eta \notin \text{dom}(p^\mathcal{X})$ then $x \notin \text{dom}(\varphi(p)^\mathcal{X})$. Also for all $n \in \mathbb{N}$ we have that $x \eta \in \text{img}((p^n)^\mathcal{X}) \iff x \in \text{img}(p^{-n})^\mathcal{X}$. So let $n \in \mathbb{N}$ be the maximum s.t. both sides of the equivalence are true. Then

$$x \eta p^\mathcal{S} = x \eta (p^{-n})^\mathcal{X} = x \varphi(p^{-n})^\mathcal{X} \eta = x \varphi(p)^\mathcal{S} \eta. \quad \square$$

Symmetries of abstract extension problems

We can also give an algebraic description of symmetric solutions.

Definition 2.7.4. A *symmetry* of a P -generated monoid M is a permutation $\varphi: P \rightarrow P$ compatible with $(\cdot)^{-1}$ s.t.

$$[u]_M = [v]_M \implies [\varphi(u)]_M = [\varphi(v)]_M$$

for all $u, v \in P^*$. A solution G of M is *symmetric* if every symmetry of M is also a symmetry of G .

If φ is a symmetry of M then φ induces an automorphism of M via $\varphi([u]_M) := [\varphi(u)]_M$.

As a corollary of Lemma 2.7.3 we get.

Corollary 2.7.5. *Every finite inverse monoid has a finite symmetric solution.*

Not surprisingly we can use symmetric solutions of abstract extension problems to solve extension problems.

Lemma 2.7.6. *Let \mathcal{X} be a plain extension problem. If G is a symmetric solution of $I(\mathcal{X})$. Then $\mathcal{X} \otimes G$ is a symmetric solution of \mathcal{X} .*

Proof. Let $\mathcal{X} = (X, (p^{\mathcal{X}})_{p \in P})$ be a plain extension problem and $\mathcal{S} = \mathcal{X} \otimes G = (S, (p^{\mathcal{S}})_{p \in P})$ its free solution. Given a symmetry (η, φ) of \mathcal{X} , we need to find a permutation $\eta': S \rightarrow S$ s.t.

$$\eta' = (\varphi(v)^{\mathcal{S}})^{-1} \eta' v^{\mathcal{S}}$$

for all $v \in P^*$ (consider the commuting diagram above). If we apply both sides to $[x, u]$ and choose v as $\varphi(u^{-1})$ we obtain the following requirement for η' which we take as its definition

$$[x, u]\eta' := [x\eta, \varphi^{-1}(u)].$$

Clearly η' extends η as $[x, \varepsilon]\eta' = [x\eta, \varepsilon]$. We are left to check that η' is well-defined and fulfils the symmetry condition.

If $(x, u) \sim (y, w)$ then there is a $v \in P^*$ s.t. $xv^{\mathcal{X}} = y$ and $[v]_G = [uw^{-1}]_G$. Then $\varphi^{-1}(v)$ shows that $(x\eta, \varphi^{-1}(u)) \sim (y\eta, \varphi^{-1}(w))$ since $(x\eta)\varphi(v^{-1})^{\mathcal{X}} = y\eta$ and $[\varphi^{-1}(v)]_G = [\varphi^{-1}(uw^{-1})]_G$. Thus η' is well-defined.

Note that the inverse of η' is given by $[x, u](\eta')^{-1} = [x\eta^{-1}, \varphi(v)]$. We see that (η', φ) is a symmetry of \mathcal{S}' .

$$\begin{aligned} [x, u]\eta' u^{\mathcal{S}} (\eta')^{-1} &= [x\eta, \varphi^{-1}(u)] p^{\mathcal{S}} (\eta')^{-1} = [x\eta, \varphi^{-1}(u)p] (\eta')^{-1} \\ &= [x, u\varphi(p)] = [x, u]\varphi(p)^{\mathcal{S}}. \end{aligned}$$

□

Note that $\text{FG}(P)$ is a symmetric solution of any P -generated inverse monoid. So by Lemma 2.4.12 (v) we get the following corollary.

