Ghent University
Faculty of Sciences
Department of Mathematics:
Algebra and Geometry

Expanding the Universe of (Restricted) Universal Groups over Right-Angled Buildings

Jens Bossaert

October, 2021

Supervisors
prof. dr. Tom De Medts
prof. dr. Hendrik Van Maldeghem

Dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Sciences: Mathematics

Table of contents

> *The man or woman who is rarely lost, rarely discovers anything new.*
>
> — Josiah Bancroft, *Arm of the Sphinx*

Acknowledgements

> *He was swimming in a sea of other people's expectations. Men had drowned in seas like that.*
>
> — Robert Jordan, *New Spring*

It requires a remarkable amount of both intelligence and stupidity to write a dissertation unaided. This PhD thesis would never have seen the light without the help of my supervisors, Tom De Medts and Hendrik Van Maldeghem. Thank you, Tom — for keeping your office door open for questions, for taking the time to read through all my rough ideas, incomplete proofs, and draft papers, and for pulling me back on track whenever I risked tumbling into another black hole. Thank you, Hendrik — for taking me to various conferences and seminars, and for helping me juggle the many approaches to building theory into a coherent framework.

Many thanks to the members of the examination board for reading this dissertation in incredible detail, and for providing pages of helpful feedback, additional insights, and encouraging remarks. Thank you, Pierre-Emmanuel Caprace, Simon Smith, Linus Kramer, Bart De Bruyn, Anneleen De Schepper, and Marnix Van Daele.

Also many thanks to all my friends and colleagues at the office, for the chats during lunch and four o'clock breaks, for the support (in all its forms), and for believing in me far more that I did myself. Thank you, Jozefien Dhaeseleer, Lisa Hernandez Lucas, Toon Baeyens, Jeroen Meulewaeter, Simon Rigby, Maarten De Boeck, Anneleen De Schepper, Magali Victoor, Paulien Jansen …

Thank you, Sara Chiers, for never letting me down whenever research or just life in general turned harsh, for sharing the joys and burdens of routine life, and for building a future together.

Thank you, Samuel Perez and Geert Vernaeve, for being the "invisible forces in the background".

Finally a thanks to my parents, for giving me the opportunity to pursue my studies in mathematics. Thank you, Kathy Timperman and Ronny Bossaert.

— Jens Bossaert,
October 2021

Introduction

> *It's a dangerous business, going out your door. You step onto the road, and if you don't keep your feet, there's no knowing where you might be swept off to.*
>
> — J.R.R. Tolkien, *The Fellowship of the Ring*

"Buildings? So … you're an architect?"

I am quite sure that most mathematicians recognise the inner struggle, when the uninitiated asks the dreaded question what your research is about. You want to make them understand — or at least appreciate — the beauty of the abstract structures and their symmetries. You want to show them illustrations like the one below and explain the construction. You want to grant them even a small glimpse of the satisfaction you feel when laying the final puzzle piece to a new result.

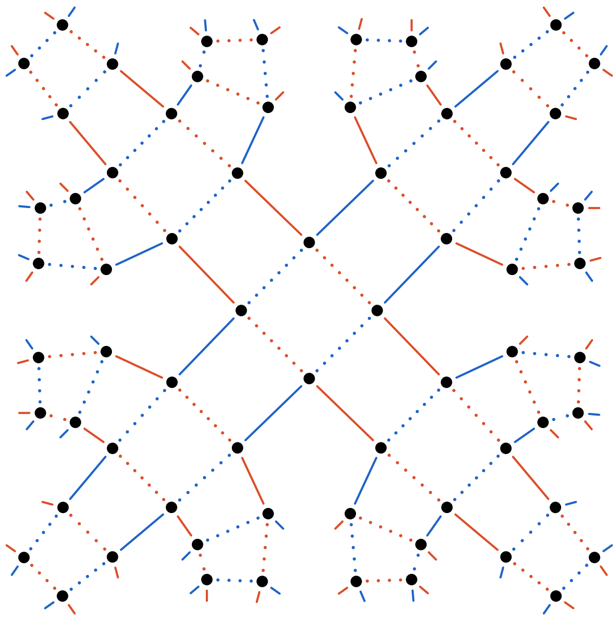

At that moment, you forget the years upon years of necessary background. You forget the frequent moments of frustration when a new approach to your tantalising research problem, again, proves useless. And you forget that, for most people, "buildings" are totally different things from the ones you're used to. (In my case, it probably also doesn't help that my very young self liked drawing — so much that, naturally, I wanted to be an architect when I grew up.)

But indeed: *buildings*. Oddly enough, throughout my mathematical curriculum at Ghent University buildings never quite stole the spotlight, but simultaneously everyone seemed to know about them. Building theory turned out to be a broad subject — encompassing various notions of group theory, incidence geometry, graph theory … into a single elegant framework. Projective and affine spaces are prime examples of buildings, as are infinite trees (a fact that is another foolproof way to invoke a giggle out of the uninitiated).

The first PhD year was mostly spent on getting to grips with this new theory of buildings and the existing literature. My first goal was to generalise results of Ana Silva ([DMdSS18] and [DMdS19]) to universal groups with local groups that are not necessarily finite or transitive — I sketch some history and motivation in the next section — and throughout the second PhD year this went rather smoothly, without major hindrances, except for one disenchanting problem that appeared to be oh so innocent. In fact, I essentially wasted the third year thinking to crack the problem "next week"! Suffice it to say, the disillusion did not work wonders for my motivation. Moreover, the fourth and final PhD year was the cursed year 2020–2021, and we all know what happened then. This thesis is the culmination of those four years, smoothing out the ups and downs of day-to-day research into one coherent document.

The main bulk of this PhD thesis consists of the results of [BDM21], accompanied with more background and written out in more detail. New results are the city products of buildings in Section 2.7, the simplicity theorem in Section 3.4, and the restricted universal groups in Chapter 4. To my regret, the results about restricted universal groups are still a bit shallow, and I mention in Chapter 5 some research questions that I would have liked to explore further.

Framing the research

The main topic of this thesis finds its origin in foundational work by Marc Burger and Shahar Mozes ([BM00a]) on the local versus global structure of groups acting on trees. The groups in question are topological groups, i.e. groups equipped with a compatible topology. Topological group theory is a very broad area on its own and has connections to numerous other mathematical areas and even other sciences. One condition preventing the topological groups from growing uncontrollably wild is *local compactness*.

Locally compact groups are interesting because they embody a huge family of naturally occurring groups. Also, they come with a natural measure called the *Haar measure*, allowing the application of various techniques from the toolbox of mathematical analysis. In general, a locally compact group decomposes into a *connected* group and a *totally disconnected* one. Those two cases can more or less be studied independently from each other.

The structure theory of *connected* locally compact groups can be called well understood since the 1950s, thanks to the solution of Hilbert's fifth problem by Andrew Gleason, Deane Montgomery, Leo Zippin, and Hidehiko Yamabe. They demonstrated that connected locally compact groups can, roughly speaking, be approximated by Lie groups; we mention the precise result in Theorem 1.2.17.

On the other hand, *totally disconnected* locally compact groups are still quite mysterious. In fact, for a long time the only general structural result was a theorem by David van Dantzig ([vD36]), roughly saying that such a group has an abundance of compact open subgroups (Corollary 1.2.19). Note that this result dates back to 1936! More recently in 1994, George Willis introduced in [Wil94] some new tools to analyse totally disconnected locally compact groups, such as the *scale function*. Today, a classification result analoguous to the Gleason–Yamabe theorem seems far out of reach.

Encouraged by van Dantzig's result, several theorems and techniques have been found that relate the global structure of a totally disconnected locally compact group to its compact open subgroups. For example, Willis showed in [Wil07] that if such a group is topologically simple and compactly generated, then its compact open subgroups cannot be solvable. The results by Burger and Mozes in [BM00a] also fall in this category: they studied the local-to-global properties of groups acting on a tree by automorphisms.

In particular, they eventually define a subgroup $\mathcal{U}(F)$ of the automorphism group of a regular tree, obtained by restricting the local actions around every vertex to a prescribed permutation group F (we give a detailed definition in Section 1.5.1). This construction provides a large family of totally disconnected locally compact groups, that are compactly generated and nondiscrete (under very mild conditions). Moreover, these "universal groups" satisfy Tits's independence criterion, and can be shown to have a simple subgroup of index two.

Another "hot topic" in locally compact group theory, are lattices: discrete subgroups $\Gamma \leq G$ such that the quotient space G/Γ is of finite invariant volume. As Burger and Mozes showed, the automorphism groups of direct products of trees are natural habitats of interesting lattices ([BM00b]). We note that Hyman Bass and Alexander Lubotzky found interesting results in the automorphism group of a single tree as well, and refer to [Car02] for a nice overview.

Automorphisms groups of trees have the additional benefit of being very flexible. Simon Smith for example proposed in [Smi17] a construction similar to the Burger–Mozes universal group over a *semiregular* tree, with two local permutation groups. This allowed Smith to construct an uncountable family of nonisomorphic, nondiscrete, compactly generated, totally disconnected, locally compact, simple groups — the existence of which was new. We give some more details in Section 1.5.2.

We also mention results by Adrien Le Boudec, who allowed for a finite number of local exceptions in the Burger–Mozes groups. His construction gives rise to interesting topological groups that do not admit any lattice, which appears to quite be an elusive phenomenon. We refer to Section 1.5.3 or [LB16].

In Section 1.5.4, we touch upon recent results by Waltraud Lederle in [Led17], who defined another variant of the Burger–Mozes group more similar to the Neretin group of spheromorphisms of trees ([BCGM12]). Her construction leads, again, to interesting topological groups that do not admit any lattices.

What mainly makes all of the above constructions valuable is the fact that there is a lot of freedom in the local data, which profoundly impacts the group's global structure.

Following a suggestion of Pierre-Emmanuel Caprace and under supervision of Tom De Medts, Ana Silva studied in her PhD thesis a generalisation of the original universal group of Burger and Mozes to *right-angled buildings*, of which trees are prototypical examples. The main reason why universal groups over trees "work" is the property that a permutation of the edges in any star can be extended to a full tree automorphism, i.e. the local data is retrievable at the global scale. Buildings in general do not enjoy this property: it is not possible to, say, extend any arbitrary permutation of the points on a line in a projective space to a full automorphism. *Right-angled* buildings, however, are a general class of buildings that do have this exceptional local-to-global property.

In this thesis we continue the work of Silva, who focused on the locally finite case with transitive local groups. We further develop the universal groups without these restrictions, and introduce a generalisation of the results of Smith and Le Boudec to this building setting. A recurring theme in this setting will be that the structure of the resulting groups not only depend on the permutational properties of the local groups, but also on the combinatorial properties of the building's diagram.

Overview of the results

We start with an introductory chapter with preliminaries in abstract and topological group theory, a refresher on graph theory, general results about automorphisms of trees, and a brief sketch of the results of Burger–Mozes, Smith, Le Boudec, and Lederle, before we introduce buildings in general.

In the next chapter we focus solely on right-angled buildings and establish some auxiliary lemmas. In that chapter, we define the concept of an *implosion* of a building: a controlled way of collapsing a building in a completely different way from projections or retractions, but by collapsing the local structure. We also introduce an operation that we call a *city product*, allowing us to glue together buildings along another building, and we establish some properties of this product.

In Chapter 3 we finally define the universal groups over right-angled buildings. We need a technical lemma for extending any partial automorphism to an automorphism “as close to a universal group element as possible” (Proposition 3.1.9). We then characterise in terms of the local groups and the diagram when the universal group is transitive, and more interestingly, primitive (Theorem 3.2.15, generalising a result of Smith). Here we can already clearly see how the diagram’s combinatorial structure affects the group’s global structure. We proceed to characterise when the universal group is generated by point stabilisers, again in terms of the local data and the diagram (Theorem 3.2.20). We proceed to endow both the local groups and the universal group with the permutation topology and study how the local and global topological structure relate. In particular, we present sufficient conditions under which the universal group is compactly generated. We conjecture these conditions to be necessary as well and present some motivational partial results. Next, we establish a general simplicity criterion, and use it to show the subgroup of the universal groups generated by chamber stabilisers to be simple. Finally, we study our city products again, and describe the universal group over a city product as a universal group of universal groups over the factor buildings.

In Chapter 4, we study an analogue of the Le Boudec groups to right-angled buildings. We quickly find restrictions on where the exceptional local actions can occur, and properties that motivate to restrict to locally finite buildings again. We endow the restricted universal groups with a topology that allows us to transfer properties of the universal groups to this setting. Finally, we generalise a characterisation of Pierre-Emmanuel Caprace, Colin Reid, and Phillip Wesolek, and describe when the groups are virtually simple.

We finish this thesis with some open questions in Chapter 5.

1 Preliminaries

> *Knowing your own ignorance is the first step to enlightenment.*
>
> — Patrick Rothfuss, *The Wise Man's Fear*

1.1 Abstract group theory

First and foremost, we cannot give a concise introduction to the vast area that is modern (abstract) group theory, but we will recall some fundamental concepts. We shall assume familiarity with these notions, and mainly take the opportunity to establish some notational conventions that we will use throughout this thesis. We refer to Page 141 for an overview of notations.

Let $(G, \cdot)$ be a group. We will always use multiplicative notation, and usually, unless the situation calls for explicit notation, we will abbreviate $g \cdot h$ as gh. The identity element will be denoted by id and the inverse of a group element g by g^{-1}. We denote conjugation of an element g by h by

$$h \cdot g \cdot h^{-1} = {}^{h}g,$$

which is a slightly awkward notation, but that is the price to pay if we want that ${}^{h_2}({}^{h_1}g) = {}^{h_2 h_1}g$. The commutator of two elements g and h will be denoted by

$$g \cdot h \cdot g^{-1} \cdot h^{-1} = [g, h].$$

Next, let G act on a set X. By convention, our groups always act on the left, and we denote by $g \mathbin{.} x$ the image of an element $x \in X$ under $g \in G$. Consequently, when G is a group of permutations, the composition $g \cdot h$ means the permutation obtained by first applying h, and then g. Then for all $g, h \in G$ and all $x \in X$,

$$(g \cdot h) \mathbin{.} x = g \mathbin{.} (h \mathbin{.} x).$$

Occasionally we will also use other common notations like gx or $g(x)$, but we will try to keep this to a minimum — we dismiss the former to avoid confusion with the group operation, and the latter to avoid an avalanche of parentheses. Extending our group action notation, the G-orbit of x will be denoted by $G \mathbin{.} x$.

For the pointwise stabiliser of a subset $Y \subseteq X$, we use the notation

$$G_{(Y)} = \{ g \in G \mid g \mathbin{.} y = y \text{ for all } y \in Y \},$$

and for the setwise stabiliser,

$$G_{\{Y\}} = \{ g \in G \mid g \mathbin{.} y \in Y \text{ for all } y \in Y \}.$$

Whenever $Y = \{y\}$, we will simply write the stabiliser as G_y since both notions obviously agree. However, in order to avoid potential confusion we will never use an abbreviation like G_Y when Y is not a singleton.

Definition 1.1.1 (transversal). A *transversal* for the action of G on a set X is any subset $\Upsilon \subseteq X$ that contains exactly one element of every orbit, or in other words, any subset $\Upsilon \subseteq X$ that makes the natural map $\Upsilon \to X/G \colon x \mapsto G\,.\,x$ to the orbit space bijective.

Definition 1.1.2. Let G be a group acting on a set X. We define the subgroup $G^+ = \langle G_x \mid x \in X\rangle$ generated by all point stabilisers.

Let us recall some familiar general properties of group actions.

Definition 1.1.3. Let G be a group acting on a set X. We say that the action is …

- *transitive* if G has only one single orbit on X;
- *faithful* if the intersection of all point stabilisers is trivial;
- *free* or *semiregular* if every point stabiliser is trivial;
- *regular* if it is both transitive and free;
- *primitive* if it is transitive and leaves no partition of X invariant, except the trivial partitions into one single class and into singletons.

Later on we will encounter other properties for groups acting on certain geometric structures, such as Definitions 1.4.1, 1.6.7, and 2.4.4. Additionally we will use the renowned primitivity criterion by Higman ([Hig67]). Here we recall some definitions and state the criterion.

Definition 1.1.4 (suborbit). Let G be acting on a set X. Then the orbits of a point stabiliser G_x are called *suborbits* of G.

Definition 1.1.5 (orbital). Let G be acting on a set X. Then G induces an action on $X \times X$, by $g\,.\,(x, y) = (g\,.\,x, g\,.\,y)$. The orbits of this action are called *orbitals* of G. To every orbital $G\,.\,(x, y)$ corresponds an *orbital graph*: the graph with vertex set X and with edge set

$$\big\{\{u, v\} \mid (u, v) \in G\,.\,(x, y) \text{ or } (v, u) \in G\,.\,(x, y)\big\}$$

(see Section 1.3 for the definition of a graph).

Suppose that G is transitive. There is one trivial or *diagonal* orbital $\{(x, x) \mid x \in X\}$, whose orbital graph simply has a loop at each vertex. No other orbital graph has loops.

Lemma 1.1.6. *Let G be acting transitively on a set X and let $x \in X$. Then there is a natural bijection between the orbits of G_x and the orbitals of G; the trivial suborbit corresponds to the diagonal orbital.*

Proof. Define

$$X \times X \to X \colon Y \to Y(x) = \{y \in X \mid (x, y) \in Y\}.$$

In particular, let Y now be an orbital of G. Then $Y(x)$ is nonempty, since the action is transitive. Let $y_1, y_2 \in Y(x)$. There exists some $g \in G$ such that $g\,.\,(x, y_1) = (x, y_2)$. It follows that $g \in G_x$ and hence $Y(x)$ is in fact an orbit of G_x. ∎

Theorem 1.1.7 (Higman). *Let G be acting transitively on a set X. Then G acts primitively on X if and only if every nondiagonal orbital graph is connected.*

Proof omitted. We refer to [Hig67, (1.12)].

We will need the following folklore result as well.

Lemma 1.1.8. *Let G be a primitive nonregular permutation group on X. Let $x, y \in X$ be two distinct elements. Then there exists a permutation $g \in G$ such that $g\,.\,x = x$ but $g\,.\,y \neq y$.*

Proof. Suppose by means of contradiction that $G_x \subseteq G_y$. Since G is primitive, G_x is a maximal subgroup, hence $G_x = G_y$. Let $g \in G$ be such that $g \,.\, x = y$. Then $g \notin G_x$ but g normalises G_x. We can hence write

$$G_x \lneq N_G(G_x) \trianglelefteq G.$$

Again, as G_x is maximal, it follows that $G_x \trianglelefteq G$. But then all point stabilisers are equal to G_x, so that G_x fixes all points of X, and G is regular — a contradiction. ∎

The next few properties and propositions involve finite index subgroups.

Definition 1.1.9 (virtually (∗)). For any property $(*)$ of groups, we say that G is *virtually* $(*)$ if there is a subgroup of finite index that satisfies property $(*)$.

Lemma 1.1.10. *Let G be a group and let $H \leq G$ be a subgroup of finite index. Then there is a normal subgroup $N \trianglelefteq G$ of finite index such that $N \leq H$.*

Proof. The action of G on the coset space G/H by left multiplication induces a permutation representation $\rho\colon G \to \mathrm{Sym}(n)$, where $n = [G : H]$. Let $N = \ker(\rho)$ be its kernel. Then indeed $N \trianglelefteq G$ and $N \leq H$, while G/N is isomorphic to a subgroup of $\mathrm{Sym}(n)$, so that $[G : N]$ is finite. ∎

Lemma 1.1.11. *A simple subgroup H of finite index in an infinite group G is a normal subgroup.*

Proof. By Lemma 1.1.10 there exists some normal subgroup $N \trianglelefteq G$ of finite index, contained in H. Then $N \trianglelefteq H$ must either be trivial (and G finite) or $N = H$. ∎

Definition 1.1.12 (monolithic). A group is called *monolithic* if the intersection of all its nontrivial normal subgroups is nontrivial. In other words, a monolithic group has a unique minimal normal subgroup, which is then called its *monolith.*

Definition 1.1.13 (commensurate). Two subgroups H_1 and H_2 of a group G are *commensurate* if their intersection is of finite index in both H_1 and H_2.

Lemma 1.1.14. *Commensuration is an equivalence relation on subgroups.*

Proof. Clearly, commensuration is reflexive and symmetric. In order to show that it is transitive as well, let $H_1, H_2, H_3 \leq G$ be subgroups such that H_1 and H_2 are commensurate and such that H_2 and H_3 are commensurate. Then

$$\begin{aligned}
\big[H_1 : H_1 \cap H_3\big] &\leq \big[H_1 : H_1 \cap H_2 \cap H_3\big] \\
&= \big[H_1 : H_1 \cap H_2\big] \cdot \big[H_1 \cap H_2 : H_1 \cap H_2 \cap H_3\big] \\
&= \big[H_1 : H_1 \cap H_2\big] \cdot \big[H_2 : H_1 \cap H_2 \cap H_3\big] \cdot \big[H_2 : H_1 \cap H_2\big]^{-1} \\
&\leq \big[H_1 : H_1 \cap H_2\big] \cdot \big[H_2 : H_1 \cap H_2\big] \cdot \big[H_2 : H_2 \cap H_3\big] \cdot \big[H_2 : H_1 \cap H_2\big]^{-1} \\
&= \big[H_1 : H_1 \cap H_2\big] \cdot \big[H_2 : H_2 \cap H_3\big],
\end{aligned}$$

which is finite. Similarly $\big[H_3 : H_1 \cap H_3\big]$ can be shown to be finite. ∎

Definition 1.1.15 (commensurator). The *commensurator* of a subgroup $H \leq G$ is the subgroup of elements of G that conjugate H to a subgroup commensurate with H, or explicitly,

$$\mathrm{Comm}_G(H) = \big\{ g \in G \;\big|\; {}^gH \cap H \text{ has finite index in both } {}^gH \text{ and } H \big\}.$$

By Lemma 1.1.14 and the observation that conjugation preserves commensuration, it follows that the commensurator $\mathrm{Comm}_G(H)$ is indeed a subgroup of G. Also note that $\mathrm{Comm}_G(H)$ contains in particular the normaliser of H.

We finish this introductory section with a definition for permutation groups (over any, not necessarily finite set X).

Definition 1.1.16 (Young subgroup). To every partition of a set X, we can associate a subgroup of $\mathrm{Sym}(X)$ of all permutations stabilising the blocks of the partition. A subgroup obtained in this fashion is called a *Young subgroup* of $\mathrm{Sym}(X)$, and is naturally isomorphic to the direct product of the symmetric groups on the blocks.

For any permutation group $G \leq \mathrm{Sym}(X)$, we have a canonical partition of X into G-orbitss. We call the Young subgroup associated to this partition the *Young overgroup of* G and denote it by $\widehat{G}$.

Note that we indeed always have the inclusions $G \leq \widehat{G} \leq \mathrm{Sym}(\Omega)$. Moreover, $G = \widehat{G}$ if and only if G is itself a Young subgroup, and $\widehat{G} = \mathrm{Sym}(\Omega)$ if and only if G is transitive.

1.2 Topological group theory

1.2.1 General topological groups

The groups that we will encounter will be interesting examples of topological groups, i.e. groups endowed with an additional topological structure that is compatible with the algebraic structure. We shall give a brief overview of the theory, mainly devoting our attention to topological groups that are totally disconnected and locally compact. For more details, we refer to one of the standard works such as [HR13]; in addition the lecture notes by Dikran Dikranjan [Dik18], by Linus Kramer [Kra20], and by Phillip Wesolek [Wes18] may very well prove to be helpful and accessible sources.

Definition 1.2.1 (topological group). A *topological group* $(G, \,\cdot\, , \tau)$ is a set G equipped with a group operation $\cdot$ and a topology τ, such that the two maps $G \times G \to G\colon (g, h) \mapsto g \cdot h$ and $G \to G\colon g \mapsto g^{-1}$ are continuous.

Note that we do not require topological groups to be *Hausdorff.* Recall that a topological space is Hausdorff if for every pair (x, y) of distinct points, there exist disjoint open sets U_x and U_y such that $x \in U_x$ and $y \in U_y$.

Just as we usually abbreviate a group $(G, \,\cdot\,)$ simply as G, we shall do the same with a topological group $(G, \,\cdot\, , \tau)$ when the topology is clear from the context. The underlying group of a topological group will be called its *abstract group.*

The continuity of the translation maps $g \mapsto g \cdot h$ and $g \mapsto h \cdot g$ (for fixed $h \in G$) intuitively means that a topological group “looks the same in every point”, that all topological behaviour is captured by a single point’s neighbourhood basis. We now make this precise.

Definition 1.2.2. Let G be a topological group.

- A subset $U \subseteq G$ is called a *neighbourhood* of a point $g \in G$ if there is some open subset V such that $g \in V \subseteq U$.
- A set $\mathfrak{B}$ of neighbourhoods of a point $g \in G$ is called a *neighbourhood basis for* g if for every neighbourhood U of g, there is some $V \in \mathfrak{B}$ such that $V \subseteq U$. A neighbourhood basis for the identity will also be called an *identity neighbourhood basis.*

Note that if U is a neighbourhood of g, then $h \cdot U$ is a neighbourhood of $h \cdot g$.

Proposition 1.2.3. *The topology of a topological group is completely determined by an identity neighbourhood basis.*

Proof omitted. We refer to [HR13, Theorem 4.5].

An identity neighbourhood basis on a topological group is a special instance of a *filter basis*. More precisely, a neighbourhood basis can be defined as a filter basis of the neighbourhood filter; though we will not need filters in general, we will mention a useful property that allows to define a group topology by means of an (abstract) filter basis.

Definition 1.2.4 (filter basis). Let X be a set. A nonempty collection $\mathfrak{B}$ of subsets of X is called a *filter basis* on X if the intersection of any two sets in $\mathfrak{B}$ is a superset of some set in $\mathfrak{B}$.

Lemma 1.2.5. *Let $\mathfrak{B}$ be a filter basis on a group G. Assume that $\mathfrak{B}$ satisfies the following properties:*

(i) *for every $U \in \mathfrak{B}$, there exists $V \in \mathfrak{B}$ such that $V \cdot V \subseteq U$;*

(ii) *for every $U \in \mathfrak{B}$, there exists $V \in \mathfrak{B}$ such that $V^{-1} \subseteq U$;*

(iii) *for every $U \in \mathfrak{B}$ and every $g \in G$, there exists $V \in \mathfrak{B}$ such that $V \subseteq {}^{g}U$.*

Then G admits a unique group topology such that $\mathfrak{B}$ is a neighbourhood basis of the identity of G.

Proof omitted. We refer to [Bou07, Proposition 1], or to be precise, the subsequent paragraph.

Example 1.2.6. (i) Every group can be interpreted as a topological group in a trivial way, by endowing it with the discrete topology (i.e. by declaring every subset to be open). We refer to such a group as a *discrete group*.

The discrete topology is the only possible Hausdorff topology on a finite group.

(ii) Every group can be made into a topological group by taking as an identity neighbourhood basis, the set of all normal subgroups of finite index. This *profinite topology* is Hausdorff if and only if the intersection of all normal subgroups of finite index is trivial. Such a group is called *residually finite*.

(iii) The metric structure of the Euclidean space $\mathbb{R}^n$ induces a canonical topology, defined by the open balls as basic open sets. This topology makes the additive group $\mathbb{R}^n$ into a topological group. A neighbourhood basis of the identity is given by the family of open balls

$$B_r(\mathbf{0}) = \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x}\| < r\},$$

where r ranges over the positive real numbers.

(iv) The circle group $\mathbb{T} = \{z \in \mathbb{C} \mid |z| = 1\} \leq \mathbb{C}^\times$ with the subspace topology inherited from $\mathbb{C}$, is a compact topological group.

(v) If G is a topological group and N a normal subgroup, then the quotient G/N is a topological group with the quotient topology. It is Hausdorff if and only if N is closed in G.

Here, $\mathbb{R}^n$ and $\mathbb{T}$ are examples of *Lie groups*. These are topological groups with far more exacting geometrical structure, namely that of a differentiable manifold, such that the group operation and inversion are smooth maps instead of merely continuous. Lie groups play a key role in the study of connected locally compact groups (see Theorem 1.2.17).

Let us mention some basic but important properties of product sets and subgroups.

Proposition 1.2.7. *Let A and B be two subsets of a topological group.*

(i) *If A or B is open (or both), then $A \cdot B$ and $B \cdot A$ are open as well.*

(ii) *If both A and B are compact, then $A \cdot B$ and $B \cdot A$ are compact as well.*

(iii) *If A is closed and B is compact, then $A \cdot B$ and $B \cdot A$ are closed as well.*

(iv) *If both A and B are closed, then $A \cdot B$ and $B \cdot A$ need not be closed.*

Proof omitted. This is [HR13, Theorem 4.4].

Proposition 1.2.8. *Let H be a subgroup of a topological group G.*

(i) *H is open if and only if H contains a nonempty open subset of G.*

(ii) *If H is open in G, then H is closed as well.*

(iii) *If H is closed and of finite index in G, then H is open as well.*

Proof. (i) First, if H contains a nonempty open subset U, then $H = U \cdot H$ is open by Proposition 1.2.7 (i). Conversely, if H is open, then H contains the nonempty open subset H.

(ii) If H is open, then so are the cosets of H. The complement $G \setminus H = \bigcup\{gH \mid g \notin H\}$ is a union of open sets, hence open, and H is closed.

(iii) If H is closed, then so are the cosets of H. The complement $G \setminus H = \bigcup\{gH \mid g \notin H\}$ is a *finite* union of closed sets, hence closed, and H is open. ■

When working with topological groups, the notions of morphisms, simple groups, … are more restrictive, as we need to take into account the topological structure as well.

Definition 1.2.9 (morphism). A map $\varphi\colon G \to H$ between topological groups is a *morphism of topological groups* is both a group morphism and a continuous map.

Since this is quite verbose, if the context is clear, usually we simply say that φ is a *morphism*. If we want to emphasise that a morphism is not necessarily continuous, we call it an *abstract morphism*.

Definition 1.2.10 (simplicity). We call a topological group *topologically simple* if it has no proper nontrivial closed normal subgroups.

A topologically simple group may very well have proper nontrivial *nonclosed* normal subgroups, but then the corresponding quotient groups will not be Hausdorff (in fact, not even Kolmogorov).

Let us recall a couple of notions from general point-set topology that are of primordial importance for topological groups.

Definition 1.2.11. Let X be a topological space. Then X is called

- *compact* if every open cover of X has a finite subcover;
- *locally compact* or *l.c.* if every point of X has a compact neighbourhood;
- *connected* if X is not the union of two nonempty disjoint open sets (i.e. no proper subset is both open and closed);
- *totally disconnected* or *t.d.* if all connected subsets are singletons.
- *totally separated* if for every two distinct points $x, y \in X$ there exists a partition of X into two open subsets U_x and U_y such that $x \in U_x$ and $y \in U_y$.

Totally disconnected locally compact groups will often be abbreviated as t.d.l.c. groups. (Evidently, whenever we call a topological group compact / locally compact / …, we mean that its underlying topological space is compact / locally compact / …)

Proposition 1.2.12. *Let G° be the connected component of the identity of a topological group G, i.e. the maximal subset containing the identity that is connected (in the subspace topology). Then G° is a closed normal subgroup of G, and G/G° is a totally disconnected group. In other words, we have a short exact sequence*

$$1 \to G^\circ \to G \to G/G^\circ \to 1,$$

and every topological group is an extension of a totally disconnected one by a connected one.

Proof omitted. We refer to [HR13, Theorems 7.1, 7.3].

Definition 1.2.13. Let G be a topological group. Then G is called *compactly generated* if there is a compact subset $K \subseteq G$ that algebraically generates G. (This should not be confused with the homonymic, but unrelated, notion for general topological spaces without a group structure.)

The last proposition of general topological groups that we mention, links topological properties of subgroups to commensuration.

Proposition 1.2.14. *Any two open compact subgroups of a topological group are commensurate.*

Proof. Let $H, H' \leq G$ be compact open subgroups. Consider the open cover of the ambient group by cosets of H' and, using compactness, let $g_1, \dots, g_k$ be a finite set of coset representatives with

$$H \subseteq g_1 H' \cup \cdots \cup g_k H'.$$

Then

$$H = g_1(H \cap H') \cup \cdots \cup g_k(H \cap H'),$$

so that $H \cap H'$ has finite index in H. Completely similarly, $H \cap H'$ has finite index in H'. ∎

1.2.2 T.d.l.c. groups

We will mainly be interested in locally compact groups. Such groups are quite well-behaved:

Proposition 1.2.15. *Let G be a topological group and H a subgroup.*

(i) *If G is locally compact and H is closed, then H is locally compact.*

(ii) *If G is locally compact and Hausdorff, and H is locally compact, then H is closed.*

(iii) *If G is locally compact, then G/H is locally compact.*

(iv) *If both H and G/H are locally compact, then G is locally compact.*

(v) *If G is compact, then G/H is compact.*

(vi) *If both H and G/H are compact, then G is compact.*

Proof omitted. Property (i) is evident; property (ii) is [HR13, Theorem 5.11]; properties (iii) and (v) are [HR13, Theorem 5.22]; properties (iv) and (vi) are [HR13, Theorem 5.25].

Proposition 1.2.16. *Let G be a locally compact group. Then G is the filtering union of its compactly generated open subgroups (where "filtering" means that any two compactly generated open subgroups are contained in a single larger one).*

Proof. Let V be a compact neighbourhood of the identity. Note that

$$\langle V, g\rangle = \bigcup_{n\in\mathbb{N}} \left(V \cup V^{-1} \cup \{g^{\pm 1}\}\right)^n$$

is an open subgroup of G. In other words, every $g \in G$ is contained in some compactly generated open subgroup of G. The filtering property is clear, as two compactly generated open subgroups again generate a compactly generated open subgroup. ∎

Propositions 1.2.12 and 1.2.15 show that — at least, in principle — the study of (Hausdorff) locally compact groups splits into the study of *connected* locally compact groups, and the study of *totally disconnected* locally compact groups. The former has been solved satisfactorily with the solution of Hilbert's fifth problem by Andrew Gleason, Deane Montgomery, Leo Zippin and Hidehiko Yamabe in the 1950s:

Theorem 1.2.17 (Gleason–Yamabe). *Let G be a connected locally compact Hausdorff group. Then every neighbourhood U of the identity contains a compact normal subgroup $K \trianglelefteq G$ such that G/K is isomorphic to a Lie group.*

Proof omitted. An excellent exposition is given in [Tao11].

This theorem is often restated by saying that connected locally compact Hausdorff groups are inverse limits of Lie groups, or more intuitively, "can be approximated by Lie groups".

The general structure of totally disconnected locally compact (t.d.l.c.) groups is far less understood. One of the strongest known structural results is the following theorem by van Dantzig.

Theorem 1.2.18 (van Dantzig). *Every t.d.l.c. group contains a compact open subgroup.*

Proof omitted. We refer to [vD36].

As such a compact open subgroup is again t.d.l.c., this immediately implies a stronger statement.

Corollary 1.2.19 (van Dantzig). *Every t.d.l.c. group admits an identity neighbourhood basis of compact open subgroups.*

The totally disconnected case hence stands in sharp contrast with the connected case: locally compact groups of the former kind have an abundance of open subgroups, while those of the latter kind have no proper open subgroups.

Another useful fact is that totally disconnected groups are automatically Hausdorff.

Proposition 1.2.20. *Let G be a totally disconnected group G. Then G is Hausdorff.*

Proof. Connected components of topological spaces are closed, hence the set $\{1\}$ is closed in G. Now consider the map $f\colon G \times G \to G\colon (g, h) \mapsto g \cdot h^{-1}$. By continuity of f, the diagonal set

$$\{(g, g) \mid g \in G\} = f^{-1}(1)$$

is closed in $G \times G$, which is equivalent to saying that G is Hausdorff. ∎

For compactly generated t.d.l.c. groups, [KM07] presents a construction of an analogue of a Cayley graph (which the authors call a *rough Cayley graph* and is more commonly called a *Cayley–Abels graph*). The crux of the construction is the existence of a so-called *good generating set* — a compact generating set of a very particular form.

Definition 1.2.21 (good generating sets). Let G be a compactly generated t.d.l.c. group. A *good generating set* for G is a compact open subgroup V together with a finite subset T which is closed under inverses, satisfying $G = \langle T \rangle \cdot V$.

Lemma 1.2.22. *Let G be a compactly generated t.d.l.c. group. Let V be any compact open subgroup. Then there exists a finite set T such that (V, T) is a good generating set for G.*

Proof omitted. This is [KM07, Lemma 2.3], which refers to [Möl03, Lemma 2].

Example 1.2.23. (i) A discrete group is trivially totally disconnected and locally compact. It is compactly generated if and only if it is finitely generated.

(ii) The additive group $\mathbb{R}^n$ is connected, locally compact, and compactly generated.

(iii) The additive group $\mathbb{Q}$, with the subspace topology inherited from $\mathbb{R}$, is totally disconnected (even totally separated).

Moreover, $\mathbb{Q}$ is compactly generated, by the compact set $\{0\} \cup \{\frac{1}{n} \mid n \in \mathbb{N} \setminus \{0\}\}$.

On the other hand, $\mathbb{Q}$ is not locally compact. Indeed, suppose by means of contradiction that 0 has a compact neighbourhood K. Then K contains a basic open set $\{x \in \mathbb{Q} \mid a < x < b\}$ for some $a \in \mathbb{R}^-$ and $b \in \mathbb{R}^+$ (by definition of the subspace topology). Let $a', b' \in \mathbb{Q}$ be such that $a < a' < 0 < b' < b$. Then $\{x \in \mathbb{Q} \mid a' \leq x \leq b'\}$ is a closed subset of K, and hence compact. But closed intervals in $\mathbb{Q}$ are not even sequentially compact; a contradiction.

(iv) For the following example, consider the countably infinite direct product

$$G = \prod_{n=0}^{\infty} \mathbb{Z}/2\mathbb{Z}.$$

Pick a bit $b_n \in \mathbb{Z}/2\mathbb{Z}$ for every $n \in \mathbb{N}$. Let $G' \leq G$ be the subgroup of all elements $(x_n)_{n\in\mathbb{N}}$ such that $x_n \in \langle b_n \rangle$ for all but finitely many n. Note that this construction reduces to the standard direct product if $b_n = 1$ for all n, and to the direct sum if $b_n = 0$ for all n. We can define a topology on G' by taking as a basis of open sets, all sets of the from

$$\prod_{n=0}^{\infty} X_n$$

with $X_n = \langle b_n \rangle$ for all but finitely many n. This makes G' into a totally disconnected locally compact topological group, whose topology is governed by the choice of bits $(b_n)_{n\in\mathbb{N}}$ in the following way:

- G' is discrete if and only if $b_n = 0$ for all but finitely many n;
- G' is compact if and only if $b_n = 1$ for all but finitely many n.

In particular, by letting for instance $b_n = n \pmod 2$, this construction provides examples of t.d.l.c. groups that are neither discrete nor compact. This is a particular instance of a *restricted* or *local direct product.* For a far more general definition and more detailed properties, we refer to [HR13, (6.16)]. We only mention that local direct products play quite an important role in the study of t.d.l.c. *abelian* groups ([Bra45, HHR18]).

(v) The next section is devoted to a natural example of a t.d.l.c. topology on permutation groups, called the *permutation topology.*

1.2.3 The permutation topology

Definition 1.2.24 (permutation topology). Let G be any group acting on a set X. The *permutation topology* on G is the topology defined by taking as a identity neighbourhood basis, all pointwise stabilisers of finite subsets of X.

Intuitively, the more points of X on which two group elements agree, the closer those elements are in the permutation topology.

It is instructive to think of the permutation topology in terms of convergence: $g \in G$ is a limit of a net $(g_n)_{n\in\mathbb{I}}$ in G if and only if for every $x \in X$ there exists some $m \in \mathbb{I}$ (depending on x) such that $g\,.\,x = g_n\,.\,x$ for all $n > m$. In other words, when we think of X as a discrete space and of permutations as maps $X \to X$, then the permutation topology on G agrees with the topology of pointwise convergence.

As the permutation topology is defined in terms of purely "algebraical" data from the group action, it should not come as a surprise that there is an even stronger interplay between algebraical and topological properties. In the next few propositions we collect a couple of such results.

Proposition 1.2.25. *Let G be acting on a set X and endow G with the permutation topology. Then a subgroup $H \leq G$ is open if and only if H contains the pointwise stabiliser of some finite subset of X.*

Proof. This follows immediately from the definition. ∎

As a group is discrete if and only if the trivial subgroup is open, this immediately implies …

Corollary 1.2.26. *Let G be acting on a set X and endow G with the permutation topology. Then the following are equivalent:*

(i) *the pointwise stabiliser of some finite subset of X is trivial;*

(ii) *G is discrete.*

Proposition 1.2.27. *Let G be acting on a set X and endow G with the permutation topology. Then the following are equivalent:*

(i) *G acts faithfully on X;*

(ii) *G is totally separated;*

(iii) *G is totally disconnected;*

(iv) *G is Hausdorff.*

Proof. For (i) ⇒ (ii), it is sufficient to show that an arbitrary element $g \in G$ can be separated from the identity in G. Using faithfulness, there exists some $x \in X$ such that $g\,.\,x \neq x$. Then G_x is an open subgroup of G, and the two sets G_x (containing the identity) and $G \setminus G_x$ (containing g) define a separation of G.

The implication (ii) ⇒ (iii) is a well-known exercise in point-set topology.

We have already established that (iii) ⇒ (iv) for general topological groups in Proposition 1.2.20.

Finally, for (iv) ⇒ (i), suppose by means of contraposition that G contains a nontrivial element g fixing X pointwise. Then by definition of the permutation topology, g is contained in every neighbourhood of the identity. Hence G is not Hausdorff. ∎

Definition 1.2.28. Following [Möl10], we say that a group G acting on X is *closed* if the image of the natural morphism $G \to \mathrm{Sym}(X)$ is closed in the permutation topology on $\mathrm{Sym}(X)$.

Lemma 1.2.29. *Let G be acting on a set X and endow G with the permutation topology. Then the following are equivalent:*

(i) *G is closed;*

(ii) *the point stabiliser G_x is closed for every $x \in X$;*

(iii) *the point stabiliser G_x is closed for some $x \in X$.*

Proof. Let $\phi\colon G \to \mathrm{Sym}(X)$ be the natural morphism. For (i) ⇒ (ii) it suffices to notice that $\phi(G_x)$ is equal to the intersection of $\mathrm{Sym}(X)_x$ and $\phi(G)$. The implication (ii) ⇒ (iii) is of course trivial.

Finally for the implication (iii) ⇒ (i), assume that G_x is closed. Consider a net $(g_n)_{n\in\mathbb{I}}$ in G such that $\phi(g_n) \to g$ in the permutation topology on $\mathrm{Sym}(X)$. Then there exists an index $m \in \mathbb{I}$ such that $g_n \,.\, x = \phi(g_n) \,.\, x = g \,.\, x$ for all $n \geq m$. Then the net $((g_m)^{-1} \cdot g_n)_{n\in\mathbb{I}'}$ over the index set $\mathbb{I}' = \{n \in \mathbb{I} \mid n \geq m\}$ is contained in the stabiliser G_x and moreover

$$\phi((g_m)^{-1} \cdot g_n) \to \phi(g_m)^{-1} \cdot g \in \phi(G_x).$$

Hence $g \in \phi(G)$, or in other words, G is closed. ■

At first sight both the terminology and the interpretation of Lemma 1.2.29 can be rather confusing. We want to explicitly point out that point stabilisers in G are always open subgroups, and hence closed, *with respect to the permutation topology on G*. Let $G \leq \mathrm{Sym}(X)$ be a (faithful) permutation group. Then note that the permutation topology on G coincides with the relative topology induced from the permutation topology on $\mathrm{Sym}(X)$. The condition that G is closed implies that, moreover, the closed sets in this topology also agree with the closed subsets of G in the permutation topology on $\mathrm{Sym}(X)$, so that there is no ambiguity possible.

Proposition 1.2.30. *Let G be acting on a set X and endow G with the permutation topology. Assume that G is closed. Then the following are equivalent:*

(i) *every G-orbit is finite;*

(ii) *G is compact.*

In particular, if $G \leq \mathrm{Sym}(X)$, a subgroup of G has compact closure if and only if all G-orbits are finite.

Proof. Suppose that G is compact and let $x \in X$ be arbitrary. Then the cosets of G_x define an open cover of G. By compactness, G_x has finite index in G, so that by the orbit-stabiliser theorem, the orbit $G \,.\, x$ is finite.

For the converse, suppose that all G-orbits are finite and let K be the kernel of the group action. Then G/K — acting faithfully on X — naturally embeds as a closed subgroup in the direct product of the finite symmetric groups on the G-orbits. The permutation topology on the direct product coincides with the product topology and is therefore compact by Tychonov's theorem. Thus, G/K is compact (in the quotient topology induced from the permutation topology on the full group G). To finish, note that K is contained in every neighbourhood of the identity and is therefore trivially compact. Proposition 1.2.15 (vi) then shows that G is compact.

For the last claim, note that taking the closure of a subgroup leaves the orbits invariant. ■

Proposition 1.2.31. *Let G be acting on a set X and endow G with the permutation topology. Assume that G is closed. Then the following are equivalent:*

(i) *the pointwise stabiliser in G of some finite subset of X has only finite orbits on X;*

(ii) *the pointwise stabiliser in G of some finite subset of X is compact;*

(iii) *G is locally compact.*

Proof. Note that (i) $\Leftrightarrow$ (ii) by Proposition 1.2.30, and (ii) $\Rightarrow$ (iii) is trivial.

For (iii) $\Rightarrow$ (ii), let K be a compact neighbourhood of the identity. Then there exists a basic open set in between $\{1\}$ and K, i.e. there exists a finite subset $Y \subseteq X$ such that $G_{(Y)} \subseteq K$. This subgroup $G_{(Y)}$ is open in G, hence closed, hence compact (as a closed subset of a compact set). ∎

It is worth making explicit that the full symmetric group $\mathrm{Sym}(X)$ on any infinite set X, endowed with the permutation topology, is hence *not* locally compact.

For the next proposition, recall that a subgroup $H \leq G$ is called *cocompact* (or sometimes tongue-in-cheekly, *mpact*) if the quotient space G/H of left cosets is compact. Note that this is equivalent to asking the space $H\backslash G$ of right cosets to be compact, as the map $G/H \to H\backslash G \colon gH \mapsto Hg^{-1}$ is a homeomorphism.

Proposition 1.2.32. *Let G be acting on a set X and endow G with the permutation topology. Assume that G is closed and that all suborbits of G are finite (so that all point stabilisers of G are compact). Let $H \leq G$ be a subgroup. Then the following are equivalent:*

(i) *some G-orbit is the union of finitely many H-orbits;*

(ii) *H is cocompact in G.*

Proof. Let $x \in X$ be arbitrary. Consider the partition of G into double cosets

$$H\backslash G/G_x = \{H \cdot g \cdot G_x \mid g \in G\}$$

and note that every such double coset is open in G. Since the natural map $\varphi \colon G \to H\backslash G$ is open by Proposition 1.2.7 (i) and continuous by definition, the sets $\varphi(H \cdot g \cdot G_x) = \varphi(g \cdot G_x)$ are open and compact in $H\backslash G$. It thus follows that $H\backslash G$ is compact if and only if the partition $H\backslash G/G_x$ has finitely many blocks. The proposition now follows from the observation that the map

$$\Lambda \colon H\backslash G/G_x \to H\backslash X \colon H \cdot g \cdot G_x \mapsto H\,.(g\,.x)$$

defines a bijection between double cosets and H-orbits contained in $G\,.x$. ∎

As noted by Pierre-Emmanuel Caprace, we can rephrase the assumptions in Proposition 1.2.32 and alternatively require that G is a locally compact group acting continuously and properly by permutations on the discrete set X. Here, a group action is continuous and proper when the shear map $G \times X \to X \times X \colon (g, x) \mapsto (g\,.x, x)$ is continuous and proper (i.e. preimages of compact sets are again compact).

1.3 Graph theory

For completeness' sake, this section is a quick refresher of the basic notions of graph theory.

Definition 1.3.1. (i) A (simple, undirected) *graph* $\Gamma = (V, E)$ consists of a set V together with a set E of unordered pairs of elements of V. The elements of V are called the graph's *vertices* while the elements of E are called the *edges*. If $\{v, w\}$ is an edge, then we say v and w are its *endpoints*. Also we say that both v and w are *incident* to $\{v, w\}$ and *adjacent* to each other, or simply *neighbours*, and we then write $v \sim w$. We will usually identity the graph (V, E) with its vertex set V.

(ii) Depending on the context, one sometimes allows for *loops*: edges with coinciding endpoints. In this thesis, this is only relevant for Higman's primitivity criterion Theorem 1.1.7.

(iii) Given a vertex v of a graph (V, E), we define its *star* as the set

$$\mathrm{st}(v) = \{e \in E \mid v \text{ is incident to } e\}.$$

The *degree* of vertex v is the cardinality $\deg(v) = |\mathrm{st}(v)|$, i.e. the number of incident edges.

(iv) A graph is called *regular* if all vertices have the same degree, and *locally finite* if all vertices have finite degree.

(v) A *morphism* between graphs (V, E) and (V', E') is a map $\varphi \colon V \to V'$ that preserves adjacency. Explicitly, for every edge $\{v, w\} \in E$, we must have that $\{\varphi(v), \varphi(w)\} \in E'$. If φ is bijective and moreover induces an bijection $: E \to E'$ then we call φ an *isomorphism.* If $V = V'$ an isomorphism is called an *automorphism.* The set of all automorphisms of a fixed graph Γ is a group under the operation of composition, and is denoted by $\mathrm{Aut}(\Gamma)$.

(vi) A *subgraph* of (V, E) is any graph (V', E') with $V' \subseteq V$ and $E' \subseteq E$. Every subset $V' \subseteq V$ defines an *induced subgraph* with edge set $E' = \{\{v, w\} \in E \mid v, w \in V'\}$.

(vii) A *path* of length n in a graph is a sequence $v_0 \sim v_1 \sim v_2 \sim \cdots \sim v_n$ of pairwise adjacent vertices. It is said to be a path *from* v_0 *to* v_n or a path *joining* v_0 *and* v_n. The notion extends to infinite or bi-infinite paths. If $v_0 = v_n$ the path is called *closed.* If all interior vertices are distinct, the path is called *simple.* Usually we will implicitly assume paths to be simple.

(viii) The *distance* between two vertices is the length of any shortest path joining them; if no such path exists, the distance is defined to be infinite. Note that the function

$$\mathrm{dist} \colon V \times V \to \mathbb{N} \cup \{\infty\} \colon (v, w) \mapsto \mathrm{dist}_\Gamma(v, w)$$

indeed turns the graph Γ into a metric space. The *diameter* of a graph is the maximal distance between two vertices.

(ix) If the diameter is finite, or in other words if every two vertices are joined by a path, the graph is said to be *connected.* In general, a *connected component* is a maximal connected subgraph.

(x) A simple closed path of length $n \geq 3$ is an n*-cycle*, or simply a *cycle.* The *girth* of a graph is the length of the shortest cycle; if the graph is acyclic, the girth is defined to be infinite.

(xi) A *tree* is a connected acyclic graph. Note that a n-regular tree with $n \geq 2$ is an infinite graph and, up to isomorphism, unique.

(xii) A graph (V, E) is *bipartite* if its vertex set admits a partition $V = V_1 \sqcup V_2$ such that every edge in E has one endpoint in V_1 and one endpoint in V_2. If the degree function is constant on both bipartition classes, then we call the graph *semiregular.*

(xiii) A *vertex cover* of a graph (V, E) is a subset $V' \subseteq V$ of the vertex set that includes at least one endpoint of every edge.

To make the bridge with the previous chapter on topological group theory, a well-known corollary of Propositions 1.2.27, 1.2.30, and 1.2.31 is the following.

Corollary 1.3.2. *Let* Γ *be a locally finite connected graph. Endow its automorphism group* $\mathrm{Aut}(\Gamma)$ *with the permutation topology. Then* $\mathrm{Aut}(\Gamma)$ *is a t.d.l.c. group with compact vertex stabilisers.*

In the next section, we will study trees and their automorphisms in some more detail. Additionally, in Section 1.7, we will essentially define chamber systems and buildings as graph-like structures. Quite some graph theory terminology has a counterpart in building theory — compare for instance the notions of paths and galleries, or stars and panels.

1.4 Automorphisms of trees

Trees will turn out to be prototypical examples of right-angled buildings. In this section, we bundle some properties of group actions on trees that we will frequently encounter throughout our study. For more detailed definitions, we refer to [Tit70] (in French) or [GGT18].

In order to obtain a rich group of automorphisms, we usually want trees to display some regularity. As an example, the infinite tree having only vertices of degree two except for one central vertex of degree $n \geq 3$, has a rather uninspiring automorphism group isomorphic to $\mathrm{Sym}(n)$. We could ask for instance that the full automorphism group leaves no nontrivial subtree invariant, leading to the following definition.

Definition 1.4.1 (minimal, geometrically dense). Let G be a group acting on a tree T by automorphisms. The action is said to be *minimal* if T has no nontrivial G-invariant subtree. The action is *geometrically dense* if it is minimal and if moreover G does not fix an end of T.

The existence of a minimally acting group is generally sufficient to exclude degenerate edge cases, but we explicitly put forward a couple of additional assumptions.

We will always assume that a tree has no vertices of degree one (or *leaves*) and has at least one vertex of degree at least three (as to exclude the bi-infinite path graph).

Definition 1.4.2. Let g be an automorphism of a tree. Then we define the *displacement* $\ell(g)$ of g as the minimal value $\min_v \operatorname{dist}(v, g \mathbin{.} v)$ over all vertices v.

Proposition 1.4.3. *Let g be an automorphism of a tree T. Then exactly one of the following holds:*

(i) *g fixes a vertex, so that $\ell(g) = 0$;*

(ii) *g inverts an edge $\{v_1, v_2\}$, i.e. we have that $g \mathbin{.} v_1 = v_2$ and $g \mathbin{.} v_2 = v_1$;*

(iii) *g leaves invariant a bi-infinite path γ in T and induces a nontrivial translation on γ.*

In (iii) *the path γ is unique, the induced translation on γ has length $\ell(g)$, and γ can be characterised as the set of all vertices v such that* $\operatorname{dist}(v, g \mathbin{.} v) = \ell(g)$.

Proof omitted. This is [Tit70, Proposition 3.2].

Definition 1.4.4. Let g be an automorphism of a tree T. Then we call g

(i) *elliptic* if g fixes a vertex;

(ii) an *inversion* if g inverts an edge;

(iii) *hyperbolic* otherwise.

The bi-infinite path left invariant by a hyperbolic element g is called its *axis* and denoted by $\mathsf{A}(g)$.

Proposition 1.4.5. *Let G act on a tree T by automorphisms.*

(i) *If $g, h \in G$ are hyperbolic elements with disjoint axes $\mathsf{A}(g), \mathsf{A}(h)$, then gh is hyperbolic as well and its axis $\mathsf{A}(gh)$ contains the unique shortest path joining $\mathsf{A}(g)$ to $\mathsf{A}(h)$. Moreover,*

$$\ell(gh) = \ell(g) + \ell(h) + 2\operatorname{dist}(\mathsf{A}(g), \mathsf{A}(h)).$$

(ii) *If $g, h \in G$ are elliptic elements with disjoint fixed point sets, then gh is hyperbolic.*

Proof. For (i), pick an edge e in the arc joining $\mathsf{A}(g)$ to $\mathsf{A}(h)$. Let T_g and T_h denote the components of $T \setminus \{e\}$ containing $\mathsf{A}(g)$ and $\mathsf{A}(h)$, respectively. Note that

$$g^{\pm 1} . (T_h \cup \{e\}) \subset T_g \qquad \text{and} \qquad h^{\pm 1} . (T_g \cup \{e\}) \subset T_h,$$

so that for every $m, n > 0$, we have $(gh)^m . e \in T_g$ and $(gh)^{-n} . e \in T_h$. Hence the orbit $\langle gh \rangle . e$ is unbounded and gh is hyperbolic. Since e is contained in the geodesics from $(gh)^{-n} . e$ to $(gh)^m . e$ for every $m, n > 0$, it follows that e is in fact contained in the axis $\mathsf{A}(gh)$. The formula for $\ell(gh)$ then easily follows from considering the effect of gh on the vertex on $\mathsf{A}(h)$ closest to $\mathsf{A}(g)$.