Corollary 2.7.7. *Free solutions are symmetric.*

The symmetric Free Extension Theorem

We can alter the reduction of the Free Extension Conjecture to Conjecture 2.5.16 slightly such that it yields symmetric solutions. For that we have to show that all construction steps do not break symmetries

If φ is a symmetry of a P -generated group G then φ defines an automorphism of the Cayley graph Γ of G which acts on the edges via $\varphi((g, p)) = (\varphi(g), \varphi(p))$. If α_u is the lift of $u \in P^*$ to g , then $\varphi(\alpha_u)$ is the lift of $\varphi(u)$ to $\varphi(g)$. This implies that the symmetries of G are symmetries of its Margolis-Meakin expansion and vice versa.

Lemma 2.7.8. *Let G be a P -generated group. Then φ is a symmetry of G if, and only if φ is a symmetry of $M(G)$.*

Proof. Let φ be a symmetry of G . If $[u]_{M(G)} = [v]_{M(G)}$, then $E(u) = E(v)$ and $[u]_G = [v]_G$. Since φ is a symmetry of G , we have $[\varphi(u)]_G = [\varphi(v)]_G$. So in order to show $[\varphi(u)]_{M(G)} = [\varphi(v)]_{M(G)}$ we are left to show that $E(\varphi(u)) = E(\varphi(v))$. Let α_u and α_v the lifts of u and v to 1. Then $\varphi(\alpha_u)$ and $\varphi(\alpha_v)$ are the lifts of $\varphi(u)$ and $\varphi(v)$ to 1. Thus $E(\varphi(u)) = E(\varphi(\alpha_u)) = \varphi(E(\alpha_u)) = \varphi(E(\alpha_v)) = E(\varphi(\alpha_v)) = E(\varphi(v))$.

Conversely let φ be a symmetry of $M(G)$. If $[u]_G = [v]_G$, then $[vv^{-1}u]_G = [u]_G = [v]_G = [uu^{-1}v]_G$. Let α_u and α_v be the lifts of u, v to 1. Then $E(vv^{-1}u) = E(\alpha_v) \cup E(\alpha_u) = E(uu^{-1}v)$. Thus $[vv^{-1}u]_{M(G)} = [uu^{-1}v]_{M(G)}$. Since φ is a symmetry of $M(G)$, $[\varphi(vv^{-1}u)]_{M(G)} = [\varphi(uu^{-1}v)]_{M(G)}$ and thus $[\varphi(u)]_G = [\varphi(vv^{-1}u)]_G = [\varphi(uu^{-1}v)]_G = [\varphi(v)]_G$ \square

Lemma 2.7.9. *If Conjecture 2.5.16 is true, then every finite inverse monoid has a finite, symmetric, ℓ -free solution.*

Proof. We proceed in two steps: First, we show that we can w.l.o.g. assume that M is a Margolis-Meakin expansion. Then we show that the solution F constructed in the proof of Lemma 2.5.17 is a symmetric solution of $M(G)$.

Let M be a finite P -generated inverse monoid. By Corollary 2.7.5, there is a symmetric solution G of M . We show that if F is a symmetric solution of $M(G)$, then it is also a symmetric solution of M . Let φ be a symmetry of M . Then, since G is a symmetric solution, it is also a symmetry of G . By Lemma 2.7.8, it is then also a symmetry of $M(G)$ and thus a symmetry of F .

Now let G be a finite P -generated group and F the solution constructed as in the proof of Lemma 2.5.17: I is the Cayley graph of G , \mathbb{H} a symmetric, component ℓ -acyclic I -groupoid, and $F = \text{sym}(\mathbb{H})$.

Let φ be a symmetry of $M(G)$. Then it is a symmetry of G and a symmetry of I .

Assume $[u]_F = [v]_F$. By Lemma 2.5.15 we have $[\alpha_u]_{\mathbb{H}} = [\alpha_v]_{\mathbb{H}}$ for the lifts of u and v to 1. Since \mathbb{H} is symmetric, we have $[\varphi(\alpha_u)]_{\mathbb{H}} = [\varphi(\alpha_v)]_{\mathbb{H}}$. The walks $\varphi(\alpha_u)$ and $\varphi(\alpha_v)$ are the lifts of $\varphi(u)$ and $\varphi(v)$ to 1 and so by Lemma 2.5.15 again, we have $[\varphi(u)]_F = [\varphi(v)]_F$. \square

We get the following strengthening of the Free Extension Conjecture.

Conjecture 2.7.10. *Every finite extension problem has a finite, symmetric, ℓ -free solution.*

3 Conclusion

The very first sentence of the introduction of this thesis states that this work is divided into two parts which treat separate topics. For the conclusion we want to do the opposite and discuss in which ways ‘acyclicity and covers’ and ‘extension problems’ are similar.

A common, unifying topic that underlies both concepts is the following question: can we approximate a free object by finite objects?

Simply connected unbranched covers are the free objects in Chapter 1. We showed in Proposition 1.6.10 that these are generally not approximable by finite means.

In Chapter 2 free solutions of extension problems take the role of the free objects. We considered properties that emulate certain behaviours of free solutions and showed that these can be achieved in finite solutions.

Another unifying aspect is the appearance of Otto’s Conjecture (see Conjecture 1.6.22 for its basic version and Conjecture 2.5.16 for its symmetric version). As shown, we can reduce the existence of finite, ℓ -acyclic branched covers (Section 1.6.2) and also to show the existence of finite, ℓ -free solutions (Section 2.5.3) to Conjecture 2.5.16.

Actually, we can connect ‘acyclicity and covers’ and ‘extension problems’ on a much more basic level. We have seen in the proof of Lemma 1.3.13 how M. Hall’s Theorem (Theorem 1.3.5) directly entails the existence of finite, highly acyclic graph covers (if we consider the Galois connection to be an elementary fact). One can also use M. Hall’s Theorem to show the existence of solutions of plain extension problems.¹

Problem 1. Can we relate ‘acyclic covers’ and ‘extension problems’ on other levels as well?

We now treat each chapter individually again. We state four open problems, two for each chapter, which we find noteworthy.

Open questions about covers of hypergraphs

We already mentioned that we cannot approximate simply connected unbranched covers of hypergraphs by finite covers in general. The only examples of finite hypergraphs for which we know that these approximations do not exist are those with undecidable word problem (see the proof of Proposition 1.6.10). So naturally we ask whether there are finite hypergraphs whose word problem is decidable but their

¹This might be a surprising statement as solving plain extension problems is essentially trivial, but so is the Theorem of M. Hall if considered from the right viewpoint.

3 Conclusion

simply connected covers cannot be approximated by finite covers. We know that the word problem for apex acyclic hypergraphs is decidable (Lemma 1.5.12) but we do not know if the simple connected cover can be approximated in the finite. Using Theorem 1.5.9, we can phrase this problem as follows.

Problem 2. Has every finite apex acyclic hypergraph a finite, ℓ -acyclic unbranched cover of arbitrary high degree $\ell \geq 3$?

The second problem is more open in nature. We note in Section 1.6.1 that simple connectivity is the strongest property that can be achieved in unbranched covers. In Section 1.6.2 we note that this connection is not true for hypergraph acyclicity and branched covers. However, the counter example provided is rather unnatural and certainly not an acyclicity notion. Can we find reasonable properties that are necessary for a hypergraph property to be an acyclicity notion? An answer to the following problem would be ideal.

Problem 3. Is hypergraph acyclicity (α -acyclicity) the strongest ‘acyclicity notion’ that can be achieved in certain covers?

Open questions about extension problems

We discussed two methods for solving extension problems. If Otto’s Conjecture were true we would obtain the most general result about extension problem the correctness of the Free Extension Conjecture. It would be interesting to know how far we can get by the method of Auinger and Steinberg. In particular, it seems quite reasonable that the following question has a positive answer (we also suspect that such a proof has the same format as the proof of Lemma 1.4.13).²

Problem 4. Is it possible to prove the Theorem of Herwig and Lascar by the construction method of Auinger and Steinberg?

It seems unlikely that we can produce ℓ -free solutions with the method of Auinger and Steinberg. On an abstract level the method of Auinger and Steinberg produces a sequence of groups G_n s.t. if some u_i are ‘good’ for G_{n+1} then we can find extensions v_i that behave freely and are ‘good’ for the ‘lower level’ G_n . However, ℓ -free solutions G do not have this descending behaviour: if we freely extend elements ‘good’ that are ‘good’ for G than extensions are also ‘good’ for G .

Since we suspect that the Herwig-Lascar result can be proven by the Auinger-Steinberg method, we also suspect that the Free Extension Conjecture is substantially stronger than the Herwig-Lascar Theorem. However, we do not know of a natural extension problem that exemplifies this difference.