For (ii) a similar proof applies, replacing $\mathsf{A}(g)$ and $\mathsf{A}(h)$ with the fixed point sets of g and h. ■

Proposition 1.4.6. *Let G act on a tree T by automorphisms and assume that G does not contain any hyperbolic elements. Then G stabilises some vertex, some edge, or some end of the tree T.*

Proof omitted. This is [Tit70, Proposition 3.4], or [GGT18, Corollary 6.6] for the special case where G acts without inversions. ■

Proposition 1.4.7. *Let G act on a tree T by automorphisms. If G contains a hyperbolic element, then there is a unique minimal G-invariant subtree, namely*

$$T_0 = \{ v \in T \mid v \text{ lies on } \mathsf{A}(g) \text{ for some hyperbolic automorphism } g \in G \}.$$

Proof. For every $g, h \in G$ with g hyperbolic, we have $h . \mathsf{A}(g) = \mathsf{A}(hgh^{-1})$. Hence T_0 is indeed G-invariant. In addition, T_0 is connected by Proposition 1.4.5. On the other hand, if $g \in G$ is hyperbolic, then any convex G-invariant subtree contains its axis $\mathsf{A}(g)$. We conclude that T_0 is an invariant subtree contained in every other invariant subtree. ■

Definition 1.4.8 (end, boundary). Let T be a tree. A *ray* is an isometric embedding $\alpha \colon \mathbb{N} \to T$ (the image of which is a semi-infinite simple path). We declare two rays α_1, α_2 to be equivalent if the function $n \mapsto \operatorname{dist}(\alpha_1(n), \alpha_2(n))$ remains bounded. Note that, in a tree, two equivalent rays necessarily eventually coincide. An equivalence class of rays is called a *point at infinity* or an *end* of T. The set of all ends is called the *boundary* of T and denoted by ∂T.

Proposition 1.4.9. *Let G act on T by automorphisms and assume that the action is minimal. Then every edge of T lies on the axis $\mathsf{A}(g)$ of some hyperbolic element $g \in G$.*

Proof. By Proposition 1.4.6, G contains a hyperbolic element. Proposition 1.4.7 and the assumption that G acts minimally yield that every vertex v lies on the axis of some hyperbolic automorphism $g_v \in G$. Now let $e = \{v, w\}$ be any edge of T and consider g_v and g_w. If $v \in \mathsf{A}(g_w)$ or $w \in \mathsf{A}(g_v)$, then we are done. Otherwise v and w both lie on the axis of $g_v g_w$ by Proposition 1.4.5 (i). ■

Recall that we excluded the bi-infinite path graph; the next few propositions motivate why.

Proposition 1.4.10. *Let G act on T by automorphisms and assume that the action is geometrically dense. Then G has no finite orbits in $T \cup \partial T$.*

Proof. Let v be a vertex of T and assume that $G.v$ is finite. Then the convex hull of $G.v$ is a finite G-invariant subtree and hence the action is not minimal. Next, let ξ be an end of T and assume that $G.\xi$ is finite, say $|G.\xi| = n$. If $n = 1$, i.e. if G fixes ξ, the action is not geometrically dense. If $n = 2$, then G stabilises the unique bi-infinite path determined by those two ends and the action is not minimal. Finally, note that every three ends of the tree determine a unique vertex. If $n \geq 3$, pick any subset $\{\xi_1, \xi_2, \xi_3\} \subseteq G.\xi$ and let v be the uniquely determined vertex. The cardinality of the orbit $G.v$ is then bounded by $\binom{n}{3}$ — a situation already handled. ■

Definition 1.4.11 (half-tree). Let T be a tree and e an edge. Denote by $T \setminus \{e\}$ the graph obtained by removing the edge e from T. Then we call the two connected components of $T \setminus \{e\}$, *half-trees* of T (determined by e).

Proposition 1.4.12. *Let G act on T by automorphisms. Assume that the action is without inversions and geometrically dense. Let $T_0 \subset T$ be any half-tree. Then G contains a hyperbolic automorphism with axis contained in T_0.*

Proof. Define the *attracting end* $a(g) \in \partial T$ of a hyperbolic automorphism g as the end of its axis in the direction of translation; the opposite end is the *repelling end* $r(g)$. Then, for arbitrary $h \in G$, we have $h.a(g) = a(hgh^{-1})$ and $h.r(g) = r(hgh^{-1})$. Hence the set $A \subseteq \partial T$ of all attracting ends of hyperbolic automorphisms in G is G-invariant, and so is the subtree of T spanned by the ends in A. Note that A is nonempty by Proposition 1.4.6. By minimality, we conclude that the boundary of every half-tree in T contains the attracting end of some hyperbolic automorphism $g \in G$.

In particular, there exists a hyperbolic $h \in G$ such that $a(h) \in \partial T_0$. Consider an automorphism g_0 of the form fhf^{-1} with $f \in G$. Then $r(g_0) = f.r(h)$ and $a(g_0) = f.a(h)$. We can choose $f \in G$ in such a way that both ends of g_0 differ from the repelling end of h — indeed, if for every $f \in G$ we would have that $f^{-1}.r(h) \in \{r(h), a(h)\}$, then either the end $r(h)$ or the set $\{r(h), a(h)\}$ would be G-invariant, in contradiction to the assumption that the action is geometrically dense.

Fix $g_0 \in G$ such that both $r(g_0)$ and $a(g_0)$ differ from $r(h)$. Define for every $n \in \mathbb{N}$ the conjugated automorphism $g_n = h^n \cdot g_0 \cdot h^{-n} \in G$. Then g_n is again hyperbolic with axis $\mathsf{A}(g_n) = h^n . \mathsf{A}(g_0)$, with repelling end $r(g_n) = h^n.r(g_0)$, and with attracting end $a(g_n) = h^n.a(g_0)$. For sufficiently large n, both ends $r(g_n)$ and $a(g_n)$ will end up in ∂T_0 and the result follows. ■

Using Proposition 1.4.12, we now establish an improvement of Proposition 1.4.9 for geometrically dense actions.

Proposition 1.4.13. *Let G act on T by automorphisms. Assume that the action is without inversions and geometrically dense. Let e_1 and e_2 be two edges. Then there exists some hyperbolic element $g \in G$, the axis of which contains both e_1 and e_2.*

Proof. We may assume that $e_1 \neq e_2$ by Proposition 1.4.9. Let T_1 be the half-tree determined by e_1 not containing e_2 and let T_2 be the half-tree determined by e_2 not containing e_1. Proposition 1.4.12 yields two hyperbolic elements g_1 and g_2 in G with axes $\mathsf{A}(g_1) \subset T_1$ and $\mathsf{A}(g_2) \subset T_2$. In particular, the axes are disjoint, and the shortest arc joining $\mathsf{A}(g_1)$ and $\mathsf{A}(g_2)$ passes through both e_1 and e_2. By Proposition 1.4.5 (i), the element $g_1\, g_2$ does the job. We refer to Figure 1.1 for an illustration. ■

For later purposes, it will be important to know that geometrical density is preserved when passing to a normal subgroup.

Proposition 1.4.14. *Let G act on T by automorphisms. Assume that the action is without inversions and geometrically dense. Let $N \trianglelefteq G$ be a nontrivial normal subgroup. Then the action of N on T is geometrically dense as well.*

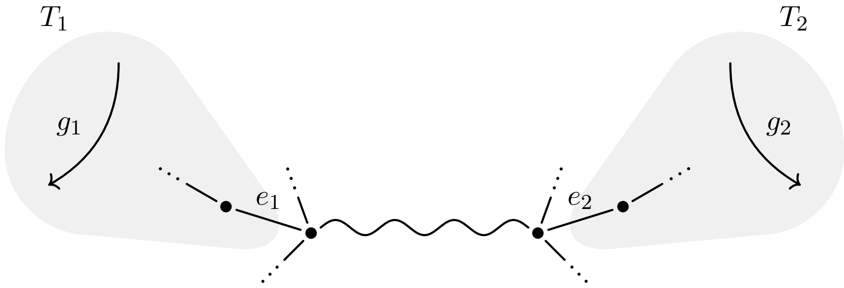

Figure 1.1. The configuration in the proof of Proposition 1.4.13.

Proof. First, we claim that N does not fix any vertex of T. Indeed, if it were the case that $N.v = v$ for some vertex v, then N fixes in fact every vertex in the orbit $G.v$. The convex hull of $G.v$ is a G-invariant subtree and is hence the full tree T by minimality. It follows that N fixes all vertices, and is hence the trivial subgroup.

Similarly, we claim that N does not fix any end of the tree. Indeed, if it were the case that $N.\xi = \xi$ for some end ξ, then N fixes every end in the orbit $G.\xi$. This orbit is infinite by Proposition 1.4.10. The convex hull of all vertices determined by three ends in $G.\xi$ is a G-invariant subtree, hence is the full tree by minimality of the action. It follows that $G.\xi = \partial T$, so that N fixes all ends, and is hence the trivial subgroup.

Because N fixes neither vertex, edge, or end, Proposition 1.4.6 yields that N contains a hyperbolic automorphism, and Proposition 1.4.7 that there is a unique minimal N-invariant subtree. This tree is stabilised by G. As the action of G is geometrically dense, we must have that in fact the full tree is N-invariant, which concludes the proof. ■

Now we finally state Tits's celebrated independence property and simplicity criterion.

Definition 1.4.15 (Tits's independence property). Let T be a tree and let γ be any path in T — either finite, infinite or bi-infinite. There is a canonical projection $\pi\colon T \to \gamma\colon v \mapsto \pi(v)$, mapping every vertex v to the vertex of γ that is closest to v. See Figure 1.2 for an illustration.

Let $G \leq \mathrm{Aut}(T)$ be a group of automorphisms of T and let $H = G_{(\gamma)}$ be the pointwise stabiliser of γ. The sets $\pi^{-1}(v)$, where v ranges over γ, define a partition of T into H-invariant subtrees. We then have a natural morphism

$$\varphi_\gamma\colon H \to \prod_{v\in\gamma} H\big|_{\pi^{-1}(v)}$$

which is injective, but not always surjective in general. We say that G satisfies *Tits's independence property* if φ_γ is an isomorphism for every choice of γ. Tits's independence property (intuitively) ensures that the actions of H on the subtrees branching from γ can be chosen independently from each other, and is also known as property (P) throughout the literature.

When G is a *closed* subgroup of $\mathrm{Aut}(T)$ in the permutation topology, the independence property can be relaxed to only take into account edges instead of paths of arbitrary lengths. More precisely we have the following.

Proposition 1.4.16. *Let G be a closed subgroup of $\mathrm{Aut}(T)$. Then G satisfies Tits's independence property if and only if for every edge e of T, the equality*

$$G_{(e)} = G_{(T_1)} \cdot G_{(T_2)}$$

holds, where T_1 and T_2 are the rooted half-trees emanating from e (so that $T = T_1 \sqcup T_2$).

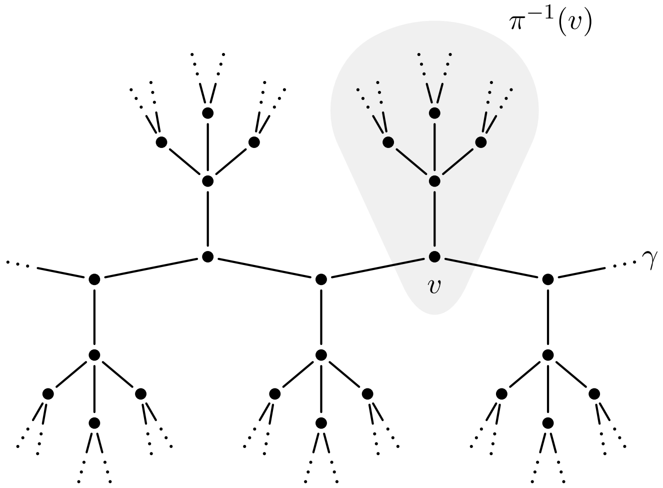

Figure 1.2. Tits's independence property for trees.

Proof omitted. This is Lemma 10 and the subsequent paragraph in [Ama03].

Theorem 1.4.17. *Let $G \leq \mathrm{Aut}(T)$ and let G^+ be the subgroup generated by pointwise stabilisers of edges of T. Assume that the action of G is geometrically dense and that G satisfies Tits's independence property. Then every nontrivial subgroup $H \leq G$ normalised by G^+ in fact contains G^+. In particular, if G^+ is nontrivial, then it is a simple group.*

Proof omitted. We refer to [Tit70, Théorème 4.5].

1.5 Universal groups over trees

In this section, we shall give an overview of the original construction by Marc Burger and Shahar Mozes in [BM00a, Section 3.2], and brush over some more recent results by Simon Smith ([Smi17]) and Adrien Le Boudec ([LB16]). The universal groups are defined as groups of automorphisms of a locally finite regular tree, restricting the local actions around every vertex to some fixed prescribed permutation group. The resulting groups, equipped with the permutation topology, provide a rich class of interesting topological groups (under very mild conditions on the local groups); see Proposition 1.5.5 below.

This section is by no means intended to be complete, contentwise nor rigourwise, but mostly aims to give an overview of recent developments since the original paper of Burger–Mozes.

We asssume that $m, n \geq 3$.

1.5.1 Burger–Mozes's groups

Definition 1.5.1 (colouring). Let $T_n = (V, E)$ be the n-regular tree. A *(legal) colouring* is a map

$$\lambda \colon E \mapsto \{1, \dots, n\}$$

assigning to every edge a label or *colour* in $\{1, \dots, n\}$, in such a way that for every vertex v of T_n the restriction of λ to the star of v is a bijection.

It is not hard to see that legal colourings always exists, and neither should the following lemma be surprising.

Lemma 1.5.2. *Let v and v' be two vertices of T_n and let λ and λ' be two legal colourings. Then there exists an automorphism g such that $g \,.\, v = v'$ and $\lambda' \circ g = \lambda$.*

Proof omitted. We refer to [LMZ94, Proposition 2.1].

Definition 1.5.3 (local action). For $v \in V$ and $g \in \mathrm{Aut}(T_n)$, the *local action* of g at v is given by

$$\sigma_\lambda(g, v) = \lambda\big|_{\mathrm{st}(g\,.\,v)} \circ g\big|_{\mathrm{st}(v)} \circ \lambda\big|_{\mathrm{st}(v)}^{-1} \in \mathrm{Sym}(n).$$

Definition 1.5.4. Let T_n be the n-regular tree, let λ be a legal colouring, and let $F \leq \mathrm{Sym}(n)$ be a permutation group. Then the *Burger–Mozes universal group* of F with respect to λ is the group

$$\mathcal{U}^\lambda(F) = \big\{ g \in \mathrm{Aut}(T_n) \;\big|\; \sigma_\lambda(g, v) \in F \text{ for every } v \in V \big\}. \tag{$*$}$$

From Lemma 1.5.2 it readily follows that the groups $\mathcal{U}^\lambda(F)$ and $\mathcal{U}^{\lambda'}(F)$ (defined using different legal colourings λ and λ') are in fact conjugate in $\mathrm{Aut}(T_n)$. Hence we can fix one colouring λ and abbreviate $\mathcal{U}^\lambda(F)$ to $\mathcal{U}(F)$.

Consider the extremal case with $F = \mathrm{Sym}(n)$. Then clearly $\mathcal{U}(F) = \mathrm{Aut}(\Delta)$, since the condition in ($*$) is void. In the other extreme case, if F is the trivial permutation group on n elements, then a straightforward application of the ping-pong lemma shows that $\mathcal{U}(F)$ is isomorphic to the free product of n copies of $\mathbb{Z}/2\mathbb{Z}$, coming from the automorphisms that invert an edge in a single star. For nontrivial F, the universal group $\mathcal{U}(F)$ interpolates between these extreme cases.

Endowing $\mathcal{U}(F)$ and $\mathrm{Aut}(T_n)$ with the permutation topology, Burger and Mozes then observe the following properties.

Proposition 1.5.5. *Let $F \leq \mathrm{Sym}(n)$ and consider the universal group $\mathcal{U}(F)$. The following hold:*

(i) *$\mathcal{U}(F)$ is vertex-transitive.*

(ii) *$\mathcal{U}(F)$ is edge-transitive if and only if F is transitive.*

(iii) *$\mathcal{U}(F)$ is a closed subgroup of $\mathrm{Aut}(T_n)$ and hence t.d.l.c.*

(iv) *$\mathcal{U}(F)$ is discrete if and only if F is a free permutation group.*

(v) *$\mathcal{U}(F)$ is compactly generated.*

Proof omitted. We refer to [BM00a] but note that these are quite straightforward.

Observing that the universal groups satisfy Tits's independence property (Definition 1.4.15), and using Theorem 1.4.17 and a result of [Tit70], the authors then remark without further proof:

Proposition 1.5.6. *Let $\mathcal{U}(F)^+$ be the subgroup generated by pointwise edge stabilisers of $\mathcal{U}(F)$.*

(i) *$\mathcal{U}(F)^+$ is either simple or trivial.*

(ii) *$\mathcal{U}(F)^+$ is of finite index in $\mathcal{U}(F)$ if and only if F is transitive and generated by stabilisers. In this case, $\mathcal{U}(F)^+ = \mathcal{U}(F) \cap \mathrm{Aut}(T_n)^+$ and $[\mathcal{U}(F) : \mathcal{U}(F)^+] = 2$.*

Proof omitted. This is [BM00a, Proposition 3.2.1]. We also refer to [GGT18] for more details.

Note, the local action of $\mathcal{U}(F)$ on any star of the tree is permutationally isomorphic to the original group F. In fact, among the groups of automorphisms which are both vertex-transitive and locally permutationally isomorphic to F, the universal group is maximal, in the following sense — this is the universality property that justifies the name.

Proposition 1.5.7. *Let $F \leq \mathrm{Sym}(n)$ be a transitive permutation group and let $H \leq \mathrm{Aut}(T_n)$ be a vertex-transitive subgroup. Assume that the local action of H on every star of T_n is permutationally isomorphic to F. Then $H \leq \mathcal{U}^{\lambda}(F)$ for some suitable legal colouring λ.*

Proof omitted. This is [BM00a, Proposition 3.2.2].

1.5.2 Smith's groups

Simon Smith generalised the construction of Burger–Mozes to groups acting on semiregular trees with two prescribed local groups. Since the construction is very similar to Definition 1.5.4 (and in fact agrees with our Definition 3.1.4 in the setting of right-angled buildings) we will not go through the details, but refer to [Smi17]. More important than relaxing to semiregular trees is the fact that Smith did not assume the trees to be locally finite, but allowed for permutation groups on sets of arbitrary cardinality. Unsurprisingly, the topology does get more subtle in this setting.

Proposition 1.5.8. *Let $M \leq \mathrm{Sym}(X_1)$ and $N \leq \mathrm{Sym}(X_2)$. Let T be the $(|X_1|, |X_2|)$-semiregular tree and consider the universal group $\mathcal{U}(M, N) \leq \mathrm{Aut}(T)$. The following hold:*

(i) *If M and N are closed, then $\mathcal{U}(M, N)$ is closed subgroup of $\mathrm{Aut}(T)$.*

(ii) *If $M = M^+$ and $N = N^+$, then $\mathcal{U}(M, N)$ is simple if and only if M or N is transitive.*

(iii) *If M and N are closed, then $\mathcal{U}(M, N)$ is locally compact if and only if every point stabiliser of M and N is compact.*

(iv) *If M and N are closed, compactly generated, have compact point stabilisers, have finitely many orbits, and M or N is transitive, then $\mathcal{U}(M, N)$ is compactly generated.*

(v) *$\mathcal{U}(M, N)$ is discrete if and only if M and N are free permutation groups.*

Proof omitted. This is [Smi17, Theorem 1], which also includes a generalisation of the universality property Proposition 1.5.7.

Using the general universal groups, Smith defined the *box product $M \boxtimes N$* of permutation groups, being the new permutation group induced by $\mathcal{U}(M, N)$ acting on one bipartition class of the semiregular tree. This box product inherits quite a few permutational and topological properties from the local groups, in particular primitivity:

Theorem 1.5.9. *Let $M \boxtimes N$ be the box product of $M \leq \mathrm{Sym}(X_1)$ and $N \leq \mathrm{Sym}(X_2)$.*

(i) *$M \boxtimes N$ is transitive if and only if M is transitive.*

(ii) *$M \boxtimes N$ is primitive if and only if M is primitive and nonregular, and N is transitive.*

Proof omitted. We refer to [Smi17, Theorem 26.].

The box product has another, even more spectacular application that we will now describe. In 1953, Ruth Camm proved the existence of a continuum of nonisomorphic simple 2-generated groups in [Cam53]. A topological analogue is the class $\mathscr{S}$ of nondiscrete compactly generated topologically simple t.d.l.c. groups. The cardinality of the isomorphism classes of $\mathscr{S}$ has long been an open question, until Smith showed in [Smi17, Theorem 38] the existence of a continuum of nonisomorphic groups of the form $\mathcal{U}(M, N)$ in $\mathscr{S}$.

1.5.3 Le Boudec's groups

In another direction of generalisation, Adrien Le Boudec in [LB16] allowed the universal groups of Burger–Mozes to have a finite number of *singularities*, where the local action does not need to be a permutation in the prescribed local group. We do ask the tree to be locally finite again.

Definition 1.5.10. Let T_n be the n-regular tree, let λ be a legal colouring, and let $F \leq \mathrm{Sym}(n)$ be a permutation group. Then the *Le Boudec group* of F with respect to λ is the group

$$\mathcal{G}^\lambda(F) = \left\{ g \in \mathrm{Aut}(T_n) \,\middle|\, \sigma_\lambda(g, v) \in F \text{ for all but finitely many } v \in V \right\}. \tag{$*$}$$

If $g \in \mathcal{G}^\lambda(F)$, then an exceptional vertex v in $(*)$ is called a *singularity* of the automorphism g.

As it turns out, the local actions of $\mathcal{G}^\lambda(F)$ exhibit some rigidity at singularities as well.

Proposition 1.5.11. *For every $v \in V$ and $g \in \mathcal{G}(F)$, the local action $\sigma_\lambda(g, v)$ is a permutation that stabilises the F-orbits. In other words, recalling Definition 1.1.16, we have an inclusion*

$$\mathcal{G}(F) \leq \mathcal{U}(\widehat{F}).$$

Proof omitted. The key observation is that whenever $\sigma_\lambda(g, v) \notin \widehat{F}$, we can find at least two neighbours of v, and hence a bi-infinite path of singularities.

Le Boudec continued with a second local group F' satisfying $F \leq F' \leq \widehat{F}$ and defined the group

$$\mathcal{G}(F, F') = \mathcal{G}(F) \cap \mathcal{U}(F').$$

This group was coined a *restricted universal group* in [CRW19]. It contains all automorphisms where the local actions are restricted to F' and, for up to a finite number of exceptions, even to F. There is a way to topologise $\mathcal{G}(F, F')$ that makes the inclusion $\mathcal{U}(F) \hookrightarrow \mathcal{G}(F, F')$ an open continous map. This allows one to transfer topological properties of $\mathcal{U}(F)$ to the restricted universal groups.

On the other hand, the groups $\mathcal{G}(F, F')$ are usually not closed in $\mathrm{Aut}(T_n)$, and they hence escape common criteria for closed subgroups of $\mathrm{Aut}(T_n)$ to contain *lattices*: a lattice in a locally compact group G is a discrete subgroup H such that the quotient space G/H carries a finite G-invariant measure. Lattices are of central interest in the theory of locally compact groups, and groups without lattices seem to be quite rare. However, as an application, Le Boudec established the following.

Theorem 1.5.12. *There exists permutation groups $F \leq F' \leq \mathrm{Sym}(n)$ such that the group $\mathcal{G}(F, F')$ does not admit any lattice, for example $F = \mathrm{PSL}(2, q)$ and $F' = \mathrm{PGL}(2, q)$ acting on the projective line $\mathrm{PG}(1, q)$ with $q \equiv 1 \pmod 4$.*

Proof omitted. This is [LB16, Theorem 1.4].

Le Boudec remarks that, at that point, the only known compactly generated simple group without lattices was the Neretin group $\mathrm{AAut}(T_n)$ of spheromorphisms of a regular tree (Theorem 1.5.18).

Theorem 1.5.13. *There exist compactly generated t.d.l.c. groups $H \leq G$ such that*

(i) *H is cocompact in G;*

(ii) *both G and H are abstractly simple;*

(iii) *G contains lattices while H does not.*

Proof omitted. This is [LB16, Theorem 1.6].

Theorem 1.5.14. *There exist nondiscrete, compactly generated, locally compact, abstractly simple groups admitting (cocompact) simple lattices.*

Proof omitted. This is [LB16, Theorem 1.7].

1.5.4 Lederle's groups

Finally, we briefly mention some results by Waltraud Lederle ([Led17]), who allowed for even more drastic departures from the Burger–Mozes universal group. First, some definitions — we note that we slightly deviate from the terminology of Lederle (speaking of *almost-automorphisms* and *spheromorphisms*, instead of *honest almost-automorphisms* and *almost-automorphisms*).

Definition 1.5.15 (almost-automorphism). Let T be the n-regular tree. A finite subtree $S \subset T$ is called *complete* if every vertex of S is either a leaf or a vertex of degree n. Given any complete subtree S, by the difference $T \setminus S$ we mean the rooted forest obtained by removing from T all edges and internal vertices of S, and declaring the leaves of S to be the roots. Note that $T \setminus S$ is a forest with as many components as S has leaves. Now an *almost-automorphism* of T is an isomorphism of rooted forests $\varphi\colon T \setminus S_1 \to T \setminus S_2$ where S_1 and S_2 are two complete finite subtrees.

Definition 1.5.16 (spheromorphism). We declare almost-automorphisms $g\colon T \setminus S_1 \to T \setminus S_2$ and $h\colon T \setminus S_1' \to T \setminus S_2'$ to be equivalent if we have

$$g\big|_{T\setminus S} = h\big|_{T\setminus S}$$

for some finite complete subtree $S \supseteq S_1 \cup S_1'$ of T. Note that equivalent almost-automorphisms induce the same homeomorphism on the boundary of the tree, called a *spheromorphism* of ∂T.

There is a natural way to compose two spheromorphisms: pick representatives $g\colon T \setminus S_1 \to T \setminus S_2$ and $h\colon T \setminus S_1' \to T \setminus S_2'$, then let $S \supseteq S_1 \cup S_2'$ be finite complete subtree, and pass to equivalent representatives $g' \approx g$ and $h' \approx h$ defined on $T \setminus S$. These representatives can then be composed. It is not too hard to convince oneself this gives a well-defined composition of equivalence classes.

Definition 1.5.17 (Neretin group). The set of all spheromorphisms is a group under the composition described above, called the *Neretin group* and denoted by $\mathrm{AAut}(T)$.

Theorem 1.5.18. *Let T be the n-regular tree. The Neretin group $\mathrm{AAut}(T)$ is compactly generated, locally compact, simple, and does not admit any lattice.*

Proof. This is the content of [BCGM12].

Similar to how the Burger–Mozes group imposes local conditions on automorphisms, we can play the same game here and restrict the initial automorphisms. Given a subgroup $G \leq \mathrm{Aut}(T)$, define a *G-almost-automorphism* as an almost-automorphism $\varphi\colon T \setminus S_1 \to T \setminus S_2$ with the property that for every component T_v of $T \setminus S_1$ there is some $g_v \in G$ with

$$\varphi\big|_{T_v} = g_v\big|_{T_v}$$

The equivalence classes of G-almost-automorphisms then define a subgroup $\mathrm{A}G \leq \mathrm{AAut}(T)$.

Theorem 1.5.19. *Let $F \leq \mathrm{Sym}(n)$. Let $\mathcal{N}(F) = \mathrm{A}\mathcal{U}(F)$ be the group obtained as described above from the Burger–Mozes group $\mathcal{U}(F)$.*

(i) *The derived subgroup of $\mathcal{N}(F)$ is open, simple, and has finite index in $\mathcal{N}(F)$.*

(ii) *$\mathcal{N}(F)$ is compactly generated.*

(iii) *If F is a Young subgroup with strictly less than n orbits, then $\mathcal{N}(F)$ does not admit any lattice. If F has precisely n orbits, then $\mathcal{N}(F)$ does not admit any cocompact lattice.*

Consequently, if F is a Young subgroup with less than n orbits, then the derived subgroup of $\mathcal{N}(F)$ is a compactly generated, non-discrete, simple group without lattices.

Proof omitted. This is [Led17, Theorem 1.2]. $\square$

1.6 $\mathrm{CAT}(0)$ geometry

$\mathrm{CAT}(0)$ spaces were introduced by Aleksandr Aleksandrov in the 1950s, but were put in the spotlight by Mikhael Gromov, who showed that the $\mathrm{CAT}(0)$ condition alone was sufficient for a surprisingly deep study of manifolds of nonpositive sectional curvature. Gromov suggested the name CAT as an acronym for Élie Cartan, Aleksandr Aleksandrov, and Victor Toponogov. For more motivation, details, and history, we refer to [BGS85] or [BH99].

Definition 1.6.1 (geodesic segment). Let (X, d) be a metric space, i.e. a set X endowed with a distance $d\colon X \times X \to \mathbb{R}$. A *geodesic segment* (or *geodesic* for short) is a continuous isometric map $\gamma\colon [s_0, s_1] \to X$. Explicitly, for every $s_0 \leq t_0 \leq t_1 \leq s_1$ we require that $d(\gamma(t_0), \gamma(t_1)) = t_1 - t_0$.

A metric space is called *geodesic* if every two points x_0 and x_1 can be joined by a geodesic segment $\gamma\colon [s_0, s_1] \to X$ (such that $\gamma(s_0) = x_0$ and $\gamma(s_1) = x_1$). We denote such a geodesic by $[x_0, x_1]$ and note that it need not be unique.

It should be noted that this definition of geodesic segments is a global one, while usually in differential geometry a geodesic only needs to minimise distances locally.

As a familiar example, the Euclidean plane $(\mathbb{R}^2, d_E)$ with the Euclidean distance is a geodesic metric space, where the geodesic segments are line segments. The intuition behind the $\mathrm{CAT}(0)$ condition is that it rules out "fat" triangles, by comparing triangles of geodesic segments to Euclidean ones. The following definition makes this precise.

Definition 1.6.2 ($\mathrm{CAT}(0)$ space). Let (X, d) be a geodesic metric space.

(i) A *geodesic triangle* in X is a triple (x, y, z) of points together with three geodesic segments $[x, y]$, $[y, z]$, $[x, z]$. A *comparison triangle* is a triple $(\hat{x}, \hat{y}, \hat{z})$ of points in the Euclidean plane such that

$$d(x, y) = d_E(\hat{x}, \hat{y}), \qquad d(y, z) = d_E(\hat{y}, \hat{z}), \qquad d(x, z) = d_E(\hat{x}, \hat{z}).$$

For every point p on the image of $[x, y]$, there is a Euclidean *comparison point* $\hat{p}$ on the line segment $[\hat{x}, \hat{y}]$ that satisfies $d(x, p) = d_E(\hat{x}, \hat{p})$ and $d(p, y) = d_E(\hat{p}, \hat{y})$.

(ii) (X, d) is a $\mathrm{CAT}(0)$ space if for every geodesic triangle (x, y, z) and every point p on $[x, y]$, the inequality $d(p, z) \leq d_E(\hat{p}, \hat{z})$ holds. See Figure 1.3 for an illustration.

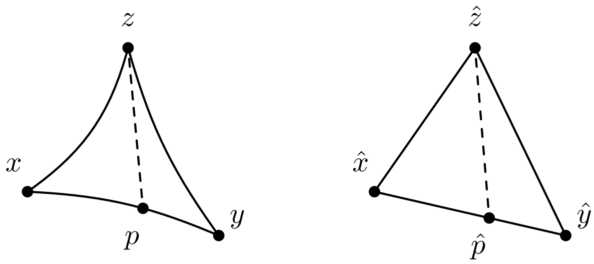

Figure 1.3. The $\mathrm{CAT}(0)$ condition on triangles and comparison triangles.

(iii) If there exists some positive constant δ such that for every geodesic triangle (x, y, z) the side $[x, y]$ is contained in the δ-neighbourhood of $[x, z] \cup [y, z]$, then we say that (X, d) is Gromov hyperbolic, or explicitly, δ-hyperbolic.

The definition immediately yields some fundamental properties.

Proposition 1.6.3. *A $\mathrm{CAT}(0)$ space is contractible and uniquely geodesic (i.e. geodesic segments joining two points are unique).*

Example 1.6.4. (i) Euclidean space $(\mathbb{R}^n, d_E)$ is a $\mathrm{CAT}(0)$ space.

(ii) A complete Riemannian manifold M, endowed with its canonical metric, is a $\mathrm{CAT}(0)$ space if and only if M has nonpositive sectional curvature. In particular real hyperbolic space $\mathbb{H}^n$ is a $\mathrm{CAT}(0)$ space.

(iii) A metric graph, obtained by replacing every edge of a graph by a line segment of unit length, is a $\mathrm{CAT}(0)$ space if and only if the original graph is a tree (i.e. a connected acyclic graph).

An example of the power of the $\mathrm{CAT}(0)$ condition is the following fixed point theorem, originally proven by François Bruhat and Jacques Tits for affine buildings.

Theorem 1.6.5 (Bruhat–Tits). *Let G be a group acting on a complete $\mathrm{CAT}(0)$ space (X, d) by isometries. If G has a bounded orbit, then the fixed point set of G is a nonempty convex subset of X.*

Proof omitted. This is [BH99, Corollary 2.8].

One defining property of $\mathrm{CAT}(0)$ spaces is that they have a well-behaved structure at infinity.

Definition 1.6.6 (visual boundary). Let (X, d) be a $\mathrm{CAT}(0)$ space. Then we define two geodesic rays $\gamma, \gamma' \colon \mathbb{R}^+ \to X$ to be equivalent if and only if the distance $\operatorname{dist}(\gamma(t), \gamma'(t))$ remains bounded (where t ranges over $\mathbb{R}^+$). The rays γ and γ' are said to be *asymptotic*. The equivalence class of γ is commonly denoted by $\gamma(\infty)$.

The visual boundary of X is then the set of all equivalence classes of geodesic rays, and is denoted by ∂X. Note that the images of asymptotic rays under isometries of X are again asymptotic. Thus, isometries of X extend to bijections of $X \cup \partial X$. For a topologification of $X \cup \partial X$ compatible with the $\mathrm{CAT}(0)$ structure on X, we refer to [BH99, Chapter II.8] or [Cap14b].

Definition 1.6.7 (minimal, geometrically dense). Let G be a group acting on a $\mathrm{CAT}(0)$ space (X, d) by isometries. The action is said to be *minimal* if there is no nontrivial G-invariant closed convex subset of X. The action is *geometrically dense* if it is minimal and if moreover G does not fix a point in ∂X.

Compare Definitions 1.4.1 and 1.6.7; every automorphism of a tree in the classical sense induces an isometry of its $\mathrm{CAT}(0)$ realisation as a simplicial tree, and the "combinatorial" notions of minimal or geometrically dense actions agree with the "geometrical" notions on the metric realisation.

1.7 Building theory

The roots of buildings lie in the field of algebraic group theory, where Jacques Tits — the architect of building theory — constructed a uniform framework for understanding certain algebraic groups in a geometrical way.

Driven by Felix Klein's *Erlangen* program from the 1870s, generations of mathematicians classified geometries in terms of their groups of automorphisms. The standard examples are the projective linear groups, whose structure captures projective geometries. Similar geometries were constructed for orthogonal, symplectic, and unitary groups, and the development of the theory of semisimple groups allowed Tits to construct geometries for the exceptional groups as well.

A central insight in the 1950s was that to a general semisimple algebraic group G, one can associate a finite reflection group called the *Weyl group* W of G. It is obtained as the quotient N/T, where T is a maximal torus and N its normaliser. Another key ingredient in understanding the structure of G is that of a *Borel subgroup* B and the parabolic subgroups. Using such a Borel subgroup, one obtains a decomposition $G = BWB$ and even more precisely, a one-to-one correspondence between W and the set $B\backslash G/B$ of double cosets of B. This is the so-called *Bruhat decomposition* of G.

Using the parabolic subgroups as building blocks, one can then define a geometry whose structure is governed by the Weyl group. One recovers e.g. the projective geometries from projective groups in this fashion, but from purely algebraic data, and in a vastly more general setting. Tits studied how the algebraic properties restricted the geometries, initially over the complex field $\mathbb{C}$ but gradually extending to more general fields, and eventually gave several axiomatisations of the resulting geometries, which he called *buildings*.

Tits presented an outline of the theory in a 1965 Bourbaki Seminar ([Tit65]) and provided a more extensive picture in the book *Buildings of spherical type and finite BN-pairs* ([Tit74]). At that time, Tits thought of buildings as simplicial complexes with a distinguished family of subcomplexes called *apartments*. In this point of view, the *chambers* are the simplices of maximal dimension. Since then other viewpoints have been developed, the most succesful one following a suggestion of Luis Puig, taking chambers as fundamental objects and founding the theory on the more abstract framework of *chamber systems*. Tits presented this approach in [Tit81]. Throughout this thesis, we will use this more modern viewpoint.

1.7.1 Chamber systems

We briefly go through the necessary definitions and some examples, and refer to one of the standard works such as [Ron09] or [Wei03] for more background. A warning however — in the literature, the precise definition of a chamber system varies somewhat, as one might or might not allow for certain degenerate situations. There should be no confusion, as we will not encounter such degeneracies, but in any case, our approach is based on [Ron09].

Definition 1.7.1 (chamber system). Let I be any index set. A *chamber system over I* is a set Δ together with, for every $i \in I$, an equivalence relation called *i-adjancence*. The elements of Δ are called *chambers*. If two chambers c and d are i-adjacent, we write $c \sim_i d$, or simply $c \sim d$ if we do not want to stress the adjacency type. The cardinality $|I|$ is called the *rank* of Δ. In this thesis, the rank will always be finite.

Just as with groups, we usually say "Δ is a chamber system" when the equivalence relations on Δ are clear from context .

Convention 1.7.2. There are a few commonly used but slightly nonconcurring ways to represent chamber systems. The original simplicial approach to building theory used the notion of a *chamber complex* — a sufficiently connected simplicial complex in which all maximal simplices have the same dimension. The chambers are these maximal simplices, and chambers are adjacent if they intersect in a face of codimension one. Unfortunately, only chamber complexes of low rank can be visualised satisfactorily using this approach. Our combinatorial approach treats the chambers as core objects. It is then more natural to visualise chamber systems as graphs, where the chambers are the vertices and the different types of adjacency are represented by edges of different colours. Both viewpoints have their benefits and we feel the need to switch occasionally. However, whichever viewpoint we use, *chambers will always be represented in black* (either as vertices, edges, simplices) and *different types of adjacency will always be represented by different colours.*

Definition 1.7.3 (gallery). A *gallery* γ is a finite sequence of pairwise adjacent chambers

$$c_0 \sim_{i_1} c_1 \sim_{i_2} \cdots \sim_{i_n} c_n$$

for certain $i_1, \ldots, i_n \in I$. We call the word $i_1 \cdots i_n$ (an element of the free monoid I^*) the *type* of γ, and the integer n the *length* of γ. If there is no strictly shorter gallery from c_0 to c_n we call γ a *minimal* gallery. Note that we colloquially say the gallery to join the chambers c_0 and c_n or to go from c_0 to c_n — this should not cause any confusion.

Chamber systems are naturally metric spaces, where the metric

$$\mathrm{dist}\colon \Delta \times \Delta \to \mathbb{N} \cup \{\infty\}$$

is defined by declaring $\mathrm{dist}(c, d)$ to be the minimal length of all galleries joining c and d (or infinity if there is no such gallery). It is clear that this distance function is positive-definite, symmetrical, and satisfies the triangle inequality.

Definition 1.7.4 (ball, sphere). Let $c \in \Delta$ and $n \in \mathbb{N}$. Then we define the sets

$$\mathsf{B}_n(c) = \{d \in \Delta \mid \mathrm{dist}(c, d) \leq n\}, \qquad \mathsf{S}_n(c) = \{d \in \Delta \mid \mathrm{dist}(c, d) = n\},$$

that we call the *ball* and the *sphere*, respectively, with *centre* c and *radius* n.

Definition 1.7.5 (convex). A subset $C \subseteq \Delta$ is called *(combinatorially) convex* if any gallery between two chambers in C lies entirely in C.

Definition 1.7.6 (panel, residue). Let $J \subseteq I$. A subset $C \subseteq \Delta$ is called *J-connected* if any two chambers in C can be joined by a gallery of type in J^*. A *residue of type J*, or simply a *J-residue*, is a J-connected components of Δ. A *panel of type j*, or simply a *J-panel*, is a residue of type $\{j\}$. A residue of type J is also said to have *cotype $I \setminus J$*.

For a given chamber c and type $J \subseteq I$, we will denote the residue of type J containing c by $\mathcal{R}_J(c)$. When $J = \{j\}$, we will denote the panel of type j containing c by $\mathcal{P}_j(c)$. The set of all J-residues of the chamber system Δ will be denoted by $\mathrm{Res}_J(\Delta)$.

Note that a J-residue is in its own right a connected chamber system over the index set J.

Definition 1.7.7 (thin, thick). A chamber system is called *thin* if every panel contains exactly two chambers, and *thick* if every panel contains at least three chambers. (Panels containing only a single chamber are degenerate cases that should not occur in any reasonable application.)

Note that a chamber system does not have to be either thin or thick.

Definition 1.7.8 (morphism). A map $\varphi\colon \Delta_1 \to \Delta_2$ between two chamber systems is a *morphism* if $\varphi(c) \sim \varphi(d)$ in Δ_2 whenever $c \sim d$ in Δ_1. As usual, an *isomorphism* is a bijective morphism, and an *automorphism* is an isomorphism to the same chamber system. Assuming that Δ_1 and Δ_2 have the same index set, a morphism is *type-preserving* if more precisely $\varphi(c) \sim_i \varphi(d)$ whenever $c \sim_i d$. In this thesis, we shall always assume morphisms to be type-preserving.

The set of all automorphisms of a chamber system Δ forms a group, denoted by $\mathrm{Aut}(\Delta)$.

Example 1.7.9. The basic notion of a chamber system has very few restrictions and examples are plentiful, but it helps to keep in mind a more substantial example like (vi) below when building up the theory.

(i) A chamber system of rank zero is simply a set of chambers with no adjacencies whatsoever.

(ii) By visualising the chambers as vertices of a graph and the unique adjacency relation as edges, a chamber system of rank one is nothing more than the disjoint union of complete graphs.

(iii) A chamber system of rank one can be visualised as the disjoint union of complete graphs, where the chambers are vertices and the adjacency is determined by the edges.

(iv) Chamber systems of rank two are essentially the same as bipartite graphs, as the following example from [Wei03, Example 1.8] shows. Let $\Gamma = (V, E)$ be a graph with a bipartition into "white" and "black" vertices, $V = V_\circ \sqcup V_\bullet$. Define a chamber system Δ_Γ over the set $\{\circ, \bullet\}$ with chamber set E by declaring two chambers (i.e. edges of Γ) to be $\circ$-adjacent if they share an endpoint in $V_\circ$ and $\bullet$-adjacent if they share an endpoint in $V_\bullet$. In this chamber system, panels in Δ_Γ naturally correspond to vertices of Γ.

    Conversely, let Δ be a chamber system of rank two, and define a graph Γ_Δ with the panels of Δ as the vertex set. Join two vertices in Γ_Δ by an edge if and only if the corresponding panels have a nonempty intersection. Note that such panels necessarily have different types, so that Γ_Δ is a bipartite graph.

    Up to isomorphism and relabeling of the index set, these two constructions are inverses.

(v) If Δ_1 and Δ_2 are two chamber systems over I_1 and I_2 respectively, then their direct product $\Delta_1 \times \Delta_2$ is naturally a chamber system over $I_1 \sqcup I_2$ by declaring $(c_1, c_2) \sim_i (d_1, d_2)$ if either $c_1 \sim_i d_1$ and $c_2 = d_2$ (for $i \in I_1$) or $c_1 = d_1$ and $c_2 \sim_i d_2$ (for $i \in I_2$). Residues in this direct product are then direct products of residues.

(vi) Usually in geometry, projective space of dimension n consists of objects of n different types (points, lines, $\ldots$) and one universal incidence relation. In order to recast this into a chamber system, we need a single set of homogeneous objects, equipped with n equivalence relations. Define the chambers to be the *maximal flags*, i.e. sets of exactly n mutually incident objects. For every $0 \leq i < n$, we declare two flags to be i-adjacent if they differ only in their element of dimension i. This results in a thick chamber system of rank n.

    A panel of type i then corresponds to a collection of subspaces of dimension i, sandwiched between two given incident subspaces of dimension $i-1$ and $i+1$. The residues of cotype i can be identified with the i-subspaces.

For the Fano plane $\mathrm{PG}(2,2)$, there are 21 flags in total. A visualisation of the corresponding chamber system is given in Figure 1.4. The Levi graph of the Fano plane is the bipartite graph with the points and lines as vertices and edges defined by incidence. The 21 edges correspond to the chambers, and there are two types of adjacency, defined by the bipartition.

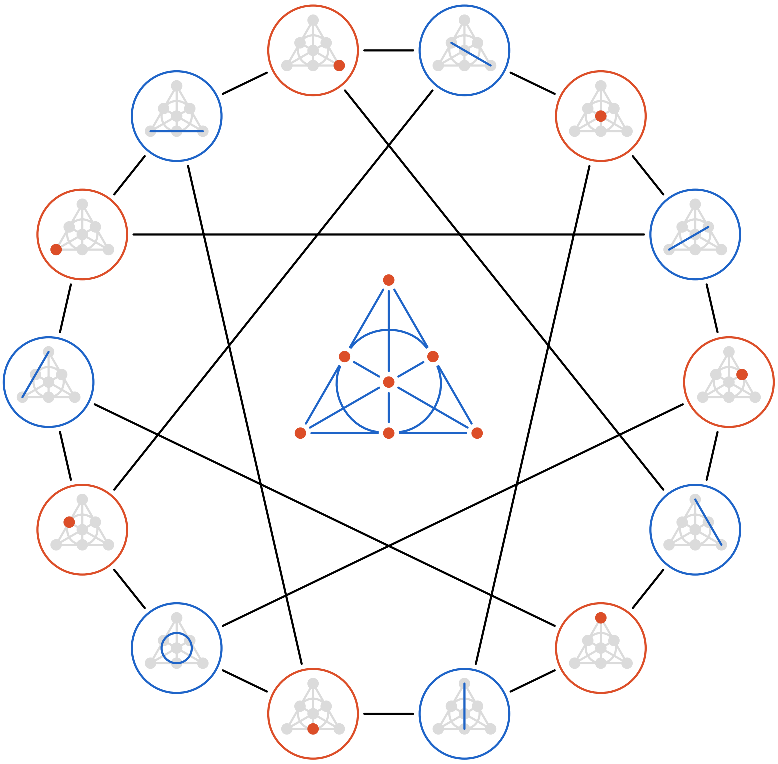

Figure 1.4. The Fano plane when interpreted as a chamber system.

1.7.2 Coxeter systems

The next step towards buildings is the definition of Coxeter complexes. These complexes will turn out to be precisely the thin buildings.

Definition 1.7.10 (Coxeter system). Let I be any index set and M a function

$$M\colon I \times I \to \mathbb{N} \cup \{\infty\}\colon (i,j) \mapsto m_{ij}$$

satisfying $m_{ii} = 1$, $m_{ij} \geq 2$, and $m_{ij} = m_{ji}$ for all $i \neq j \in I$. Then the *Coxeter group of type M* is the group defined by the presentation

$$W = \left\langle s_i \;\middle|\; (s_i s_j)^{m_{ij}} = 1 \text{ for all } i,j \in I \right\rangle.$$

When $m_{ij} = \infty$, this means that no relation on $s_i s_j$ should be imposed. Note that the assumption that $m_{ii} = 1$ for all $i \in I$ immediately implies that the generators s_i are involutions. Additionally, note that when $m_{ij} = 2$, the generators s_i and s_j commute.

Together with the generating set $S = \{s_i \mid i \in I\}$, the pair (W,S) is called the *Coxeter system of type M*. The *rank* of (W,S) is the cardinality of I.

We can represent M be means of its *Coxeter matrix* (m_{ij}), or more commonly its *Coxeter diagram*: the nodes of the diagram are the elements of I (sometimes with explicit labels), and two nodes are connected by a decorated edge according to the following rules:

One caveat is that the set S in a Coxeter *system* is not uniquely determined by the Coxeter *group* alone; for instance, the two diagrams below give rise to isomorphic Coxeter groups. When dealing with general Coxeter systems, it is hence of importance to specify the generating set (although this might be superfluous for certain specific families; see Theorem 2.1.2 below).

A priori from the presentation, the order of the product $s_i s_j$ of two generators is bounded by m_{ij}. Tits constructed a linear representation $W \to \mathrm{GL}(\mathbb{R}^I)$ from which one can read off more.

Proposition 1.7.11. *Let (W, S) be a Coxeter system. Then the order of $s_i s_j$ equals m_{ij}.*

Proof omitted. We refer to [Ron09, Lemma 2.1], [Wei03, Theorem 2.3], [Dav12, Corollary 6.12.6], or the original article by Tits [Tit69, Corollaire 2].

Let us mention a few important families.

Definition 1.7.12 (irreducible). We call a Coxeter system (W, S) *irreducible* if the underlying graph of its Coxeter diagram is connected, and *reducible* otherwise.

If a diagram has two or more connected components, the Coxeter group is isomorphic to the direct product of the Coxeter groups associated to the individual components, since generators corresponding to different components pairwise commute. Usually, we may thus restrict our attention to the irreducible case.

Definition 1.7.13 (spherical). We call a Coxeter system (W, S) *spherical* if W is a finite group.

As briefly sketched in Section 1.7, Weyl groups of semisimple linear algebraic groups are spherical Coxeter groups. Historically, these groups were realised as reflection groups of the unit sphere in Euclidean space of dimension n, where n is the rank of (W, S) — hence the name.

Definition 1.7.14 (right-angled). We call a Coxeter system (W, S) *right-angled* if for all $i \neq j$, we have that $m_{ij} \in \{2, \infty\}$.

Definition 1.7.15 (Coxeter complex). Let (W, S) be a Coxeter system of type M over I. Define a chamber system over I with the elements of W as chambers, and declare two group elements v and w to be i-adjacent if and only if $vs_i = w$. Then i-adjacency is indeed an equivalence relation since s_i is an involution for all $i \in I$. The resulting chamber system is called the *Coxeter complex of type M*.

It is worth emphasising that Coxeter complexes are always connected and thin: every chamber is i-adjacent to exactly one other chamber for every $i \in I$.

The Coxeter complex associated to a Coxeter system (W, S) is nothing more than the (undirected) Cayley graph of W with respect to the generating set S — vertices corresponding to chambers and i-adjacency corresponding to edges with label s_i. In particular, W acts as the automorphism group of the Coxeter complex.

Example 1.7.16. It helps to have a few examples in mind.

(i) Consider the following diagram with n nodes, commonly denoted by A_n.

This corresponds to the group presentation $\langle s_1, \ldots, s_n \mid \mathcal{R}\rangle$, where $\mathcal{R}$ is the set of relators

$$\begin{aligned} {s_i}^2 &= 1 \quad \text{for all } 1 \leq i \leq n, \\ (s_i s_{i+1})^3 &= 1 \quad \text{for all } 1 \leq i < n, \\ (s_i s_j)^2 &= 1 \quad \text{for all } 1 \leq i < j \leq n \text{ such that } |i-j| \geq 2. \end{aligned}$$

With a bit of effort, one can show this is a presentation for the symmetric group $\mathrm{Sym}(n+1)$ on $n+1$ elements. We briefly sketch a proof. There exists a well-defined morphism from W to the symmetric group taking generator s_i to transposition $(i \ \ i+1)$. Indeed: transpositions are involutions, *disjoint* transpositions commute, and products of two *adjacent* transpositions have order 3. Since these n transpositions generate $\mathrm{Sym}(n+1)$, the morphism is surjective. Finally, with an inductive argument one shows that the order of W is bounded by $(n+1)!$, so we have, in fact, an isomorphism.

The associated Coxeter complex is a skeleton of a *permutohedron*, an n-dimensional polytope whose k-dimensional faces correspond to rank k residues for all $0 \leq k \leq n$. Figure 1.5 gives a visualisation for $n = 2$, where the permutohedron is a hexagon, and for $n = 3$, where the permutohedron is a truncated cuboctahedron.

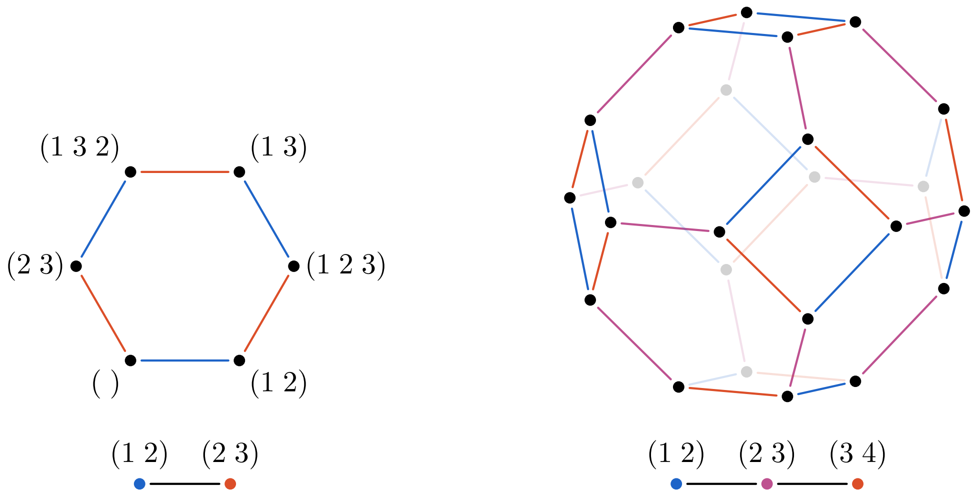

Figure 1.5. The Coxeter complexes of type A_2 and A_3.

(ii) Consider the following diagram with $n \geq 2$, commonly denoted by $\mathsf{I}_2(n)$. For $n \in \{2, 3, 4\}$, instead of a label n our notational convention would dictate either no edge, a single edge, or a double edge, but no confusion should be possible.

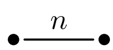

For finite n, this corresponds to the presentation $\langle s, t \mid s^2 = t^2 = (st)^n = 1\rangle$ — a familiar presentation for the finite dihedral group $\mathbf{D}_{2n}$ (the group of symmetries of a regular n-gon). Geometrically, s and t can be represented as reflections in Euclidean space through lines that make an angle of π/n. Figure 1.6 shows the associated Coxeter complex for $n = 4$; in general the complex is a cycle of length $2n$.

Note that diagrams $\mathsf{I}_2(3)$ and A_2 are identical, explaining the isomorphism $\mathrm{Sym}(3) \cong \mathbf{D}_6$.

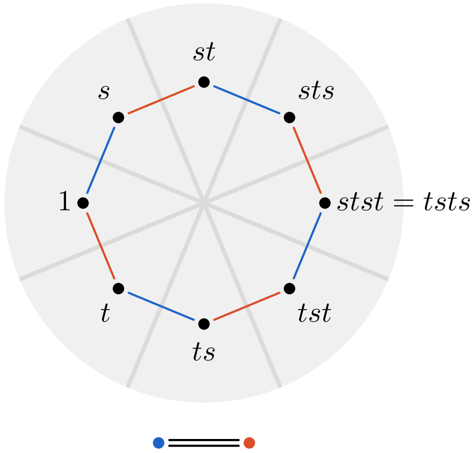

Figure 1.6. The Coxeter complex of type $\mathsf{I}_2(4)$.

For infinite n, we find a presentation $\langle s, t \mid s^2 = t^2 = 1 \rangle$ for the infinite dihedral group $\mathbf{D}_\infty$ (the group of symmetries of the integers, or of a regular apeirogon). Geometrically, s and t can be represented as Euclidean reflections through parallel lines. The Coxeter complex is a two-way infinite path, as depicted in Figure 1.7.

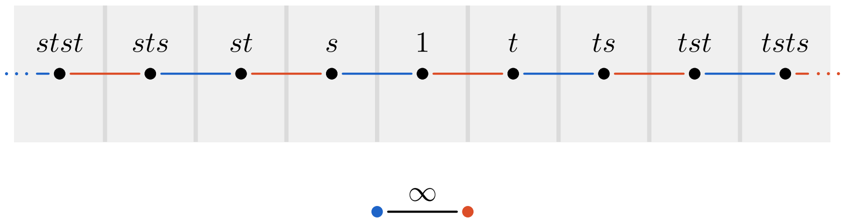

Figure 1.7. The Coxeter complex of type $\mathsf{I}_2(\infty)$.

(iii) An important family of infinite Coxeter groups is given by the affine ones. Just like *spherical* Coxeter groups of rank $n+1$ are characterised as groups of reflections of an n-sphere, *affine* Coxeter groups of rank $n+1$ are groups of reflections of n-dimensional affine space. These groups are extensions of an abelian group by a spherical Coxeter group.