Problem 5. Can we establish EPPA for some natural class C of structures by the Free Extension Conjecture which cannot be obtained by the Herwig-Lascar Theorem?

²It appears that such a method is applied in [4]

Open question about coherent solutions

In a comment after the statement of the Free Extension Conjecture (Conjecture 2.1.11) we said that this would be the strongest result so far except for results about coherent solutions. We briefly discuss coherence.

Coherent solutions are introduced by Solecki [43]. A solution \mathcal{S} of a P -extension problem \mathcal{X} is coherent if equations of the partial bijections on \mathcal{X} are reflected in \mathcal{S} .

A pair (f, g) of partial bijections of X are coherent if $\text{img}(f) = \text{dom}(g)$. A word $u = p_1 \dots p_n \in P$ is coherent w.r.t. \mathcal{X} if all $(p_i^{\mathcal{X}}, p_{i+1}^{\mathcal{X}})$ are coherent. A solution \mathcal{S} is coherent if

$$u^{\mathcal{X}} = v^{\mathcal{X}} \implies u^{\mathcal{S}} = v^{\mathcal{S}}$$

for all $u, v \in P^*$ that are coherent w.r.t. \mathcal{X} . We can give an equivalent definition.

Definition 3.0.1. A solution \mathcal{S} of \mathcal{X} is coherent if

$$u^{\mathcal{X}} \leq \text{id} \implies u^{\mathcal{S}} = \text{id}$$

for all $u \in P^*$ coherent w.r.t. \mathcal{X} .

In [43] it is shown how the proof of the Herwig-Lascar Theorem can be adapted such that it yields solutions that are also coherent. By Lemma 2.3.10 this shows the existence of finite, symmetric, cluster ℓ -free solutions. In [41] a variant of the theorem of Hodkinson and Otto (Theorem 2.1.8) in this sense is proven as well.

However, the constructions of Auinger and Steinberg and of Otto both do not preserve coherence of solutions. Even simple constructions of solutions for simple extension problems (such as Lemma 2.7.3) do not produce coherent solutions.

Problem 6. If the Free Extension Conjecture is true, can it even be strengthened to yield coherent solutions?

We want to point out that coherence is in most cases diametral to freeness.

Lemma 3.0.2. *The free solution of a P -extension problem \mathcal{X} is coherent if, and only if, all $p^{\mathcal{X}}$ are total or there are no non-trivial $u \in P^*$ that are coherent w.r.t. \mathcal{X} with $u^{\mathcal{X}} \leq 1$.*

Proof. Let \mathcal{U} be the free solution of \mathcal{X} .

Clearly, if one of the two conditions above are true, then \mathcal{U} is coherent.

So assume that not all $p^{\mathcal{X}}$ are total and that there is a non-trivial $u \in P^*$ coherent w.r.t. \mathcal{X} with $u^{\mathcal{X}} \subseteq \text{id}$. Let u be minimal, i.e., for no subword u' we have $(u')^{\mathcal{X}} \subseteq \text{id}$. Let $q \in P$ and $x \notin \text{dom}(q^{\mathcal{X}})$. Let $a = xq^{\mathcal{U}}$. We show that $au^{\mathcal{S}} \neq a$ and thus that \mathcal{U} is not coherent. If $au^{\mathcal{U}} = a$ then $x(pup^{-1})^{\mathcal{U}} = x$. So, since \mathcal{U} is free, $x[pup^{-1}]_{\text{FG}(P)}^{\mathcal{X}} = x$. Since $x \notin p^{\mathcal{X}}$ we have $u = p^{-1}u'p$. Then either u was not minimal or trivial. \square