Type $\mathsf{I}_2(\infty)$ is one example. Figures 1.8 and 1.9 give two other examples, namely $\tilde{\mathsf{A}}_2$ and $\tilde{\mathsf{G}}_2$. The Coxeter complex of the former corresponds to the hexagonal tiling of the Euclidean plane (or its dual, the triangular tiling); the complex of the latter to the truncated trihexagonal tiling (or its dual, the kisrhombille tiling).

(iv) Figure 1.10 is the direct product of two systems of type $\mathsf{I}_2(\infty)$ and an example of a reducible right-angled Coxeter system. Note that the product is again an affine Coxeter group; the Coxeter complex corresponds to a simple (self-dual) square tiling.

(v) Figure 1.11 is an example of an irreducible right-angled Coxeter system.

(vi) Figure 1.12 visualises an irreducible right-angled Coxeter complex embedded in the hyperbolic plane as the dual of a tessellation with regular right-angled pentagons. This Coxeter system of rank five is the building block for the *Bourdon building* of type $I_{5,2}$ as defined in [Bou97].

Whenever a group is given by a presentation, a natural question is to understand its combinatorial behaviour. Coxeter groups turn out to be quite tame: not only did Tits find a satisfying theoretical solution to the word problem, but also surprisingly efficient practical algorithms for computations in Coxeter groups have been developped as well, building further upon Tits's geometrical approach (see, for instance, [Cas02]). In order to introduce this, we need a couple of definitions.

Definition 1.7.17. Let (W, S) be a Coxeter system over some index set I. Then there is a natural surjective *evaluation morphism* of monoids from the free monoid on I to the group W defined by

$$\varsigma \colon I^* \to W : i \mapsto s_i.$$

Definition 1.7.18. For every $i \neq j$ such that m_{ij} is finite, define in I^* the word

$$p(i,j) = \begin{cases} (ij)^k & \text{if } m_{ij} = 2k \text{ is even,} \\ j(ij)^k & \text{if } m_{ij} = 2k+1 \text{ is odd.} \end{cases}$$

In other words, $p(i,j)$ is the word with m_{ij} alternating letters i and j, ending in j. When $m_{ij} = \infty$, $p(i,j)$ remains undefined.

We can now describe a couple of elementary operations on words in I^* that will nicely relate to the group structure of W.

Definition 1.7.19 (homotopy). Let $i, j \in I$ and $w_1, w_2 \in I^*$.

(i) An *elementary homotopy* (or also a *braid relation*) is a transformation of a word $w_1\, p(i,j)\, w_2$ into the word $w_1\, p(j,i)\, w_2$. Two words w and w' are *homotopic* if w can be transformed into w' by a sequence of elementary homotopies; we denote this by $w \simeq w'$. Clearly, homotopy is an equivalence relation and preserves the length of the words.

(ii) An *elementary contraction* is a transformation of a word $w_1\, ii\, w_2$ into the word $w_1\, w_2$.

(iii) An *elementary expansion* is a transformation of a word $w_1\, w_2$ into a word $w_1\, ii\, w_2$.

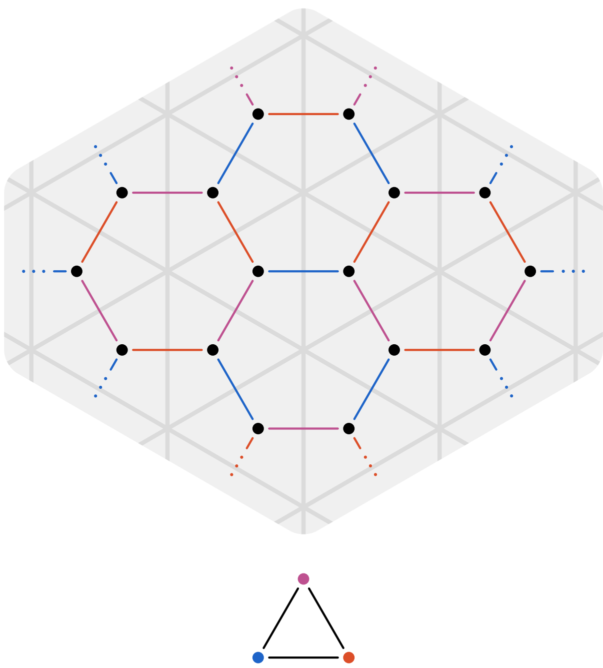

Figure 1.8. The Coxeter complex of type $\tilde{\mathrm{A}}_2$.

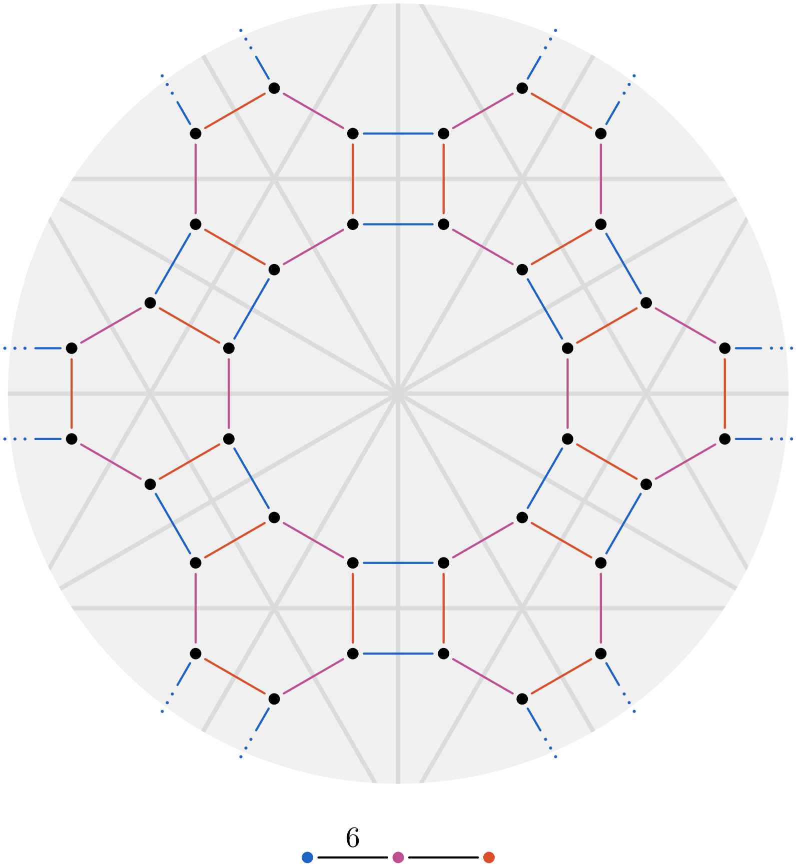

Figure 1.9. The Coxeter complex of type $\tilde{\mathsf{G}}_2$.

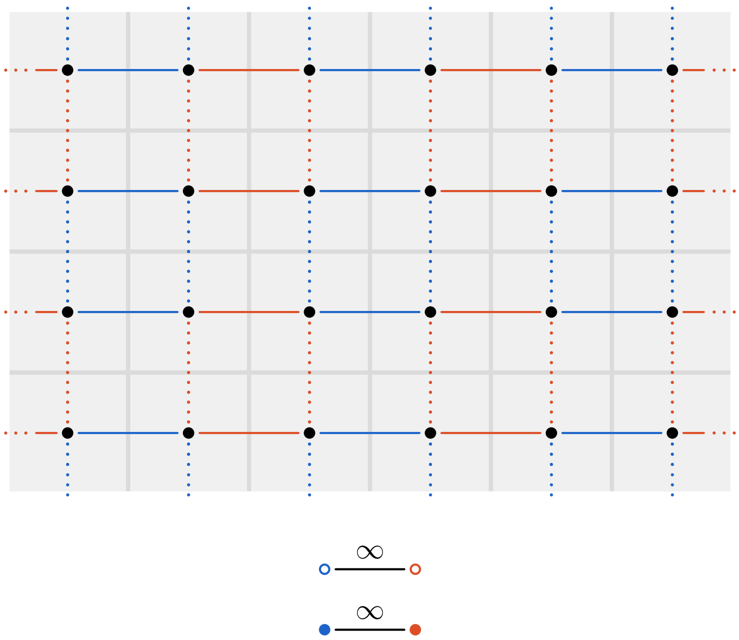

Figure 1.10. The Coxeter complex of type $\mathsf{I}_2(\infty) \times \mathsf{I}_2(\infty)$.
Nodes ● and ● correspond to solid edges, nodes ○ and ○ to dotted edges.

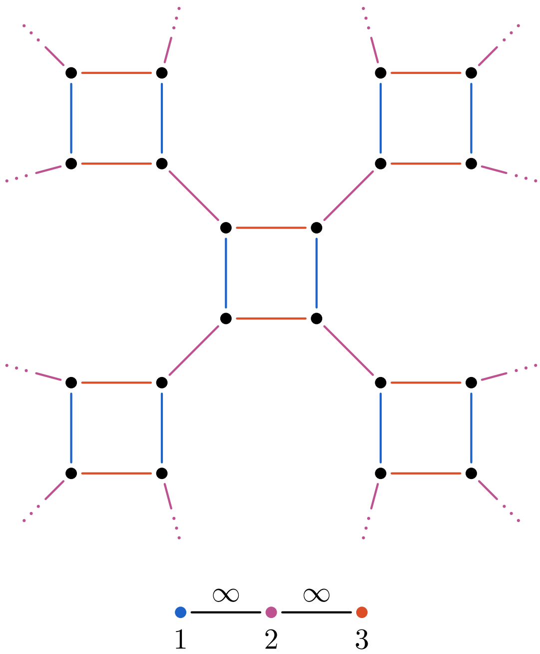

Figure 1.11. A right-angled Coxeter complex.

A word is called *reduced* if it is not homotopic to a word of the form $w_1\, ii\, w_2$ (for some $i \in I$). Two words w and w' are called *equivalent* if w can be transformed into w' by a sequence of elementary homotopies, contractions, and expansions. Clearly, every equivalence class contains some reduced word.

It helps to think of these operations in geometrical terms. Choosing any initial chamber, a word w in I^* can be thought of as a gallery of type w in the Coxeter complex. An elementary contraction then removes a backtracking $g \sim_i gs_i \sim_i g$, while an elementary expansion introduces one. Moreover, residues of type $\{i, j\}$ correspond to a cycle of length $2m_{ij}$ (by Proposition 1.7.11), and an elementary homotopy transforms a gallery of length m_{ij} in this cycle (i.e. going to an opposite chamber) into the gallery that "goes the other way around".

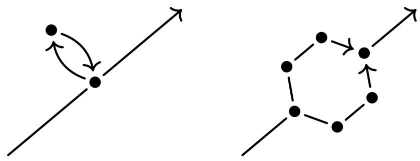

Theorem 1.7.20. (i) *Two words w and w' are equivalent if and only if $\varsigma(w) = \varsigma(w')$.*

(ii) *Two reduced words w and w' are equivalent if and only if they are homotopic.*

(iii) *A gallery in the Coxeter complex is minimal if and only if its type is reduced.*

(iv) *If for some $i \in I$ a word w is reduced but iw (or wi) is not, then w is homotopic to a word that begins (or ends, respectively) with i.*

Proof. By the defining relations $(s_i s_j)^{m_{ij}} = 1$ in the presentation, $p(i, j)$ and $p(j, i)$ have the same image under ς, and $\varsigma(ii)$ is the identity. Statement (i) follows immediately. For (ii) and (iii), we refer to [Ron09, Theorem 2.11]. Statement (iv) is [Ron09, Corollary 2.13]. ∎

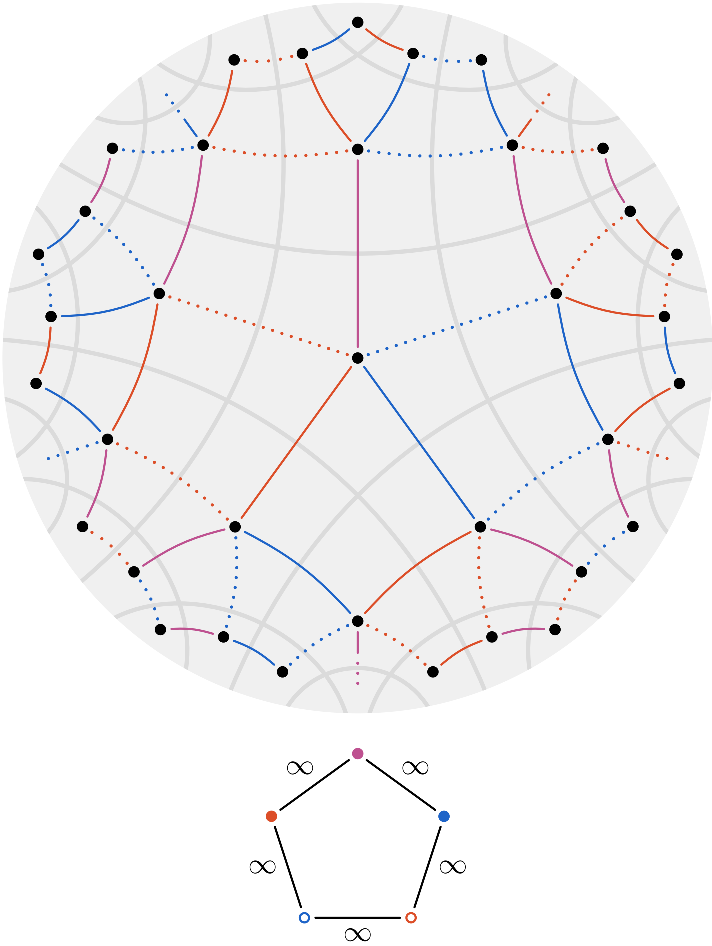

Figure 1.12. A right-angled Coxeter complex in the hyperbolic plane. Nodes ● and ● correspond to solid edges, nodes ○ and ○ to dotted edges.

As a corollary, Theorem 1.7.20 shows the word problem in Coxeter groups to be decidable: in order to check whether two words w and w' represent the same group element, it suffices to reduce them and check whether or not the reduced words are homotopic. This is decidable, since homotopies preserve word lengths; a brute-force algorithm could (in principle at least) finish the job. Obviously for practical purposes a more efficient implementation is desirable, but the geometry continues to play an important role (see e.g. [Cas02]).

Coxeter systems have a natural notion of subsystems, as the following proposition shows.

Proposition 1.7.21. *Let (W, S) be a Coxeter system of type M over I. Let $J \subseteq I$ be any subset, let $S_J = \{s_j \mid j \in J\} \subseteq S$ and $W_J = \langle S_J \rangle \leq W$, and let M_J the subdiagram of M induced by J.*

(i) *If $s_i \in W_J$ for some $i \in I$, then $i \in J$. In other words, $S_J = S \cap W_J$.*

(ii) *If $w \in I^*$ is a reduced word with $\varsigma(w) \in W_J$, then actually $w \in J^*$.*

(iii) *(W_J, S_J) is a Coxeter system of type M_J over J.*

Proof omitted. We refer to [Ron09, Lemma 2.1, Lemma 2.10, Corollary 2.14]. $\not\square$

Corollary 1.7.22. *Reusing the notation from Proposition 1.7.21, let $\mathcal{R}$ be a J-residue in the Coxeter complex of (W, S). Then $\mathcal{R}$ is convex, and $\mathcal{R}$ is (isomorphic to) a Coxeter complex of type M_J.*

We conclude this section with another important property of residues in Coxeter complexes, known as the *gate property.*

Proposition 1.7.23. *Let $\mathcal{R}$ be a residue and c be a chamber of a Coxeter complex. Then there exists a unique chamber $d \in \mathcal{R}$ such that $\operatorname{dist}(c, d)$ is minimal. Moreover, for every chamber $c' \in \mathcal{R}$, there exists a minimal gallery from c to c' via d.*

Proof omitted. We refer to [Ron09, Theorem 2.9]. $\not\square$

1.7.3 General buildings

We can finally state the definition of a building (using the "combinatorial" approach from [Tit81]).

Definition 1.7.24 (building). Let (W, S) be a Coxeter system of type M over some index set I. A *building* (Δ, δ) *of type* M is a chamber system Δ over I such that every panel contains at least two chambers, and equipped with a map $\delta\colon \Delta \times \Delta \to W$ satisfying the following property for every reduced word $w \in I^*$:

$$\delta(c, d) = \varsigma(w) \text{ if and only if } c \text{ and } d \text{ can be joined by a gallery of type } w.$$

Such a gallery is automatically minimal by Theorem 1.7.20. In particular, the distance between two chambers c and d is exactly the length of $\delta(c, d)$ in the word metric of W (w.r.t. generating set S).

Intuitively, the map δ not only measures a notion of "distance" between chambers, but in addition a notion of "direction" as well, with values in the group W. This group is called the *Weyl group* of the building, and the map δ is hence also called the *W-distance* or *Weyl distance function.*

For brevity we shall usually identify the building with its chamber set and abbreviate (Δ, δ) to Δ.

Why do we need to restrict w to *reduced* words in the definition? Suppose that $\gamma = c_0 \sim_i c_1 \sim_i c_2$ is a gallery of type ii. Then two possibilities arise: either $c_0 = c_2$ in which case we can replace γ by the null gallery, or $c_0 \neq c_2$ in which case we can replace γ by the gallery $c_0 \sim_i c_2$ of type i. Hence an elementary contraction on the word level does not lift to a meaningful operation on the gallery level: in general, a gallery of type $w_1\, ii\, w_2$ cannot be replaced by one of type $w_1\, w_2$.

We bundle a couple of immediate consequences of the definition.

Proposition 1.7.25. *Let (Δ, δ) be a building.*

(i) *Δ is connected and δ is surjective;*

(ii) *$\delta(c, d) = \delta(d, c)^{-1}$ for all chambers $c, d \in \Delta$;*

(iii) *$\delta(c, d) = 1$ if and only if $c = d$;*

(iv) *$\delta(c, d) = s_i$ if and only if $c \neq d$ and $c \sim_i d$;*

(v) *i-adjacency and j-adjacency are mutually exclusive for $i \neq j$;*

(vi) *a gallery in Δ of type w is minimal if and only if w is reduced;*

(vii) *if w is a (not necessarily reduced) word such that $\delta(c, d) = \varsigma(w)$, then there exists a gallery of type w from c to d;*

(viii) *if γ_1 is a (not necessarily minimal) gallery of type w_1 from c to d and w_1 is homotopic to w_2, then there exists a gallery γ_2 of type w_2 from c to d;*

(ix) *if w is reduced, then a gallery of type w between two chambers is unique.*

Definition 1.7.26 (semiregular). A building Δ over I is called *semiregular with parameters* $(q_i)_{i \in I}$ if for every $i \in I$ the panels of type i all have the same (possibly infinite) cardinality $q_i \geq 2$. Note that the thin buildings are precisely the semiregular buildings with parameters $q_i = 2$.

Example 1.7.27. The associated Coxeter complex of a Coxeter system (W, S) of type M is itself a building of type M, by setting $\delta(g, h) = g^{-1}h$ for all $g, h \in W$. This is a thin building.

Since distances in a building are governed by δ and the underlying Coxeter system, the definition of an *isometry* should take into account the map δ as well.

Definition 1.7.28 (isometry). Let (Δ_1, δ_1) and (Δ_2, δ_2) be two buildings of the same type, and let $C \subseteq \Delta_1$ be a subset. A map $\varphi\colon C \to \Delta_2$ is called an *isometry* if $\delta_1(c, d) = \delta_2(\varphi(c), \varphi(d))$ for all $c, d \in C$, or in other words, if the following diagram commutes.

$$\begin{tikzcd} C \times C \arrow[r, "\varphi\times\varphi"] \arrow[d, "\delta_1"] & \Delta_2 \times \Delta_2 \arrow[d, "\delta_2"] \\ W \arrow[r, equal] & W \end{tikzcd}$$

Definition 1.7.29 (isomorphism). Let (Δ_1, δ_1) and (Δ_2, δ_2) be two buildings of the same type. A *building isomorphism* $\Delta_1 \to \Delta_2$ (or shortly an *isomorphism* when clear from context) is an isomorphism of chamber systems that is moreover an isometry.

While chamber system isomorphisms only take into account “local” data (i-adjacencies) and building isometries more “global” data (images of δ), there is still a strong correspondence between the two notions when considering full-scale buildings, as the following proposition shows.

Proposition 1.7.30. *Let (Δ_1, δ_1) and (Δ_2, δ_2) be buildings of the same type. Then $\varphi\colon \Delta_1 \to \Delta_2$ is an isometry if and only if φ is a (type-preserving) chamber system isomorphism from Δ_1 to $\varphi(\Delta_1)$.*

Proof omitted. This is [Wei03, Proposition 8.2]. ⧄

Definition 1.7.31 (apartment). Let Δ be a building of type M and let (W, S) be the underlying Coxeter system of type M. Using Example 1.7.27, we can then consider isometries $\varphi\colon W \to \Delta$. An isometric image $\varphi(W)$ of W is called an *apartment* of Δ.

By Proposition 1.7.30, an apartment on its own is a thin building of type M. The following proposition and its corollary show that apartments are plentiful in buildings.

Proposition 1.7.32. *An isometry of a subset $C \subseteq W$ into Δ extends to an isometry of W into Δ.*

Proof omitted. This is [Ron09, Theorem 3.6]. ⧄

Corollary 1.7.33. *Every two chambers of a building are contained in a common apartment.*

Proof. Let c and d be two distinct chambers in Δ. Set $g = \delta(c, d)$. Then the map $\varphi\colon \{1, g\} \to \Delta$ with $\varphi(1) = c$ and $\varphi(g) = d$ can be extended to an isometry $\bar{\varphi}\colon W \to \Delta$ by Proposition 1.7.32. The image $\bar{\varphi}(W)$ is an apartment containing both c and d. ■

Proposition 1.7.34. *Let $\mathcal{A}$ be an apartment of Δ and let $c \in \mathcal{A}$ be a chamber. Then there is a unique isometry $\varphi\colon W \to \Delta$ such that $\varphi(1) = c$ and $\varphi(W) = \mathcal{A}$.*

Proof. Suppose φ' is another such isometry. Then $\varphi^{-1}\varphi'$ is an isometry of W and fixes the identity, hence $\varphi^{-1}\varphi'$ is the identity map. ■

Definition 1.7.35 (retraction). Let c be a chamber in an apartment $\mathcal{A}$ of Δ. Let φ be the isometry from Proposition 1.7.34. Then we define the map

$$\rho_{c,\mathcal{A}}\colon \Delta \to \mathcal{A}\colon d \mapsto \varphi(\delta(c, d)),$$

called the *rectraction* of Δ *onto* $\mathcal{A}$ with *centre* c.

Using retractions one can then show the following.

Proposition 1.7.36. *Apartments are convex.*

Proof omitted. This is a special case of [Ron09, Theorem 3.8]. ⧄

As a corollary, we can lift Proposition 1.7.23 to the realm of buildings.

Corollary 1.7.37 (gate property). *Let $\mathcal{R}$ be a residue and c be a chamber of a building. Then there exists a unique chamber $d \in \mathcal{R}$ such that $\operatorname{dist}(c, d)$ is minimal. Moreover, for every chamber $c' \in \mathcal{R}$, there exists a minimal gallery from c to c' via d.*

Definition 1.7.38 (projection). Let $\mathcal{R}$ be a residue and c a chamber of Δ. The unique chamber in $\mathcal{R}$ at minimal distance to c is called the *projection* of c onto $\mathcal{R}$ and denoted by $\operatorname{proj}_{\mathcal{R}}(c)$.

Similarly, Corollary 1.7.22 has an analogue in the realm of buildings.

Proposition 1.7.39. *Let Δ be a building of type M and let $\mathcal{R}$ be a J-residue. Then $\mathcal{R}$ is convex, and $\mathcal{R}$ is (isomorphic to) a building of type M_J.*

Together with the gate property, this yields a useful characterisation of convexity in buildings.

Corollary 1.7.40. *A set $C \subseteq \Delta$ is convex if and only if for every $c \in C$ and every residue $\mathcal{R}$ with $C \cap \mathcal{R} \neq \emptyset$, we have that $\operatorname{proj}_{\mathcal{R}}(c) \in C$.*

Before giving some examples, we collect some more useful properties of the projection map.

Proposition 1.7.41. *Let Δ be any building.*

(i) *If c and d are adjacent chambers in Δ, and $\mathcal{R}$ is any residue, then $\operatorname{proj}_{\mathcal{R}}(c)$ and $\operatorname{proj}_{\mathcal{R}}(d)$ are adjacent as well (and possibly coincide).*

(ii) *The projection map is distance-decreasing, i.e. for any pair of chambers c, c' and any residue $\mathcal{R}$ we have that $\operatorname{dist}(\operatorname{proj}_{\mathcal{R}}(c), \operatorname{proj}_{\mathcal{R}}(c')) \leq \operatorname{dist}(c, c')$;*

(iii) *Let $\mathcal{S} \subseteq \mathcal{R}$ be two residues. Then for any chamber c, we have $\operatorname{proj}_{\mathcal{S}}(c) = \operatorname{proj}_{\mathcal{S}}(\operatorname{proj}_{\mathcal{R}}(c))$. In particular, if $\operatorname{proj}_{\mathcal{R}}(c) \in \mathcal{S}$, then $\operatorname{proj}_{\mathcal{S}}(c) = \operatorname{proj}_{\mathcal{R}}(c)$.*

Proof. For (i), we may assume without loss of generality that $\operatorname{dist}(c, \operatorname{proj}_{\mathcal{R}}(c)) \leq \operatorname{dist}(d, \operatorname{proj}_{\mathcal{R}}(d))$ so that by the gate property,

$$\begin{aligned} \operatorname{dist}(\operatorname{proj}_{\mathcal{R}}(c), \operatorname{proj}_{\mathcal{R}}(d)) &= \operatorname{dist}(c, \operatorname{proj}_{\mathcal{R}}(d)) - \operatorname{dist}(c, \operatorname{proj}_{\mathcal{R}}(c)) \\ &\leq \operatorname{dist}(c, \operatorname{proj}_{\mathcal{R}}(d)) - \operatorname{dist}(d, \operatorname{proj}_{\mathcal{R}}(d)) \leq 1. \end{aligned}$$

Property (ii) follows immediately from (i), and (iii) from the gate property (with $c' = \operatorname{proj}_{\mathcal{S}}(c)$). ■

Let us now look at a few more interesting examples of buildings.

Example 1.7.42. (i) A building of rank one is nothing more than a complete graph. Indeed, the Coxeter group of rank one is generated by a single involution s and hence isomorphic to $\mathbb{Z}/2\mathbb{Z}$. A building Δ of this type needs to have $\delta(c, d) = s$ for all $c \neq d$, or in other words, the adjacency relation on Δ is the universal relation. Interpreting the chambers as vertices of a graph, Δ is simply a complete graph and vice versa. The apartments are the edges.

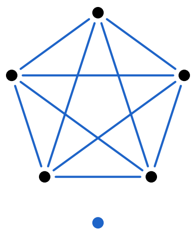

Figure 1.13. A building of rank one.
The diagram is a single node (in blue).

(ii) Recall that a *generalised m-gon* is a connected, bipartite graph of diameter m and girth $2m$, in which every vertex is incident to at least two edges; in this example, we shall sketch why the generalised m-gons are precisely the rank two buildings.

Let Δ be a building of type $\mathsf{I}_2(m)$. From Example 1.7.16 (ii) we know the associated Coxeter complex is a cycle of length $2m$. The reduced words are of the form $iji\cdots$ or $jij\cdots$ of length at most m; they give different group elements except for the words $p(i,j)$ and $p(j,i)$. Define a graph Γ_Δ where the vertices are the i-panels and j-panels of Δ, and join two vertices by an edge if the corresponding panels intersect nontrivially (in a necessarily unique chamber). Then Γ_Δ is connected, comes with a natural bipartition, has diameter m (since every gallery can be reduced to one of length at most m), and girth $2m$ (since the shortest circuits are galleries of type $(ij)^{m_{ij}}$).

Conversely, if Γ is a generalised m-gon, we can construct a chamber system Δ_Γ of rank two over $I = \{\circ, \bullet\}$ as in Example 1.7.9 (iv). Let (W, S) be the Coxeter system of type $\mathsf{I}_2(m)$ with generating set $S = \{s_\circ, s_\bullet\}$, subject to the relations $(s_\circ)^2 = (s_\bullet)^2 = (s_\circ s_\bullet)^m = 1$. Then the only reduced words of length m are $p(\circ, \bullet)$ and $p(\bullet, \circ)$, and any reduced word of length smaller than m is homotopic only to itself. We can then define $\delta\colon \Delta_\Gamma \times \Delta_\Gamma \to W$ by setting $\delta(c, d) = \varsigma(w)$ if $\operatorname{dist}(c, d) < m$ and w is the type of the unique minimal gallery from c to d, and $\delta(c, d) = \varsigma(p(\circ, \bullet)) = \varsigma(p(\bullet, \circ))$ if $\operatorname{dist}(c, d) = m$. The generalised polygon axioms then easily imply that (Δ_Γ, δ) is a building of type $\mathsf{I}_2(m)$.

The apartments of a generalised m-gon are exactly the circuits of length $2m$.

As a concrete illustration, a generalised triangle is nothing more than an (axiomatic) projective plane. Recall from Example 1.7.9 (vi) the interpretation of the Fano plane as a chamber system. By the discussion above, this chamber system is a building of type $\mathsf{I}_2(3) = \mathsf{A}_2$. We have already met the Coxeter complex of this type — a hexagon — in Figure 1.5. And indeed, in Figure 1.4 we can see an abundance of hexagons: every two chambers (edges) are contained in a common circuit of length six.

In a completely similar way, generalised quadrangles are buildings of type $\mathsf{I}_2(4) = \mathsf{B}_2$ with octagons as apartments (as in Figure 1.6). The smallest nontrivial (i.e. thick) generalised quadrangle is commonly known as the *doily*, illustrated in Figure 1.14.

Generalised polygons are well-studied. Despite their apparent simplicity, the axioms enforce rather strong restrictions. For example, if n is finite, a generalised n-gon is automatically semiregular — the tuple $(q_\circ - 1, q_\bullet - 1)$ is called its *order*. We mention a couple of celebrated theorems and refer to [VM98] for proofs.

Theorem 1.7.43 (Feit–Higman). *A finite thick n-gon exist only for $n \in \{2, 3, 4, 6, 8\}$.*

Theorem 1.7.44. *Let (s, t) be the order of a finite thick generalised n-gon.*

- *If $n = 2$, then s and t can be arbitrary.*
- *If $n = 3$, then $s = t$.*
- **(Bose).** *If $n = 4$, then $s \le t^2$ and $t \le s^2$.*
- **(Haemers–Roos).** *If $n = 6$, then st is a square, $s \le t^3$ and $t \le s^3$.*
- **(Feit–Higman, Bose).** *If $n = 8$, then $2st$ is a square, $s \le t^2$ and $t \le s^2$.*

(iii) We explicitly mention that a generalised ∞-gon, or in other words a building of type $\mathsf{I}_2(\infty)$, is nothing more than an infinite tree without leaves (vertices of degree one). The apartments are the two-way infinite paths as in Figure 1.7.

It is worth emphasising the chambers are the *edges* of the tree, and the panels are (stars of) the *vertices*. Recalling our Convention 1.7.2, we point out that a tree looks slightly unfamiliar when drawn as a chamber complex with the chambers as vertices; in fact, we obtain the line graph of the tree. Figure 1.15 compares the two viewpoints.

Note, in graph theory a semiregular tree is usually called *biregular* — which is slightly unfortunate, since such a tree is only “half as regular” as a regular tree, not “twice as regular”!

(iv) Let (W_1, S_1) and (W_2, S_2) be two Coxeter systems with diagrams M_1 and M_2, respectively. Let M be the disjoint union of the two diagrams. This corresponds to a new Coxeter system $(\langle W_1, W_2\rangle, S_1 \sqcup S_2)$ where generators in S_1 and in S_2 pairwise commute. Then a building Δ of type M is (isomorphic to) a direct product of a building Δ_1 of type M_1 and a building Δ_2 of type M_2.

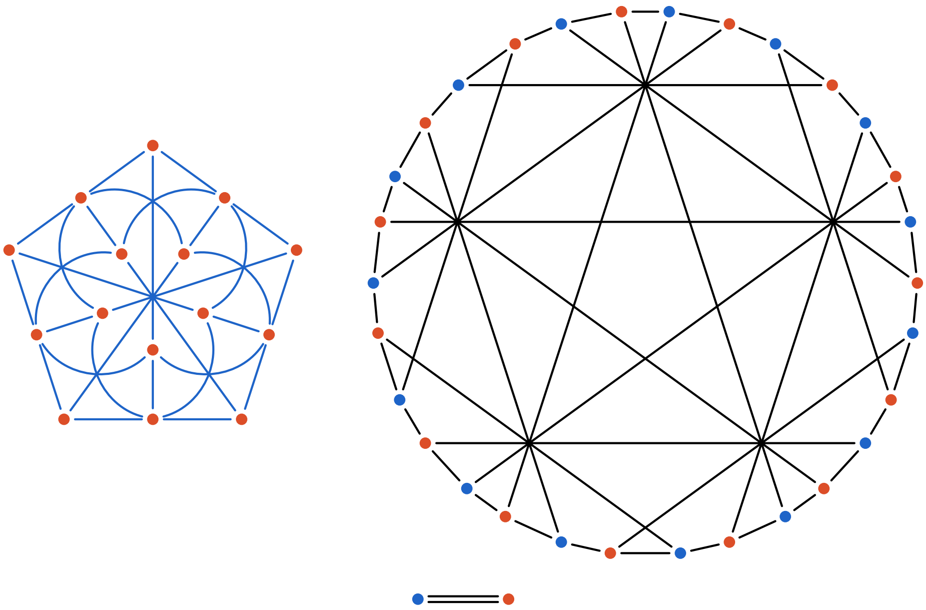

Figure 1.14. The *doily*, the smallest nontrivial generalised quadrangle.
Left: as a point-line incidence structure. Right: as a chamber system.

In fact, we can easily construct an explicit isomorphism. Let $\mathcal{R}_1$ and $\mathcal{R}_2$ be two residues of types I_1 and I_2, respectively, and recall that $\mathcal{R}_1$ and $\mathcal{R}_2$ are buildings in their own right by Proposition 1.7.39. Then the map

$$\varphi\colon \Delta \to \mathcal{R}_1 \times \mathcal{R}_2 \colon d \mapsto (\mathrm{proj}_{\mathcal{R}_1}(d), \mathrm{proj}_{\mathcal{R}_2}(d))$$

is a building isomorphism. For details and proof, we refer to [Ron09, Theorem 3.10].

1.7.4 Equivalent definitions of buildings

In this section, we (briefly) present two alternative but equivalent definitions of buildings — partly out of historical interest, partly because certain properties might be easier to prove when using a different characterisation. We refer to the literature for proofs of the equivalence.

The original definition of Tits presupposed the existence of apartments in a simplicial complex. The formulation that we give is again in terms of chamber systems and is due to Ronan.

Definition 1.7.45. A *building* is a chamber system that can be expressed as the union of certain subsystems $\mathcal{A}$, called *apartments* and satisfying the following axioms:

(i) each apartment $\mathcal{A}$ is (isomorphic to) a Coxeter complex;

(ii) any two chambers c and c' lie in a common apartment $\mathcal{A}$;

(iii) if $\mathcal{A}$ and $\mathcal{A}'$ are apartments containing a common chamber c and chamber c' (or panel $\mathcal{P}$), then there is an isomorphism $\mathcal{A} \to \mathcal{A}'$ fixing c and c' (or $\mathcal{P}$ pointwise).

Proof omitted. For a proof of equivalence, we refer to [Ron09, Theorem 3.11].

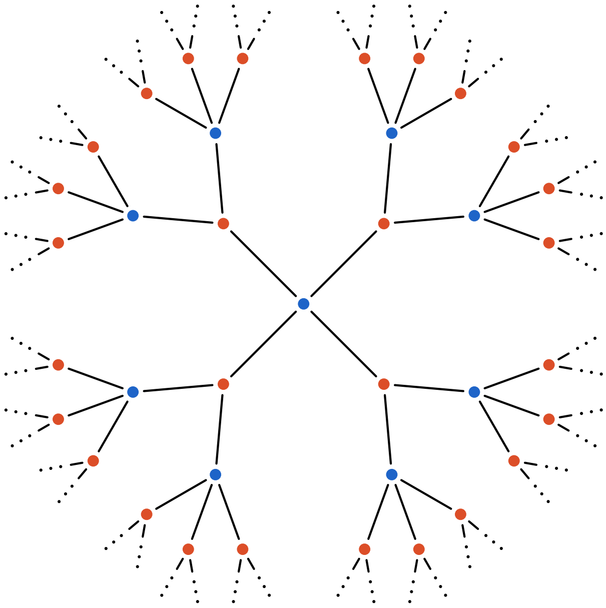

(a) The familiar viewpoint.

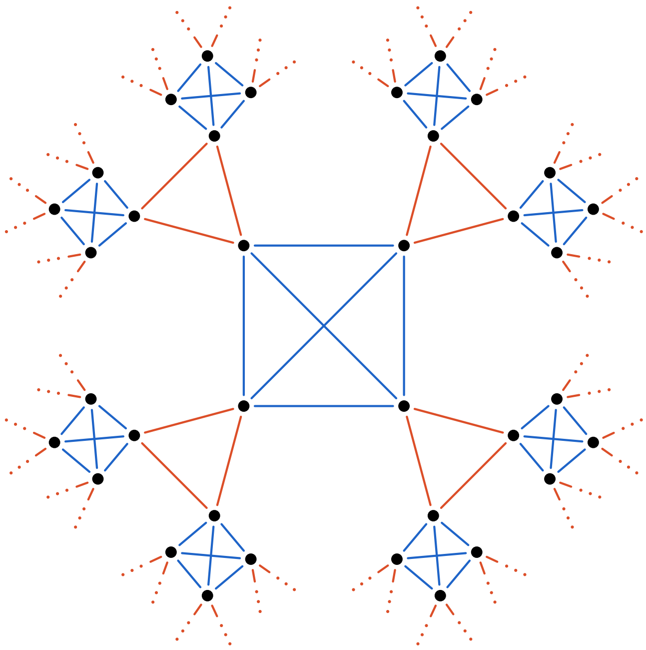

(b) The dual viewpoint.

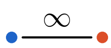

Figure 1.15. An infinite tree when interpreted as a chamber system.
In both viewpoints, the chambers are drawn in black (either as edges or as vertices).

Within this setting, the Weyl distance between two chambers can be defined by looking at a common apartment $\mathcal{A}$, reading off the type of a minimal gallery contained in $\mathcal{A}$, and noting that this is independent of the choice of $\mathcal{A}$ by axiom (iii).

Another equivalent definition was given in [Tit92] and superficially resembles our definition, but defines the Weyl distance in a more technical way from local data.

Definition 1.7.46. Let (W, S) be a Coxeter system of type M over an index set I. A *building of type* M is a chamber system Δ over I together with a map $\delta\colon \Delta \times \Delta \to W$ satisfying the following axioms:

(i) $\delta(c, d) = 1$ if and only if $c = d$;

(ii) if $\delta(c, d) = w$ and c' satisfies $\delta(c', c) = s \in S$, then $\delta(c', d) \in \{sw, w\}$, and if in addition $|sw| = |w| + 1$ in the word length on W, then $\delta(c', d) = sw$;

(iii) if $\delta(c, d) = w$, then for any $s \in S$ there exists a chamber c' such that both $\delta(c', c) = s$ and $\delta(c', d) = sw$.

Proof omitted. For a proof of equivalence, we refer to [AB08, Chapter 5].

Axiom (ii) is vaguely similar to the triangle inequality in a metric space, and axiom (iii) intuitively says that one can always move away from any chamber in any direction. It takes a bit of work to show that δ is "symmetrical" (in the sense that $\delta(c, d)^{-1} = \delta(d, c)$), but the set of axioms is more robust; in fact, four initial exercises in [AB08, Chapter 5] ask to show that one can replace one or more axioms by certain other properties and still retain an equivalent definition.

1.7.5 Buildings as $\mathrm{CAT}(0)$ spaces

In this section, we describe the classical result of Michael Davis ([Dav98]) that every building can be realised as a complete $\mathrm{CAT}(0)$ metric space (or a *Hadamard space*).

In order to construct a geometrical realisation of an arbitrary building, we can essentially choose a topological space Z that will fulfill the role of a chamber, choose a subspace Z_i for every $i \in I$, and glue together copies of Z along those subspaces. The following definition makes this formal.

Definition 1.7.47 (mirror space, Z-realisation). Let Δ be a building of type M over I.

(i) A *mirror space* over I is a topological space Z together with a family $\{Z_i \mid i \in I\}$ of nonempty closed subspaces called *mirrors*. For every point $z \in Z$, define $I(z) = \{i \in I \mid z \in Z_i\}$.

(ii) Define an equivalence relation on $\Delta \times Z$ by declaring $(c, z) \sim (c', z')$ if and only if $z = z'$ and $\delta(c, c') \in W_{I(z)}$ (i.e. the subgroup generated by all s_i such that $z \in Z_i$). Equip Δ with the discrete topology, $\Delta \times Z$ with the product topology, and the quotient

$$Z(\Delta) = (\Delta \times Z)/{\sim}$$

with the quotient topology. The space $Z(\Delta)$ is called the *Z-realisation* of Δ.

It should not come as a surprise that topological properties of the mirror space Z are reflected in the Z-realisation; we mention one illustration.

Proposition 1.7.48. *If Z is (path-) connected, then $Z(\Delta)$ is (path-) connected.*

Proof. For Coxeter complexes, we refer to [Dav12, Corollary 5.2.4]. For general buildings, it suffices to notice that every two points of $Z(\Delta)$ are contained in a common homeomorphic image of the Z-realisation of a Coxeter complex (by Corollary 1.7.33). ■

Of more interest is the fact that group actions lift to Z-realisations as well. Let G be a group acting on Δ by automorphisms. Then there is a natural action on $\Delta \times Z$, defined by $g\,.\,(c, z) = (g\,.\,c, z)$. This preserves the equivalence $\sim$ and hence induces an well-defined action on $Z(\Delta)$, by

$$g\,.\,[(c, z)] = [(g\,.\,c, z)].$$

Proposition 1.7.49. *Let Z be a mirror space. Let G act on Δ by automorphisms and consider the induced action on $Z(\Delta)$.*

(i) *The stabiliser of the point $[(c, z)]$ equals the stabiliser of the residue of type $I(z)$ containing c.*

(ii) *Assume that Z is not the union of its mirrors. If the action on Δ is faithful, then so is the action on $Z(\Delta)$.*

Proof. For (i), we can straightforwardly calculate

$$\begin{aligned} G_{[(c,z)]} &= \{g \in G \mid g\,.\,(c, z) \sim (c, z)\} \\ &= \{g \in G \mid \delta(g\,.\,c, c) \in \langle s_i \mid z \in Z_i\rangle\} \\ &= \{g \in G \mid g\,.\,c \text{ and } c \text{ are contained in the same } I(z)\text{-residue}\}. \end{aligned}$$

For (ii), let $z_0 \in Z$ be a point not contained in any mirror. Then for every $c \in \Delta$, the stabiliser of $[(c, z_0)]$ coincides with the stabiliser of c, and the result follows. ■

Example 1.7.50. (i) As a trivial degenerate example, when Z is a single point, the Z-realisation of a building collapses into a single point as well.

(ii) The tessellation in Figure 1.10 is the Z-realisation obtained by taking a square as the mirror space with its sides as mirrors (opposites sides corresponding to noncommuting generators). Similarly, from a pentagon as the mirror space one can obtain the tessellation in Figure 1.12.

(iii) Let the mirror space Z be a simplex on n vertices (where n is the rank of Δ), with its n faces of codimension one as mirrors. This way, one recovers exactly the simplicial viewpoint on buildings; we call this the *simplicial realisation* of a building. A simplex comes with a natural Euclidean metric and hence so does the simplicial realisation. Unfortunately, though $Z(\Delta)$ is $\mathrm{CAT}(0)$ for an affine building, $Z(\Delta)$ is a triangulation of a sphere for a spherical building and hence the simplicial realisation is not $\mathrm{CAT}(0)$ in general.

Example 1.7.50 (i) shows that every building Δ can be realised as a $\mathrm{CAT}(0)$ space in a trivial way. We now construct a mirror space Z such that the resulting Z-realisation will be $\mathrm{CAT}(0)$ without destroying the building's structure, i.e. there will be a canonical injection $\mathrm{Aut}(\Delta) \hookrightarrow \mathrm{Iso}(Z(\Delta))$.

The construction is based on Example 1.7.50 (iii), but builds a more refined simplicial complex Z — intuitively, the main obstruction for the $\mathrm{CAT}(0)$ property to hold in the simplicial realisation are the triangulated spheres, so the goal is to add simplices that "fill up" these spheres into triangulated solid balls. The spherical data of a Coxeter system is captured in the next few definitions.

Definition 1.7.51 (nerve, Davis realisation). Let Δ be a building over I and let (W, S) be the associated Coxeter system.

(i) A subset $J \subseteq I$ is *spherical* if the subsystem (W_J, S_J) is spherical (as in Proposition 1.7.21), i.e. if the generators s_j with $j \in J$ generate a finite subgroup. Clearly, a subset of a spherical subset is again spherical, and all singletons are spherical.

(ii) Define the poset $\mathbb{L}(W, S)$ of nonempty spherical subsets of I, partially ordered by inclusion. This is an abstract simplicial complex where the simplices are the spherical subsets. $\mathbb{L}(W, S)$ is called the *nerve* of (W, S). The nerve can be realised geometrically as a simplicial complex with vertex set S, such that a set of generators spans a simplex if and only if they generate a finite subgroup.

(iii) Define $\mathbb{L}'(W, S)$ to be the barycentric subdivision of (the geometric realisation of) the nerve. This is essentially a new simplicial complex with a vertex for every simplex in $\mathbb{L}(W, S)$.

(iv) Let $\mathbb{K}$ be the cone over $\mathbb{L}'(W, S)$. It follows that vertices of $\mathbb{K}$ are in natural correspondence to spherical subsets of I; the cone point corresponds to the empty subset. Note that $\mathbb{K}$ has a natural piecewise Euclidean metric. For every $i \in I$, let $\mathbb{K}_i$ be the closed star in $\mathbb{L}'(W, S)$ of the vertex s_i (i.e. the union of all simplices in $\mathbb{L}'(W, S)$ containing s_i).

(v) The Z-realisation obtained from the mirror space $\mathbb{K}$ with the mirrors $\mathbb{K}_i$ is called the *Davis realisation* $\mathbb{K}(\Delta)$ of Δ.

We give an example of the Davis realisation of a thin right-angled building in Figure 1.16. Note that the original Coxeter complex is illustrated in Figure 1.11.

Theorem 1.7.52. *The Davis realisation* $\mathbb{K}(\Delta)$ *of any building* Δ *is a complete* $\mathrm{CAT}(0)$ *space.*

Proof omitted. We refer to [Dav98, Theorem 11.1] or [Dav12, Theorem 18.3.1].

As a corollary, the Bruhat–Tits fixed point theorem (Theorem 1.6.5) is applicable.

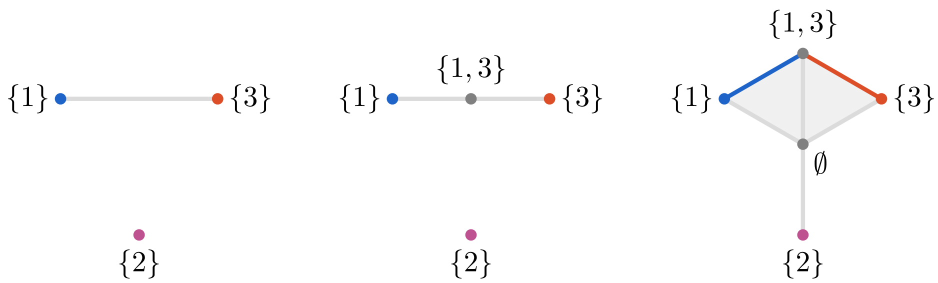

(a) From left to right: the nerve $\mathbb{L}$, its barycentric subdivision $\mathbb{L}'$, and the mirror space $\mathbb{K}$.

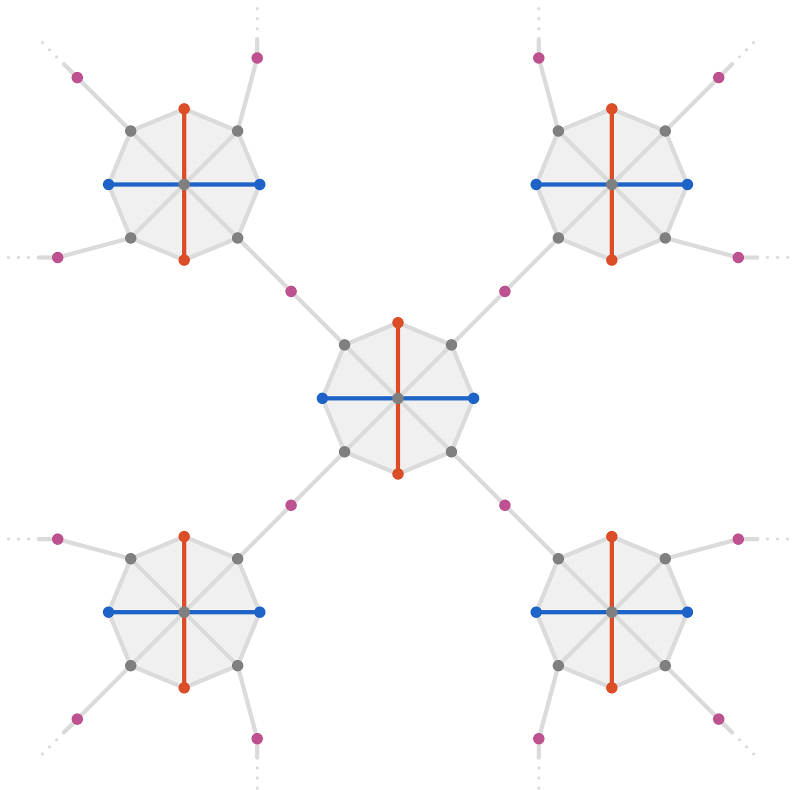

(b) The quotient space of glued-together copies of $\mathbb{K}$.

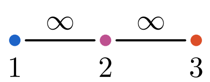

Figure 1.16. A thin right-angled example of a Davis realisation.

2 Right-angled buildings

> *Oh-ho-ho! Enlightenment! When it comes, it comes like a brick to the head, doesn't it?*
>
> — Scott Lynch, *The Lies of Locke Lamora*

In this introductory section we establish some general facts about right-angled buildings and their automorphisms. We will only work with *semiregular* buildings. The main motivation for restricting our attention should be clear: automorphisms preserve cardinalities of panels, and semiregularity is hence a necessary condition for the automorphism group to act transitively on the chambers.

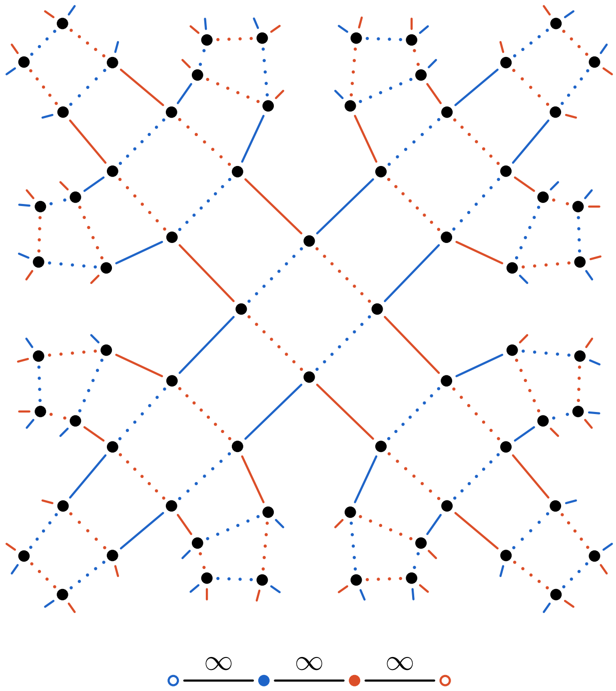

Figure 2.1. A right-angled Coxeter complex of rank four.
Nodes ● and ● correspond to solid edges, nodes ○ and ○ to dotted edges.

2.1 Global results

Recall that a rank two residue of an arbitrary building is essentially a generalised n-gon. Recall also that Theorems 1.7.43 and 1.7.44 enforce quite strong restrictions on the parameters of this residue — unless $n \in \{2, \infty\}$. For $n = 2$, a rank two residue is a direct product of two arbitrary complete graphs; for $n = \infty$, a rank two residue is (the line graph of) an infinite tree with arbitrary valencies. The following theorem is due to Frédéric Haglund and Frédéric Paulin ([HP03]), and shows that this local freedom in the rank two residues of a right-angled building extends to the global building.

Theorem 2.1.1 (Haglund–Paulin). *For any choice of (possibly infinite) cardinal numbers* $(q_i)_{i\in I}$ *with* $q_i \geq 2$*, there exists a semiregular right-angled building* Δ *with these parameters. Moreover,* Δ *is unique up to isomorphism, the automorphism group* $\mathrm{Aut}(\Delta)$ *acts transitively on the chambers, and every automorphism of a residue of* Δ *extends to an automorphism of* Δ*.*

Proof omitted. This is [HP03, Proposition 1.2].

Haglund and Paulin note that Theorem 2.1.1 was already known to Mark Globus (but unpublished), Michael Davis and Gabor Moussong, and Tadeusz Januszkiewicz and Jacek Świątkowski. For right-angled buildings embedded in the hyperbolic plane (for instance, Figure 1.12), the first results were due to Marc Bourdon.

Another nice property of right-angled buildings is that the underlying Coxeter systems are “rigid”.

Theorem 2.1.2. *If a right-angled Coxeter group* W *admits two Coxeter systems* (W, S) *and* (W, S')*, then these Coxeter systems are isomorphic (i.e. there is a diagram-preserving bijection* $S \to S'$*).*

Proof omitted. We refer to [Rad02] or [Hos03].

Theorem 2.1.2 is important in the following sense. Recall that to any Coxeter system (W, S) — in fact, to any building — we have associated a $\mathrm{CAT}(0)$ space in Definition 1.7.51, defined by the data in S. As a corollary of Theorem 2.1.2, for right-angled Coxeter systems, this geometry is intrinsic to the Coxeter *group*. We remark that for Coxeter systems in general the “large scale geometry” of this $\mathrm{CAT}(0)$ realisation is conjectured to be determined by the group alone; in particular Dranishnikov’s *rigidity conjecture* claims that isomorphic Coxeter groups have homeomorphic boundaries ([Dra01]).

We will not immediately need the associated $\mathrm{CAT}(0)$ geometries, as we can build up the theory in a more combinatorial fashion using the Coxeter system and galleries in the chamber system.

2.2 Minimal galleries

A couple of important results allow us to modify minimal galleries in right-angled buildings in a controlled way by “closing squares” — see Figure 2.2. We note that Lemmas 2.2.1 and 2.2.2, Proposition 2.2.3, and Corollary 2.2.4 first appeared in [DMdSS18].