Bibliography

- [1] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of databases*. Addison-Wesley Publishing Company, Inc., 1995.
- [2] A. Aranda, D. Bradley-Williams, E. K. Hng, J. Hubička, M. Karamanlis, M. Kompatscher, M. Konečný, and M. Pawliuk. Completing graphs to metric spaces. *Electronic Notes in Discrete Mathematics*, 61:53 – 60, 2017. The European Conference on Combinatorics, Graph Theory and Applications (EUROCOMB’17).
- [3] E. Artin. *Geometric algebra*, volume 3 of *Interscience Tracts in Pure and Applied Mathematics*. Interscience publishers, 1957.
- [4] C. J. Ash. Inevitable graphs: A proof of the type II conjecture and some related decision procedures. *International Journal of Algebra and Computation*, 01(01):127–146, 1991.
- [5] K. Auinger and B. Steinberg. On the extension problem for partial permutations. *Proc. Amer. Math. Soc.*, 131(9):2693 – 2703, 2003.
- [6] K. Auinger and B. Steinberg. A constructive version of the ribes-zalesskiproduct theorem. *Mathematische Zeitschrift*, 250(2):287–297, 2005.
- [7] K. Auinger and M. Szendrei. On F-inverse covers of inverse monoids. *Journal of Pure and Applied Algebra*, 204(3):493 – 506, 2006.
- [8] C. Beeri, R. Fagin, D. Maier, and M. Yannakakis. On the desirability of acyclic database schemes. *J. ACM*, 30(3):479–513, 1983.
- [9] W. W. Boone. The word problem. *Annals of Mathematics*, 70(2):207–265, 1959.
- [10] J. Brault-Baron. Hypergraph acyclicity revisited. *ACM Comput. Surv.*, 49(3):54:1–54:26, 2016.
- [11] G. Conant. *Model Theory and Combinatorics of Homogeneous Metric Spaces*. PhD thesis, University of Illinois at Chicago, 2015.
- [12] A. Dawar and M. Otto. Modal characterisation theorems over special classes of frames. In *Proceedings of 20th IEEE Symposium on Logic in Computer Science LICS’05*, pages 21–30, 2005.
- [13] R. H. Fox. Covering spaces with singularities. In *Symposium in Honor of S. Lefschetz*, page 243–257. Princeton University Press, 1957.

Bibliography

- [14] R. Fraïsse. Sur certaines relations qui généralisent l'ordre des nombres rationnels. *C. R. Acad. Sci.*, 52(237):540–542, 1953.
- [15] M. H. Graham. On the universal relation. Technical report, University of Toronto, 1979.
- [16] M. Hall. Coset representations in free groups. *Trans. Amer. Math. Soc.*, 67(2):421–432, 1949.
- [17] A. Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.
- [18] K. Henckell, S. Margolis, J.-E. Pin, and J. Rhodes. Ash's type II theorem, profinite topology and Malcev products. *International Journal of Algebra and Computation*, 1:411–436, 1991.
- [19] K. Henckell and J. Rhodes. The theorem of Knast, the $PG = BG$ and type-II conjectures. In J. Rhodes, editor, *Monoids and Semigroups with Applications*, pages 453–463. World Scientific, 1991.
- [20] B. Herwig. Extending partial isomorphisms on finite structures. *Combinatorica*, 15(3):365–371, 1995.
- [21] B. Herwig and D. Lascar. Extending partial automorphisms and the profinite topology on free groups. *Trans. Amer. Math. Soc.*, 352:1985–2021, 2000.
- [22] G. Higman, B. H. Neumann, and H. Neuman. Embedding theorems for groups. *Journal of the London Mathematical Society*, s1-24(4):247–254, 1949.
- [23] W. Hodges, I. Hodkinson, D. Lascar, and S. Shelah. The small index property for ω -stable ω -categorical structures and for the random graph. *Journal of the London Mathematical Society*, s2-48(2):204–218, 1993.
- [24] I. Hodkinson and M. Otto. Finite conformal hypergraph covers and Gaifman cliques in finite structures. *The Bulletin of Symbolic Logic*, 9(3):387–405, 2003.
- [25] E. Hrushovski. Extending partial isomorphisms of graphs. *Combinatorica*, 12(4):411–416, 1992.
- [26] M. V. Lawson. *Inverse semigroups*. World Scientific Publishing, 1998.
- [27] M. V. Lawson, S. W. Margolis, and B. Steinberg. Expansions of inverse semigroups. *Journal of the Australian Mathematical Society*, 80(2):205–228, 2006.
- [28] C. Löh. *Geometric Group Theory*. Universitext. Springer, 2017.
- [29] D. Macpherson. A survey of homogeneous structures. *Discrete Mathematics*, 311(15):1599 – 1634, 2011.