Lemma 2.2.1 (closing squares (1)). *Let* c_0 *be a fixed chamber in a right-angled building* Δ*. Let* $c_1, c_2, c_3 \in \Delta$ *be such that* $\mathrm{dist}(c_0, c_1) = \mathrm{dist}(c_0, c_3) = n$*,* $\mathrm{dist}(c_0, c_2) = n + 1$*, and* $c_1 \sim_i c_2 \sim_j c_3$ *for* $i \neq j$*. Then* $m_{ij} = 2$ *and there exists* $d \in \Delta$ *such that* $\mathrm{dist}(c_0, d) = n - 1$ *and* $c_1 \sim_j d \sim_i c_3$*.*

Proof. Let w_1 and w_3 be the types of a minimal gallery joining c_0 to c_1 and to c_3 respectively. Then the words $w_1\, i$ and $w_3\, j$ of length $n + 1$ are both reduced representations of $\delta(c_0, c_2)$, hence they are homotopic. It follows that $m_{ij} = 2$ and that w_1 is homotopic to a word of the form wj with w reduced of length $n - 1$. Let d be the chamber such that $\delta(c_0, d) = \varsigma(w)$ and $\delta(d, c_1) = \varsigma(j)$. Then

$$\varsigma(w\, i) \cdot \varsigma(j) = \varsigma(w\, ij) = \varsigma(w\, ji) = \varsigma(w_1\, i) = \varsigma(w_3\, j) = \varsigma(w_3) \cdot \varsigma(j)$$

and since all words involved are reduced, we finally obtain that the gallery of type w from c_0 to d can be extended to c_3 by a single i-adjacency. In other words, $d \sim_i c_3$. ∎

Lemma 2.2.2 (closing squares (2)). *Let* c_0 *be a fixed chamber in a right-angled building* Δ*. Let* $c_1, c_2, c_3 \in \Delta$ *be such that* $\mathrm{dist}(c_0, c_1) = \mathrm{dist}(c_0, c_2) = n + 1$*,* $\mathrm{dist}(c_0, c_3) = n$*, and* $c_1 \sim_i c_2 \sim_j c_3$ *for* $i \neq j$*. Then* $m_{ij} = 2$ *and there exists* $d \in \Delta$ *such that* $\mathrm{dist}(c_0, d) = n$ *and* $c_1 \sim_j d \sim_i c_3$*.*

Proof. Let w_1 and w_2 be the types of a minimal gallery joining c_0 to c_1 and to c_2 respectively. Then since c_1 and c_2 are i-adjacent and lie at equal distance to c_0, types w_1 and w_2 are both homotopic to a word of the form $w\,i$ with w reduced of length n. Let c' be the chamber such that $\delta(c_0, c') = \varsigma(w)$ and $\delta(c', c_1) = \delta(c', c_2) = \varsigma(i)$. Then by Lemma 2.2.1, applied to the chambers $\{c', c_2, c_3\}$, it follows that $m_{ij} = 2$ and that there is a chamber d' such that $\mathrm{dist}(c_0, d') = n - 1$ and $c' \sim_j d' \sim_i c_3$. Now note that there is a gallery $d' \sim_j c' \sim_i c_1$ of type ji; hence there is also a gallery $d' \sim_i d \sim_j c_1$ of type ij for some chamber d. Then $\mathrm{dist}(c_0, d) = n$ and $d \sim_i d' \sim_i c_3$. ■

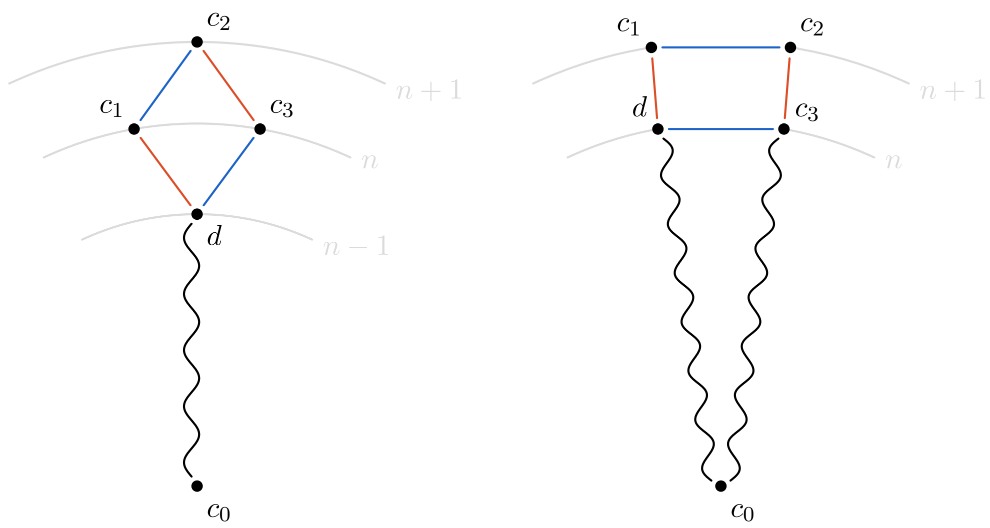

(a) The configuration from Lemma 2.2.1. (b) The configuration from Lemma 2.2.2.

Figure 2.2. The “closing squares” lemmas.

As a corollary, given any minimal gallery γ and any reference point c_0, we can always transform γ into a “concave” gallery with the same extremities.

Proposition 2.2.3. *Let c_0, c, c' be three chambers in a right-angled building. Then there exists a minimal gallery $\gamma = (d_0, \dots, d_\ell)$ from $d_0 = c$ to $d_\ell = c'$ of length ℓ, and indices $0 \le m \le n \le \ell$, satisfying the following:*

- $\mathrm{dist}(c_0, d_{k-1}) > \mathrm{dist}(c_0, k)$ *for all* $k \in \{1, \dots, m\}$;
- $\mathrm{dist}(c_0, d_{k-1}) = \mathrm{dist}(c_0, k)$ *for all* $k \in \{m+1, \dots, n\}$;
- $\mathrm{dist}(c_0, d_{k-1}) < \mathrm{dist}(c_0, k)$ *for all* $k \in \{n+1, \dots, \ell\}$.

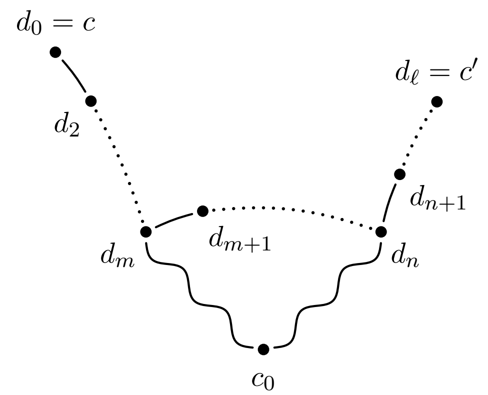

Proof. Let $\gamma = (d_0, \dots, d_\ell)$ be any minimal gallery from $d_0 = c$ to $d_\ell = c'$. We will prove the result by closing squares in γ whenever possible. Note that the required concavity can be rephrased more colloquially as follows: *once γ stops going downwards, it can never go down again, and once γ starts going upwards, it must continue going up.* Observe then that this translates into the local condition that γ has no length two subgalleries of the following form (see also Figure 2.3):

(a) (d_k, d_{k+1}, d_{k+2}) with $\mathrm{dist}(c_0, d_k) = \mathrm{dist}(c_0, d_{k+1}) > \mathrm{dist}(c_0, d_{k+1})$;

(b) (d_k, d_{k+1}, d_{k+2}) with $\mathrm{dist}(c_0, d_k) < \mathrm{dist}(c_0, d_{k+1})$ and $\mathrm{dist}(c_0, d_{k+1}) > \mathrm{dist}(c_0, d_{k+2})$;

(c) (d_k, d_{k+1}, d_{k+2}) with $\mathrm{dist}(c_0, d_k) < \mathrm{dist}(c_0, d_{k+1}) = \mathrm{dist}(c_0, d_{k+2})$.

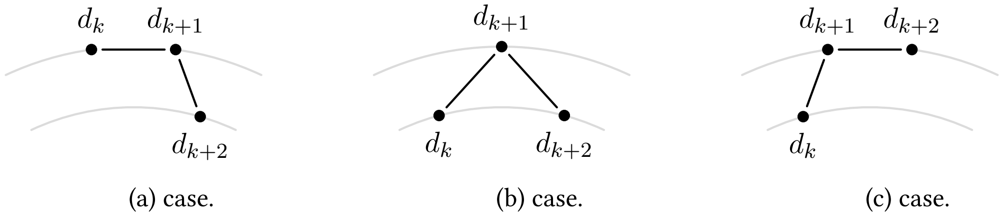

Figure 2.3. The forbidden cases in a “concave” gallery.

Define the *total height* $h(\gamma)$ of the gallery γ as the sum of distances

$$h(\gamma) = \sum_{k=0}^{\ell} \mathrm{dist}(c_0, d_k).$$

If γ has some length two subgallery of the form (a) or (c), then apply Lemma 2.2.2 to replace d_{k+1} by another chamber in order to decrease the total height by one. Similarly, if γ has a subgallery of the form (b), then apply Lemma 2.2.1 in order to decrease the total height by two. In any case, we obtain a new minimal gallery γ' with the same extremities but of strictly smaller total height.

Since the total height $h(\gamma)$ is a natural number, we cannot repeat this process indefinitely, and at some point we end up with a minimal gallery with the required properties. ■

In a right-angled building, balls are not convex in general, but Proposition 2.2.3 immediately yields a second best possibility.

Corollary 2.2.4. *Let c_0 be a fixed chamber in a right-angled building and let $c, c' \in \mathsf{B}_n(c_0)$. Then there exists a minimal gallery that joins c to c' and is contained in $\mathsf{B}_n(c_0)$.*

Proof. A “concave” gallery as in Proposition 2.2.3 (with reference chamber c_0) suffices. ■

2.3 Projections, parallelism, wings

Recall our Definition 1.7.38 of combinatorial projections: given a chamber c and a residue $\mathcal{R}$ of Δ, the projection $\mathrm{proj}_{\mathcal{R}}(c)$ is the unique chamber in $\mathcal{R}$ at minimal distance to c. We first give a rather technical lemma that relates projections onto panels to the Coxeter diagram, and that will prove useful in step-by-step constructions of automorphisms.

Lemma 2.3.1. *Let c be a chamber and $\mathcal{P}$ an i-panel in a right-angled building Δ. Let $c' = \mathrm{proj}_{\mathcal{P}}(c)$ and $\mathrm{dist}(c, c') = n$. Let $d \in \mathsf{B}_{n+1}(c) \setminus \mathcal{P}$, let $d' = \mathrm{proj}_{\mathcal{P}}(d)$, and assume that $d' \neq c'$. Then $\mathrm{dist}(c, d') = n + 1$ and d' is j-adjacent to some chamber in $\mathsf{S}_n(c)$ with $m_{ij} = 2$.*

Proof. Since c' is the unique chamber in $\mathcal{P}$ at minimal distance to c, the first claim follows immediately from the assumption that $d' \neq c'$. By Corollary 2.2.4, there exists a concave minimal gallery $\gamma = (d \sim \cdots e \sim d')$ with respect to c. Let $w = i_1 \cdots i_\ell$ be the type of γ. Then $i_\ell \neq i$ because $\mathrm{proj}_{\mathcal{P}}(d) = d' \neq e$. There are two possibilities: either $\mathrm{dist}(c, e) = n$ and the result follows from Lemma 2.2.1, or $\mathrm{dist}(c, e) = n + 1$ and the result follows from Lemma 2.2.2. ■

We can define projections between residues in a natural way: for residues $\mathcal{R}$ and $\mathcal{S}$, we set

$$\mathrm{proj}_{\mathcal{S}}(\mathcal{R}) = \{\mathrm{proj}_{\mathcal{S}}(c) \mid c \in \mathcal{R}\}.$$

Note that this set is again a residue of Δ contained in $\mathcal{S}$, and that its rank is bounded by the ranks of both $\mathcal{S}$ and $\mathcal{R}$.

Definition 2.3.2 (parallelism). Two residues $\mathcal{R}_1$ and $\mathcal{R}_2$ are called *parallel* if $\mathrm{proj}_{\mathcal{R}_1}(\mathcal{R}_2) = \mathcal{R}_1$ and $\mathrm{proj}_{\mathcal{R}_2}(\mathcal{R}_1) = \mathcal{R}_2$. In that case, both projection maps define bijections between $\mathcal{R}_1$ and $\mathcal{R}_2$.

Since a projection map between residues does not increase the rank, it follows immediately that parallel residues have the same rank.

Lemma 2.3.3. *Let $\mathcal{R}_1$ and $\mathcal{R}_2$ be parallel residues and let $c_1 \in \mathcal{R}_1$ and $c_2 \in \mathcal{R}_2$. Then*

$$\mathrm{dist}(c_1, \mathcal{R}_2) = \mathrm{dist}(c_2, \mathcal{R}_1) = \mathrm{dist}(\mathcal{R}_1, \mathcal{R}_2).$$

Proof omitted. This is [Cap14a, Lemma 2.4].

There is a useful criterion for detecting parallelism among panels, that can then be used to detect parallelism among higher rank residues as well.

Lemma 2.3.4.
(i) *Let $\mathcal{P}_1$ and $\mathcal{P}_2$ be panels. Then $\mathcal{P}_1$ and $\mathcal{P}_2$ are parallel if and only if there exist two chambers in $\mathcal{P}_1$ with distinct projections on $\mathcal{P}_2$.*

(ii) *Let $\mathcal{R}_1$ and $\mathcal{R}_2$ be residues. Then $\mathcal{R}_1$ and $\mathcal{R}_2$ are parallel if and only if the projection of every panel in $\mathcal{R}_1$ to $\mathcal{R}_2$ is a panel and vice versa.*

(iii) *Let $\mathcal{R}_1$ and $\mathcal{R}_2$ be residues. Then $\mathrm{proj}_{\mathcal{R}_1}(\mathcal{R}_2)$ and $\mathrm{proj}_{\mathcal{R}_2}(\mathcal{R}_1)$ are parallel.*

Proof omitted. This is [Cap14a, Lemma 2.5, Lemma 2.6, and Lemma 2.7].

Example 2.3.5.
(i) Recall the Coxeter complex of type A_2 from Figure 1.5, which is a hexagon. There two panels are parallel if and only if they correspond to opposite sides of the hexagon. In a full-scale building of type A_2 — a projective plane — pairs of parallel panels correspond to nonincident point-line pairs. Note in particular that this notion of parallelism is not transitive in the projective plane.

There is a more general notion of opposition in spherical buildings (two residues are *opposite* if they lie at maximal distance, in a certain precise way) and any two such opposite residues are in fact parallel.

(ii) Let $J_1, J_2 \subset I$ be two disjoint nonempty subsets with $m_{j_1 j_2} = 2$ for all (j_1, j_2) in $J_1 \times J_2$. Then a residue $\mathcal{R}$ of type $J_1 \cup J_2$ is the direct product of two residues $\mathcal{R}_1$ and $\mathcal{R}_2$ of types J_1 and J_2 (recall Example 1.7.42 (iv)). Moreover, the canonical projection maps

$$\pi_1 \colon \mathcal{R}_1 \times \mathcal{R}_2 \to \mathcal{R}_1 \qquad \text{and} \qquad \pi_2 \colon \mathcal{R}_1 \times \mathcal{R}_2 \to \mathcal{R}_2$$

coincide with the restrictions of the geometrical projection maps

$$\mathrm{proj}_{\mathcal{R}_1}\big|_{\mathcal{R}} \colon \mathcal{R} \to \mathcal{R}_1 \qquad \text{and} \qquad \mathrm{proj}_{\mathcal{R}_2}\big|_{\mathcal{R}} \colon \mathcal{R} \to \mathcal{R}_2.$$

In particular, any two residues of type J_1 contained in $\mathcal{R}$ are parallel, as are any two residues of type J_2 contained in $\mathcal{R}$.

Example 2.3.5 (i) shows that parallel panels do not necessarily have the same type. In right-angled buildings, however, Example 2.3.5 (ii) is essentially the only possible occurence of parallelism, as the following proposition from [Cap14a] shows. In particular, parallelism in the right-angled case will be a type-preserving equivalence relation. First we need a definition.

Definition 2.3.6. Let $J \subseteq I$. Then we define the set

$$J^\perp = \{i \in I \mid m_{ij} = 2 \text{ for all } j \in J\} = \{i \in I \setminus J \mid ij = ji \text{ for all } j \in J\}.$$

When $J = \{j\}$ is a singleton, we abbreviate

$$j^\perp = \{i \in I \mid m_{ij} = 2\} = \{i \in I \setminus J \mid ij = ji\}.$$

Proposition 2.3.7. *Let Δ be a right-angled building over I.*

(i) *Two parallel residues have equal type.*

(ii) *Two residues of equal type J are parallel if and only if they are contained in a common residue of type $J \cup J^\perp$.*

(iii) *Parallelism of residues is an equivalence relation.*

Proof omitted. We refer to [Cap14a, Proposition 2.8 and Corollary 2.9].

It is worth mentioning that [Cap14a, Proposition 2.10] characterises right-angled thick buildings in terms of parallelism: a thick building turns out to be right-angled if and only if panel parallelism is an equivalence relation.

Definition 2.3.8 (tree-wall). Let Δ be a right-angled building and let $i \in I$. An equivalence class of parallel i-panels is called an *i-tree-wall* of Δ. By slight abuse of notation we often identify an i-tree-wall $\mathcal{T}$ with the set of chambers $\{c \in \mathcal{P} \mid \mathcal{P} \in \mathcal{T}\}$.

The name "tree-wall" might be quite mysterious, since we did not define *walls* in buildings. Let us settle here for a brief sketch (and the remark that walls are easier to define in the simplicial setting). In the Coxeter complex, one defines *reflections* as nontrivial Coxeter group elements that stabilise some panel (edge), and after a moment of thought, a reflection is nothing more than a conjugate of a generator of the Coxeter group. A maximal set of panels invariant under some reflection is called a *wall* of the complex. Then, for a general building, a *wall* is by definition nothing more than a wall in some apartment.

One then observes that the intersection of a tree-wall with an apartment in a right-angled building is either empty or a wall in that apartment. In the literature, tree-walls have also been called *wall-trees* by ([Bou97]) or *wall-residues* (by [Cap14a]).

As a corollary of Proposition 2.3.7, the i-tree-walls are in one-to-one correspondence to (the sets of i-panels contained in) the residues of type $i \cup i^\perp$. In particular, it makes sense to define projections on tree-walls.

Definition 2.3.9 (wings). Let $J \subseteq I$ and let $c \in \Delta$ be a chamber. Then the set of chambers

$$X_J(c) = \{d \in \Delta \mid \operatorname{proj}_{\mathcal{R}}(d) = c\}$$

where $\mathcal{R}$ is the J-residue containing c, is called the *J-wing of c*. Again, if $J = \{j\}$ is a singleton, we usually simply write $X_j(c)$ and call it the *j-wing of c*.

Lemma 2.3.10. *Let c be a chamber in a right-angled building Δ and let $J \subseteq I$. Then*

(i) *$X_J(c)$ is convex;*

(ii) *$X_J(c)$ is equal to the intersection of all wings $X_j(c)$ with $j \in J$.*

Let $\mathcal{R}$ be the residue of type $J \cup J^\perp$ containing c. Then moreover

(iii) *the intersection $X_J(c) \cap \mathcal{R}$ is the residue of type $J^\perp$ containing c;*

(iv) *for all d in the intersection $X_J(c) \cap \mathcal{R}$, we have $X_J(c) = X_J(d)$;*

Proof omitted. This is [Cap14a, Lemma 3.1 and Proposition 3.2].

Lemma 2.3.10 implies that every tree-wall induces a well-defined partition of Δ. Indeed, let $\mathcal{P}$ be any i-panel in an i-tree-wall $\mathcal{T}$; then the i-wings $X_i(d)$ with $d \in \mathcal{P}$ partition the chambers of Δ into $|\mathcal{P}|$ subsets, and this partition is independent of the choice of $\mathcal{P}$ in the tree-wall.

We collect some more lemmas from [Cap14a], without proof.

Lemma 2.3.11. *Let $c \in \Delta$ and let $\mathcal{R}$ be the residue of type $i \cup i^\perp$ containing c. Let d and d' be two chambers such that $d \in X_i(c)$ but $d' \notin X_i(c)$. Then the concatenation of a minimal gallery from d to $\operatorname{proj}_{\mathcal{R}}(d)$, a minimal gallery from $\operatorname{proj}_{\mathcal{R}}(d)$ to $\operatorname{proj}_{\mathcal{R}}(d')$, and a minimal gallery from $\operatorname{proj}_{\mathcal{R}}(d')$ to d', is a minimal gallery from d to d'.*

Proof omitted. This is [Cap14a, Lemma 3.3].

Lemma 2.3.12. *Let $i, j \in I$ and $c, d \in \Delta$. Assume $c \in X_j(d)$ but $d \notin X_i(c)$, and moreover either $i = j$ or $m_{ij} = \infty$. Then $X_i(c) \subseteq X_j(d)$.*

Proof omitted. This is [Cap14a, Lemma 3.4 (a)].

Lemma 2.3.13. *Let $J \subseteq I$, let $\mathcal{R}$ be a residue of type J, and let c be a chamber in $\mathcal{R}$. Then for every $i \in I \setminus J$ we have that $\mathcal{R} \subseteq X_i(c)$.*

Proof omitted. This is [Cap14a, Corollary 3.7].

We now use the data in the i-tree-walls to associate a tree to Δ (for every $i \in I$). It will turn out that the collection of these trees captures enough information about group actions on the building to be able to go back and forth between buildings and trees and apply the theory of Section 1.4.

Definition 2.3.14 (tree-wall tree). Let V_1 be the set of all i-tree-walls of Δ, and let V_2 be the set of all residues of Δ of type $I \setminus \{i\}$. Define a bipartite graph Γ_i with vertex set $V_1 \sqcup V_2$ where an i-tree-wall $\mathcal{T}$ is adjacent to a residue $\mathcal{R}$ of type $I \setminus \{i\}$ if and only if $\mathcal{T} \cap \mathcal{R} \neq \emptyset$. We call this graph the *i-tree-wall tree* of Δ.

Note that if the intersection $\mathcal{T} \cap \mathcal{R}$ is nonempty, it equals a residue of type $i^\perp$ of Δ. Hence there is a one-to-one correspondence between such residues and edges of Γ_i.

Proposition 2.3.15. *For every $i \in I$, the i-tree-wall tree Γ_i of Δ is a semiregular tree.*

Proof. Clearly Γ_i is connected, hence in order to show that it is a tree, it suffices to show that there are no nontrivial cycles. Suppose that

$$\mathcal{T}_1 \sim \mathcal{R}_1 \sim \mathcal{T}_2 \sim \cdots \sim \mathcal{R}_n \sim \mathcal{T}_1$$

is a simple cycle in Γ_i ($n \geq 2$). For every $\ell \in \{1, \dots, n\}$, let c_ℓ be a chamber in $\mathcal{T}_\ell \cap \mathcal{R}_\ell$ and d_ℓ a chamber in $\mathcal{R}_{\ell-1} \cap \mathcal{T}_\ell$ (where the indices are considered modulo n). Then $c_\ell \in X_i(d_{\ell+1})$, which is disjoint from the wing $X_i(c_{\ell+1})$. However, since c_ℓ and $c_{\ell+1}$ do not lie in parallel panels, we do have that $c_{\ell+1} \in X_i(c_\ell)$. Consequently, by Lemma 2.3.12, it follows that

$$X_i(c_1) \supseteq X_i(c_2) \supseteq \cdots \supseteq X_i(c_n) \supseteq X_i(c_1).$$

Then all chambers c_ℓ lie in a common tree-wall — a contradiction.

The fact that Γ_i is moreover semiregular is a direct consequence of Δ being semiregular. Explicitly, the parameters of Γ_i are equal to the number of residues of type $i^\perp$ in an i-tree-wall (which is q_i) and the number of residues of type $i^\perp$ in a common residue of type $I \setminus \{i\}$. ∎

It is worth explicitly mentioning the shape of the i-tree-wall tree in the case where i is an isolated node of the diagram. Then the building has only a single i-tree-wall (namely the full building), and Γ_i is the star graph with q_i leaves. If i is not an isolated node, then the i-tree-wall tree is unbounded.

We finish this section with a couple of generalisations of earlier results. First, a definition.

Definition 2.3.16 (panel-closed). A set C of chambers is called *panel-closed* if for every panel $\mathcal{P}$ the intersection $\mathcal{P} \cap C$ is either empty, a single chamber, or the full panel $\mathcal{P}$.

The following lemma is a generalisation of the gate property for residues (Corollary 1.7.37). Note that residues are indeed convex panel-closed subsets of a right-angled building Δ.

Lemma 2.3.17. *Let C be a convex, panel-closed subset of Δ. Let $c \in \Delta$ be any chamber. Then there is a unique chamber $d \in C$ such that $\operatorname{dist}(c,d)$ is minimal. Moreover, the gate property holds: for every chamber $e \in C$, there exists a minimal gallery from c via d to e.*

Proof. Assume that d and d' are two distinct chambers in C that minimise the distance to c. Let γ be the minimal gallery that joins d and d' and is concave with respect to c, as in Proposition 2.2.3. By convexity of C and concavity of γ, every chamber on γ is a chamber in C at the same distance to c as d and d'. Henceforth we may assume that d and d' are i-adjacent for some $i \in I$. Let $\mathcal{P}$ be the i-panel containing d and d' and let $c' = \operatorname{proj}_{\mathcal{P}}(c)$. Then $\operatorname{dist}(c,c') < \operatorname{dist}(c,d)$, while $c' \in C$ — a contradiction.

For the gate property, we use induction on $\operatorname{dist}(c,C)$. If $c \in C$, then there is nothing to show, and if $\operatorname{dist}(c,C) = 1$, then we can immediately apply Corollary 1.7.40. Now assume $\operatorname{dist}(c,C) \geq 1$ and let $c' \in \Delta$ and $i \in I$ be such that $\operatorname{dist}(c',C) < \operatorname{dist}(c,C)$ and $c' \sim_i c$. Denote $\mathcal{P} = \mathcal{P}_i(c)$. Then by construction $\operatorname{proj}_{\mathcal{P}}(d) = c'$ and there exists a minimal gallery γ from c' via d to e by the induction hypothesis. Minimality yields that also $\operatorname{proj}_{\mathcal{P}}(e) = c'$. Thus $\operatorname{dist}(c,e) > \operatorname{dist}(c',e)$ and the gallery obtained by prepending $c \sim_i c'$ to γ is again minimal. Our conclusion follows by induction. ∎

We will evidently denote the unique chamber in C closest to c by $\operatorname{proj}_C(c)$.

Finally, Lemma 2.3.1 can be generalised to convex panel-closed sets as well, although in a slightly more technical fashion. Lemma 2.3.18 will prove useful for Proposition 3.1.9 later on.

Lemma 2.3.18. *Let C be a convex, panel-closed set. Let $\mathcal{P}$ be an i-panel that intersects both $\mathsf{S}_n(C)$ and $\mathsf{S}_{n+1}(C)$ nontrivially, where $n \geq 1$. Then the intersection $\mathcal{P} \cap \mathsf{S}_n(C)$ is a single chamber c.*

Moreover, let $d \in \mathsf{B}_{n+1}(C) \setminus \mathcal{P}$, let $d' = \operatorname{proj}_{\mathcal{P}}(d)$, and assume that $d' \in \mathsf{S}_{n+1}(C)$. Then either d is contained in an i-panel parallel to $\mathcal{P}$, or d' is j-adjacent to some chamber in $\mathsf{S}_n(C)$ with $m_{ij} = 2$.

Proof. Assume by means of contraposition that there is a pair of chambers $\{c_1, c_2\} \subseteq \mathcal{P} \cap \mathsf{S}_n(C)$ with $c_1 \neq c_2$. If $\operatorname{proj}_C(c_1) \neq \operatorname{proj}_C(c_2)$, then those projections are i-adjacent, so that $\mathcal{P}$ is parallel to a panel contained in C, and $\mathcal{P} \subseteq \mathsf{S}_n(C)$. Hence it must be the case that $\operatorname{proj}_C(c_1) = \operatorname{proj}_C(c_2)$; let d denote the common projection. Then the projection $\operatorname{proj}_{\mathcal{P}}(d)$ of d onto $\mathcal{P}$ must be a chamber in $\mathsf{S}_{n-1}(C)$, and the intersection $\mathcal{P} \cap \mathsf{S}_{n+1}(C)$ is again empty. Our first claim follows.

For the second claim, write $c_0 = \operatorname{proj}_C(c)$ and $d_0 = \operatorname{proj}_C(d)$. When $c_0 = d_0$ we can immediately apply Lemma 2.3.1. Otherwise, let $d' \sim e \sim \cdots \sim d$ be a minimal gallery, concave with respect to c_0 as in Proposition 2.2.3. If $\operatorname{dist}(e, c_0) \leq n+1$, we can again apply Lemma 2.3.1. In the other case, concavity implies that $\operatorname{dist}(c_0, d) = \operatorname{dist}(c_0, d') + \operatorname{dist}(d', d)$. Together with the gate property (Lemma 2.3.17) we now have two minimal galleries joining c_0 to d — one passing through d' and one passing through d_0.

The types of those galleries represent the same Coxeter group element, and are hence homotopic. In particular, we can find two chambers e and e' on the minimal gallery

$$c \sim \cdots \sim e' \sim_i e \sim \cdots \sim d$$

that passes through d_0, such that the panel containing $e \sim_i e'$ is parallel to $\mathcal{P}$ (containing $c \sim_i d'$). If $d = e$, then indeed d is contained in a panel parallel to $\mathcal{P}$. Otherwise $\{e, e'\} \subseteq \mathsf{B}_n(C)$, and then all chambers $\{c, d', e, e'\}$ lie in a common residue $\mathcal{R}$ of type $i \cup i^{\perp}$. In this case, the intersection $\mathcal{R} \cap \mathsf{S}_n(C)$ contains some chamber j-adjacent to d' with $j \in i^{\perp}$. ■

2.4 Automorphisms

Now that we have associated for every $i \in I$ a certain tree Γ_i (Definition 2.3.14) to the right-angled building Δ, on which automorphisms of Δ have a natural induced action, the following definition should not be surprising.