- [30] S. Margolis, M. Sapir, and P. Weil. Closed subgroups in pro-V topologies and the extension problem for inverse automata. *International Journal of Algebra and Computation*, 11(04):405–445, 2001.
- [31] D. B. McAlister. Groups, semilattices and inverse semigroups. *Trans. Amer. Math. Soc.*, 192(192):227–244, 1974.
- [32] J. W. Moon. Tournaments with a given automorphism group. *Canad. J. Math.*, 16:485–489, 1964.
- [33] P. S. Novikov. On the algorithmic unsolvability of the word problem in group theory. *Trudy Mat. Inst. Steklov.*, 44:3–143, 1952. (in Russian).
- [34] M. Otto. Highly acyclic groups, hypergraph covers and the guarded fragment. *Journal of the ACM*, 59, 2012.
- [35] M. Otto. Finite groupoids, finite coverings and symmetries in finite structures. *ArXiv e-prints*, 2014. arXiv:1404.4599.
- [36] M. Otto. Amalgamation and symmetry: From local to global consistency in the finite. *arXiv e-prints*, 2017. arXiv:1709.00031.
- [37] V. G. Pestov. A theorem of Hrushovski-Solecki-Vershik applied to uniform and coarse embeddings of the Urysohn metric space. *Topology and its Applications*, 155(14):1561 – 1575, 2008. Special Issue: Workshop on the Urysohn space.
- [38] E. Rosen. Modal logic over finite structures. *Journal of Logic, Language and Information*, 6:95–27, 1995.
- [39] C. Rosendal. Finitely approximable groups and actions part I: The Ribes-Zalueskiĭ property. *J. Symbolic Logic*, 76(4):1297–1306, 2011.
- [40] I. Singer and J. Thorpe. *Lecture Notes on Elementary Topology and Geometry*. Undergraduate Texts in Mathematics. Springer, 1976.
- [41] D. Siniora and S. Solecki. Coherent extension of partial automorphisms, free amalgamation, and automorphism groups. *ArXiv e-prints*, 2017. arXiv:1705.01888.
- [42] S. Solecki. Extending partial isometries. *Israel Journal of Mathematics*, 150(1):315–331, 2005.
- [43] S. Solecki. Notes on a strengthening of the Herwig-Lascar extension theorem, 2009.
- [44] J. R. Stallings. Topology of finite graphs. *Inventiones mathematicae*, 71:551–566, 1983.
- [45] J. Stillwell. *Classical Topology and Combinatorial Group Theory*. Graduate Texts in Mathematics. Springer, 1993.

Bibliography

- [46] N. Szakács. On the graph condition regarding the F-inverse cover problem. *Semigroup Forum*, 92(3):551–558, 2016.
- [47] N. Szakács and M. B. Szendrei. On F-inverse covers of finite-above inverse monoids. *Journal of Algebra*, 452(Supplement C):42 – 65, 2016.
- [48] J. K. Truss. Generic automorphisms of homogeneous structures. *Proceedings of the London Mathematical Society*, s3-65(1):121–141, 1992.
- [49] J. F. A. K. van Benthem. *Modal Logic and Classical Logic*. Distributed in the U.S.A. By Humanities Press, 1983.
- [50] A. Vershik. Globalization of the partial isometries of metric spaces and local approximation of the group of isometries of Urysohn space. *Topology and its Applications*, 155(14):1618 – 1626, 2008. Special Issue: Workshop on the Urysohn space.
- [51] J.-F. Wang. On axioms constituting the foundation of hypergraph theory. *Acta Mathematicae Applicatae Sinica*, 21(3):495–498, 2005.
- [52] E. Witt. Theorie der quadratischen Formen in beliebigen Körpern. *Journal für die reine und angewandte Mathematik*, 176:31–44, 1937.
- [53] C. T. Yu and M. Z. Ozsoyoglu. An algorithm for tree-query membership of a distributed query. In *COMPSAC 79. Proceedings. Computer Software and The IEEE Computer Society's Third International Applications Conference, 1979.*, pages 306–312, 1979.

Wissenschaftlicher Werdegang

1987 *Geboren am* 6. Oktober 1987.

2008–2012 *Technische Universität Darmstadt*, BSc. Mathematics with Computer Science.

2012–2014 *Technische Universität Darmstadt*, MSc. Mathematics.

2014–2019 *Technische Universität Darmstadt*, Wissenschaftlicher Mitarbeiter am Fachbereich Mathematik, Promotion in endlicher Modeltheorie.