Definition 2.4.1. Let Δ be a right-angled building over I and let $i \in I$. An automorphism g of Δ is called *i-hyperbolic* or *i-elliptic*, if the automorphism of Γ_i induced by g is hyperbolic or elliptic, respectively.

Note that an i-elliptic automorphism stabilises a residue of Δ of type $I \setminus \{i\}$ or of type $\{i\} \cup \{i\}^{\perp}$ (setwise), but not necessarily a chamber of Δ.

The main goal in this section is to find analogues to various results of Section 1.4 for right-angled buildings, that will eventually lead us to a simplicity criterion similar to the one by Tits for groups acting on trees (Theorem 1.4.17).

Definition 2.4.2 (i-distance). For each $i \in I$, define the i-distance $\operatorname{dist}_i(c_1, c_2)$ between any two chambers c_1 and c_2 as the number of occurences of type i in a minimal gallery joining c_1 and c_2. Note that this number does not depend on the chosen minimal gallery. This distance function is only a pseudometric, since chambers in a common residue of type $I \setminus \{i\}$ are at i-distance zero.

The i-distance function is quasi-isometric to the graph-theoretical edge distance in the i-tree-wall tree: this is the content of the following lemma.

Lemma 2.4.3. *Let c_1 and c_2 be any two chambers of Δ. Let $\mathcal{R}_1$ and $\mathcal{R}_2$ be the residues of type $i^{\perp}$ containing c_1 and c_2 respectively, viewed as edges in the i-tree-wall tree Γ_i (Definition 2.3.14). Then*

$$\operatorname{dist}_{\Gamma_i}(\mathcal{R}_1, \mathcal{R}_2) = 2 \cdot \operatorname{dist}_i(c_1, c_2) + \epsilon \leq \operatorname{dist}(c_1, c_2), \qquad \textit{where } \epsilon \in \{-1, 0, 1\}.$$

Proof. The inequality is clear. For the equality, pick any minimal gallery γ from c_1 to c_2 in Δ and consider the induced path in the i-tree-wall tree, which is a path without backtracking. Notice that we have a one-to-one correspondence between i-adjacencies on γ on the one hand, and pairs of adjacent edges in Γ_i sharing a common vertex in the bipartition class of residues of type $i \cup i^\perp$ on the other hand. Consequently, if the path in Γ_i has even length, say $2n$, there are n adjacencies of type i on γ, and if the path has odd length, say $2n+1$, there are either n or $n+1$ adjacencies of type i on γ. Our conclusion follows. ■

Definition 2.4.4 (cobounded, minimal, combinatorially dense). Let G be a group acting on a right-angled building Δ by automorphisms.

- The action is *cobounded* if there exists a constant $r \in \mathbb{N}$ and a chamber $c \in \Delta$ such that

$$\Delta = \bigcup_{g \in G} \mathsf{B}_r(g\,.\,c),$$

  i.e., if every chamber lies at a uniformly bounded distance to some G-orbit.

- The action is *minimal* if G leaves no nontrivial convex subsystem of Δ invariant.
- The action is *combinatorially dense* if it is minimal and moreover for every $i \in I$, the induced action on the i-tree-wall tree Γ_i has no fixed point at infinity.

More generally, a group action on some metric space X is called *cobounded* (or more descriptively, *coarsely transitive*) if there is a constant $r > 0$ such that the r-neighbourhood of every orbit in X covers X completely. In symbols,

$$\big(\exists\, r > 0\big)\big(\forall\, x, y \in X\big)\big(\exists\, g \in G\big)\big(\mathrm{dist}(g\,.\,x, y) \leq r\big).$$

By now we have encountered three definitions of minimality and density: Definition 1.4.1 for actions on trees, Definition 2.4.4 for actions on general buildings, and Definition 1.6.7 for actions on $\mathrm{CAT}(0)$ spaces. Evidently these notions are not completely unrelated.

Proposition 2.4.5. *Let G be a group acting on a right-angled building Δ by automorphisms and assume that the diagram has no isolated nodes. Then the following are equivalent:*

(i) *the action of G on Δ is minimal (in the sense of Definition 2.4.4);*

(ii) *for every $i \in I$, the induced action of G on the i-tree-wall tree Γ_i is minimal (Definition 1.4.1).*

As a corollary, the following are equivalent:

(i) *the action of G on Δ is combinatorially dense;*

(ii) *for every $i \in I$, the induced action of G on the i-tree-wall tree Γ_i is geometrically dense.*

Proof. First assume G does not act minimally on Δ, i.e. there exists a nontrivial convex G-invariant subset C of chambers. Let $c \sim_i d$ be such that $c \in C$ and $d \notin C$. Using absence of isolated nodes in the diagram, take $j \in I$ such that $m_{ij} = 2$, and let $e \sim_j d$. By convexity, $e \notin C$. Now define C' as the set of vertices in Γ_i that correspond to residues of Δ having nonempty intersection with C,

$$C' = \big\{\mathcal{R} \in \mathrm{Res}_{I\setminus\{i\}}(\Delta) \;\big|\; \mathcal{R} \cap C \neq \emptyset\big\} \cup \big\{\mathcal{R} \in \mathrm{Res}_{\{i\}\cup\{i\}^\perp}(\Delta) \;\big|\; \mathcal{R} \cap C \neq \emptyset\big\}.$$

Note that C' is a nonempty G-invariant subtree of Γ_i. Moreover, we claim that e is not contained in any residue in C'. Indeed, if it were the case that $e \in \mathcal{R} \in C'$, let $e' \in \mathcal{R} \cap C$. The projection $\mathrm{proj}_{\mathcal{R}}(c)$ is either d or e, depending on the type of $\mathcal{R}$. By the gate property (Corollary 1.7.37), there is a minimal gallery joining c to e' passing through either d or e, contradicting the convexity of C. Hence, C' is a *nontrivial* G-invariant subtree, and the action of G on Γ_i is not minimal.

Conversely, assume that G does not act minimally on Γ_i for some $i \in I$, i.e. there exists a nontrivial G-invariant subtree D. Define D' as the set of chambers of Δ that lie in a residue corresponding to a vertex in D,

$$D' = \big\{ c \in \Delta \;\big|\; c \in \mathcal{R} \in D \text{ for some residue } \mathcal{R} \text{ of type } I \setminus \{i\} \text{ or type } \{i\} \cup \{i\}^{\perp} \big\}.$$

Clearly D' is a nontrivial G-invariant subset of chambers. Let $c, c' \in D'$ and consider a minimal gallery γ from c to d in Δ. Then γ descends to a path in Γ_i between the residues in D containing c and d. This path is contained in the subtree D, hence γ is contained in D'. In other words, D' is a nontrivial *convex* G-invariant subsystem, and the action of G on Δ is not minimal. ■

Remark 2.4.6. Unfortunately, the notions of minimal actions on buildings and their tree-wall trees on the one hand and $\mathrm{CAT}(0)$ spaces on the other hand appears to be more murky in comparison, and we do not have a simple analogue of Proposition 2.4.5 taking into account the building's Davis realisation. As an example, consider the direct product of two bi-infinite lines — an example already encountered in Example 1.7.16 (iv). Let g be the automorphism of this thin building Δ that shifts every chamber two spaces right and two spaces up, as in Figure 2.4, and let $G = \langle g \rangle$. Then for an arbitrary $c \in \Delta$, every chamber lies on some minimal gallery joining two chambers in $G\,.\,c$. Hence the action of G is minimal (in the sense of Definition 2.4.4). However, the Davis realisation $\mathbb{K}(\Delta)$ is a tessellation of the Euclidean plane, by the prototile in Figure 2.5, and the induced action of G leaves invariant every diagonal line of unit slope.

Even when G acts minimally on an *irreducible* building Δ, it may happen that the induced action on $\mathbb{K}(\Delta)$ is not minimal in the sense of Definition 1.6.7. An example is the action of the full automorphism group on the Davis realisation in Figure 1.16 (b). Indeed, the spherical rank two residues of the thin building Δ correspond to octagonal subcomplexes of $\mathbb{K}(\Delta)$. However, the convex hull of the Davis chambers's cone points (the black vertices in Figure 1.16) is obtained by "cutting off" four protruding corners of every such octagon, and is an invariant closed convex subspace.

Figure 2.4 additionally serves as an example of a minimal group action that is not *hereditarily* minimal, in the following sense.

Definition 2.4.7 (hereditarily $(*)$). Let G act on a right-angled building Δ by automorphisms. For any property $(*)$ of group actions, we say that the action of G is *hereditarily* $(*)$ if for every residue $\mathcal{R}$ of Δ, the induced action of the setwise stabiliser $G_{\{\mathcal{R}\}}$ on $\mathcal{R}$ again has property $(*)$.

For example, note that a transitive action on Δ is always hereditarily transitive.

Except in degenerate cases, coboundedness is the strongest of the properties in our Definition 2.4.4. In the following proof, the argument for minimality is due to the anonymous referee of [BDM21], and improves our first proof of the statement.

Proposition 2.4.8. *Let G be a group acting on a right-angled building Δ by automorphisms and assume that the diagram of Δ has no isolated nodes. If the action of G is cobounded, then it is combinatorially dense.*

Proof. First, we show the action of G to be minimal. Suppose by means of contraposition that Γ is a nontrivial convex subsystem stabilised by G. Let $c \sim_i d$ be such that $c \in \Gamma$ but $d \notin \Gamma$. Since the action is cobounded, there is some constant m such that every chamber lies within distance m of some chamber in the orbit $G\,.\,c$.

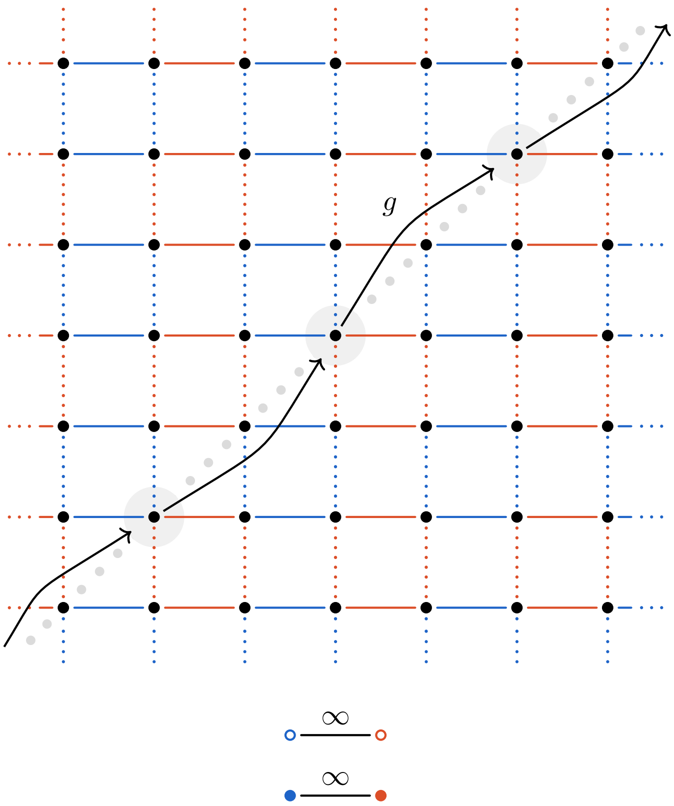

Figure 2.4. An example of a nonhereditarily minimal action.

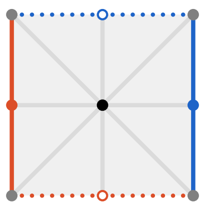

Figure 2.5. The Davis chamber associated to the example in Figure 2.4.
The Davis realisation of the thin building is a tessellation of the Euclidean plane.

Using the absence of isolated nodes, pick $j \in I$ such that $m_{ij} = \infty$ and let $\mathcal{R}$ be the $\{i, j\}$-residue containing the chamber c. Then $\mathcal{R}$ is a semiregular tree. Since Γ is convex, the projection of Γ to $\mathcal{R}$ is a connected subset of $\mathcal{R}$, which contains c but not d. As $\mathcal{R}$ is a tree, we find that the entire set Γ is contained in the wing $X_j(d)$.

Choose any geodesic ray in $\mathcal{R}$ of the form $(e_0, e_1, e_2, e_3, e_4, \dots)$ with $c = e_0$ and $d = e_1$. It follows by Lemma 2.3.12 and induction on n that the distance between e_{n+1} and any chamber of Γ is at least n, for all $n \in \mathbb{N}$. In particular, the distance

$$\operatorname{dist}(G\,.\,c, e_n) \geq \operatorname{dist}(\Gamma, e_n) \geq n$$

can be made arbitrarily large, contradicting the assumption that the action is cobounded.

Next, let $i \in I$ and consider the i-tree-wall tree Γ_i. By Lemma 2.4.3, the induced action of G on Γ_i is cobounded as well. Note that Γ_i is unbounded, as i is not an isolated node of the diagram. Thus we can take a bi-infinite path γ in Γ_i (defining two points at infinity). Since edges of γ cut Γ_i into two unbounded subtrees, a cobounded action cannot fix a point at infinity of the i-tree-wall tree.

In conclusion, the action of G is combinatorially dense. ■

Finally we mention a couple of important extension results. Recall first from Haglund and Paulin's Theorem 2.1.1 that every automorphism of a residue $\mathcal{R}$ of Δ can be extended to an automorphism of Δ (stabilising $\mathcal{R}$). In other words, the canonical morphism $\operatorname{Aut}(\Delta)_{\{\mathcal{R}\}} \to \operatorname{Aut}(\mathcal{R})$ is surjective. Pierre-Emmanuel Caprace demonstrated the following more refined version for panels.

Proposition 2.4.9. *Let $\mathcal{P}$ be a panel of a right-angled building Δ and let π be a permutation of the chambers in $\mathcal{P}$. Then there exists an automorphism $g \in \operatorname{Aut}(\Delta)$ satisfying the following properties:*

(i) *g stabilises $\mathcal{P}$;*

(ii) *the restriction $g|_{\mathcal{P}}$ is equal to π;*

(iii) *g fixes every chamber of Δ whose projection onto $\mathcal{P}$ is fixed by π.*

Proof omitted. This is [Cap14a, Proposition 4.2]; we illustrate the technique in Lemma 3.2.8.

Definition 2.4.10. We say that a group action on a building Δ is *strongly transitive* if it is transitive on the chamber-apartment pairs $(c, \mathcal{A})$ with $c \in \mathcal{A}$.

Proposition 2.4.11. *The automorphism group of a semiregular right-angled building Δ acts strongly transitively on Δ.*

Proof omitted. This is [KT12, Theorem B]; we also refer to [Cap14a, Proposition 6.1].

2.5 Colourings

In order to keep track of the local behaviour of a building automorphism, we introduce the notion of a colouring of the building. Throughout this section, Δ is a semiregular right-angled building with parameters $(q_i)_{i \in I}$.

Definition 2.5.1 (colouring). Consider a set Ω_i of cardinality q_i for every $i \in I$, the elements of which we call *i-colours* or *i-labels*. Then, a *legal colouring* of Δ is a map

$$\lambda \colon \Delta \to \prod_{i \in I} \Omega_i \colon c \mapsto (\lambda_i(c))_{i \in I}$$

satisfying the following properties for every $i \in I$ and for every i-panel $\mathcal{P}$:

(i) the restriction $\lambda_i|_{\mathcal{P}} \colon \mathcal{P} \to \Omega_i$ is a bijection;

(ii) for every $j \neq i$, the restriction $\lambda_j|_{\mathcal{P}} \colon \mathcal{P} \to \Omega_j$ is a constant map.

Example 2.5.2. Figure 2.6 provides an example of a legal colouring of the semiregular tree with parameters 3 and 4. Every chamber (i.e. edge) of the tree is allocated both a rank from $\{\mathsf{J}, \mathsf{Q}, \mathsf{K}\}$ and a suit from $\{\spadesuit, \heartsuit, \clubsuit, \diamondsuit\}$, in such a way that in each panel (i.e. star of a vertex), either the ranks are all equal and the suits are all different, or vice versa, depending on the type of the panel.

Lemma 2.5.3. *Let λ be a legal colouring of a right-angled building Δ of rank n. Let $c \in \Delta$ be any chamber and, for every $i \in I$, let $x_i \in \Omega_i$ be an i-colour. Then there exists a chamber $d \in \Delta$ such that $\operatorname{dist}(c, d) \leq n$ and $\lambda_i(d) = x_i$ for all $i \in I$.*

Proof. Consider the set $J = \{i \in I \mid \lambda_i(c) \neq x_i\}$. For any $j \in J$, consider the panel $\mathcal{P}_j$ of type j containing c. By definition of a colouring, $\mathcal{P}_j$ then contains a chamber c' with $\lambda_j(c') = x_j$ and $\lambda_i(c') = \lambda_i(c)$ for every $i \neq j$. Repeating this for every $j \in J$, we eventually obtain a gallery of length $|J| \leq n$, joining c to a chamber d with $\lambda_i(d) = x_i$ for all $i \in I$. ∎

The following proposition is Proposition 2.44 in [DMdSS18] and shows that a colouring is essentially unique, up to automorphism. The proof is worth repeating here.

Proposition 2.5.4. *Let λ and λ' be two legal colourings of a right-angled building Δ using identical colour sets. Let c and c' be two chambers such that $\lambda(c) = \lambda'(c')$. Then there exists an automorphism $g \in \operatorname{Aut}(\Delta)$ such that $g \,.\, c = c'$ and $\lambda' \circ g = \lambda$.*

Proof. For all $n \in \mathbb{N}$, consider the set

$$G_n = \{g \in \operatorname{Aut}(\Delta) \mid g \,.\, c = c' \text{ and } \lambda'(g \,.\, d) = \lambda(d) \text{ for all } d \in \mathsf{B}_n(c)\}.$$

We will inductively construct a sequence of elements g_n (with $n \in \mathbb{N}$), such that $g_n \in G_n$ for all n and that g_n and g_m agree on the ball $\mathsf{B}_m(c)$ whenever $m < n$. For $n = 0$, any automorphism g_0 with $g_0 \,.\, c = c'$ suffices; note that such an automorphism exists by Proposition 2.4.11.

Now assume that $n > 0$ and that we have constructed g_n with all the required properties. In order to define g_{n+1} we will construct an automorphism h_n that stabilises $\mathsf{B}_n(c)$ pointwise and that fixes the mismatching colours at $\mathsf{S}_{n+1}(c)$. We can then set $g_{n+1} = h_n \circ g_n$.

Let $\mathcal{P}$ be any i-panel that intersects both $\mathsf{S}_n(c)$ and $\mathsf{S}_{n+1}(c)$ nontrivially. Denote $\mathcal{P}' = g_n \,.\, \mathcal{P}$ and let $\pi_{\mathcal{P}}$ be the permutation of $g_n \,.\, \mathcal{P}$ that makes the following diagram commute.

$$\begin{tikzcd} \mathcal{P} \arrow[r, "g_n"] \arrow[dr, "\lambda_i"'] & \mathcal{P}' \arrow[r, "\pi_{\mathcal{P}}"] & \mathcal{P}' \arrow[dl, "\lambda'_i"] \\ & \Omega_i & \end{tikzcd}$$

By Proposition 2.4.9, $\pi_{\mathcal{P}}$ extends to an automorphism $\widetilde{\pi}_{\mathcal{P}}$ that fixes all chambers whose projection onto $\mathcal{P}$ is fixed by $\pi_{\mathcal{P}}$. We claim that $\widetilde{\pi}_{\mathcal{P}}$ fixes the set $\mathsf{B}_{n+1}(c') \setminus \mathcal{P}'$. Let $d \in \mathsf{B}_{n+1}(c') \setminus \mathcal{P}'$ and consider $d' = \operatorname{proj}_{\mathcal{P}'}(d)$. If $\operatorname{dist}(c', d') = n$, or in other words $d' = \operatorname{proj}_{\mathcal{P}'}(c')$, then by assumption on g_n we have that d' is fixed by $\pi_{\mathcal{P}}$. Hence d is fixed by $\widetilde{\pi}_{\mathcal{P}}$. Suppose now that $\operatorname{dist}(c', d') = n+1$. Then by Lemma 2.3.1, we have that $d' \sim_j e'$ for some chamber $e' \in \mathsf{S}_n(c')$ and $j \in I$ with $m_{ij} = 2$. Writing $e = g_n^{-1} \,.\, e'$, it follows that $d \sim_j e$. Hence

$$\lambda'_i(d') = \lambda'_i(e') = \lambda_i(e) = \lambda_i(d),$$

so that again, d' is fixed by $\pi_{\mathcal{P}}$ and d is fixed by $\widetilde{\pi}_{\mathcal{P}}$.

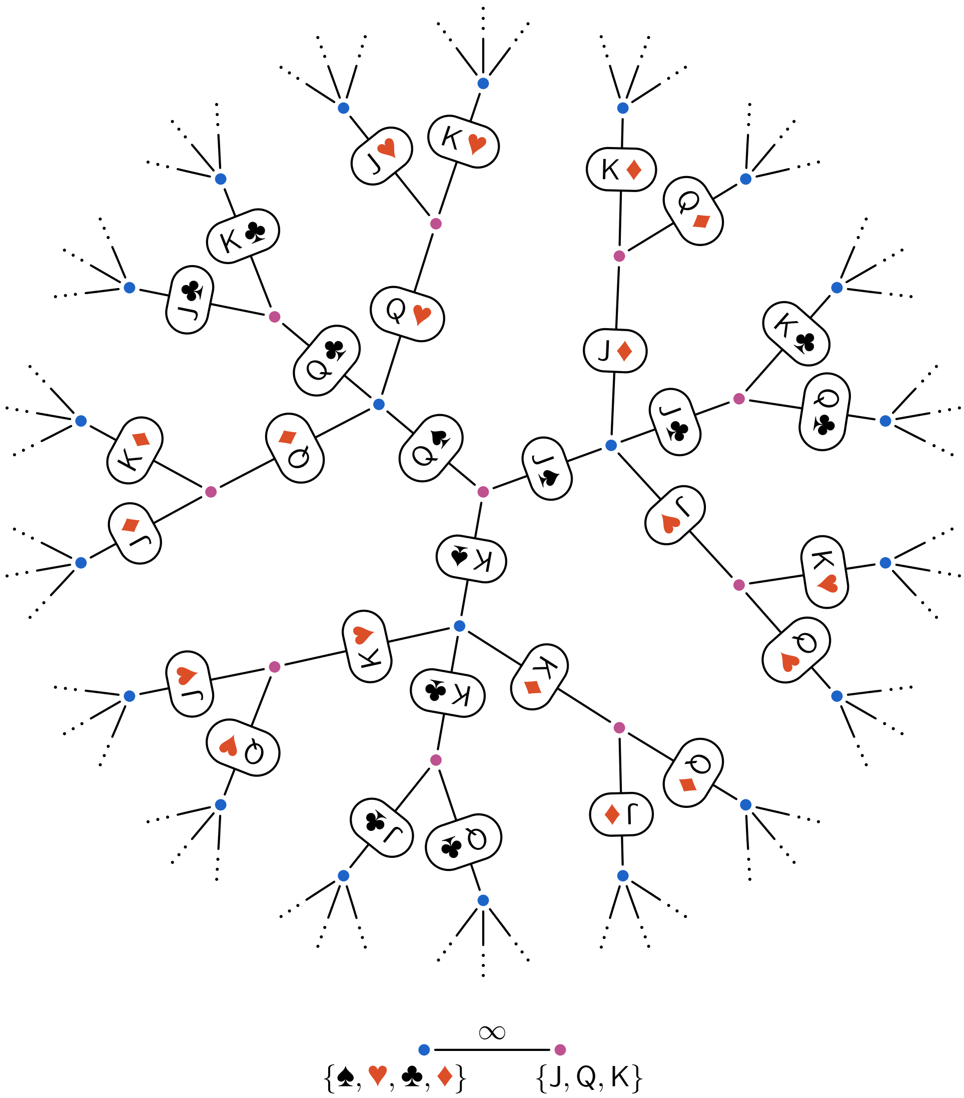

Figure 2.6. A colouring of a semiregular tree by cards suits and ranks.

We have thus constructed, for every panel $\mathcal{P}$ that intersects both $\mathsf{S}_n(c)$ and $\mathsf{S}_{n+1}(c)$, an automorphism $\widetilde{\pi}_{\mathcal{P}}$ of the building with the property that all chambers of $\mathsf{B}_{n+1}(c')$ that are moved by $\widetilde{\pi}_{\mathcal{P}}$ are contained in $\mathsf{S}_{n+1}(c') \cap \mathcal{P}$.

Now consider an arbitrary chamber $d \in \mathsf{S}_n(c')$ and let $i \in I$. There are two options: either $\mathcal{P}_i(d)$ intersects $\mathsf{S}_{n+1}(c')$ nontrivially, in which case we set $\alpha_{d,i} = \widetilde{\pi}_{\mathcal{P}_i(d)}$ — or $\mathcal{P}_i(d) \subseteq \mathsf{B}_n(c')$, in which case we set $\alpha_{d,i}$ equal to the identity. We can then define

$$\beta_d = \prod_{i \in I} \alpha_{d,i}$$

where the product is taken in arbitary order. Even though β_d might depend on the chosen order, its action on $\mathsf{B}_{n+1}(c')$ does not, since the sets of chambers of $\mathsf{B}_{n+1}(c')$ moved by different elements $\alpha_{d,i}$ are disjoint. Note that β_d fixes $\mathsf{B}_{n+1}(c') \setminus \mathsf{S}_1(d)$.

We now vary d along $\mathsf{S}_n(c')$. First we claim that if $d_1, d_2 \in \mathsf{S}_n(c')$ are distinct, then β_{d_1} and β_{d_2} restricted to $\mathsf{B}_{n+1}(c')$ have disjoint support. The only case we need to check is when $\mathsf{S}_1(d_1)$ and $\mathsf{S}_1(d_2)$ have a chamber $e \in \mathsf{S}_{n+1}(c')$ in common, say $d_1 \sim_i e \sim_j d_2$. Then

$$\lambda'_j(e) = \lambda'_j(d_1) = \lambda_j(g^{-1} . d_1) = \lambda_j(g^{-1} . e),$$

so that β_{d_1} fixes e. Similarly β_{d_2} fixes e, and our claim follows.

We now consider the product

$$h_n = \prod_{d \in \mathsf{S}_n(c')} \beta_d = \prod_{d \in \mathsf{S}_n(c')} \prod_{i \in I} \alpha_{d,i} \in \mathrm{Aut}(\Delta),$$

where the product is again taken in any order, and the action on $\mathsf{B}_{n+1}(c')$ is again independent of the chosen order. This automorphism h_n stabilises the ball $\mathsf{B}_n(c')$ pointwise. We now finally show that $g_{n+1} = h_n \circ g_n$ is an automorphism in G_{n+1} — it is immediately clear that $g_{n+1} . c = c'$ and that $\lambda'(g_{n+1} . d) = \lambda(d)$ for all $d \in \mathsf{B}_n(c)$, so we only need to check for $d \in \mathsf{B}_{n+1}(c)$ that

$$\lambda'_i(g_{n+1} . d) = \lambda'_i((\beta_d \circ g_n) . d) = \lambda'_i((\alpha_{d,i} \circ g_n) . d) = \lambda'_i((\pi_{\mathcal{P}_i(d)} \circ g_n) . d) = \lambda_i(d)$$

or

$$\lambda'_i(g_{n+1} . d) = \lambda'_i((\beta_d \circ g_n) . d) = \lambda'_i((\alpha_{d,i} \circ g_n) . d) = \lambda'_i(g_n . d) = \lambda_i(d),$$

depending on whether $\mathcal{P}_i(d)$ contains a chamber in $\mathsf{S}_n(c')$ or not. In any case, we conclude that $\lambda'_i(g_{n+1} . d) = \lambda_i(d)$, so that $g_{n+1} \in G_{n+1}$. Moreover, by construction, g_{n+1} agrees with g_n on the ball $\mathsf{B}_n(c)$. The sequence of elements $g_0, g_1, g_2, \ldots$ thus obtained, converges to an automorphism that satisfies the desired properties. ∎

The next property states that any finite set of chambers is contained in a single wing, and moreover we have control over the colour of the base chamber.

Proposition 2.5.5. *Let $C \subseteq \Delta$ be a finite subset of chambers and let $x \in X_i$ be an i-colour. Assume that the diagram of Δ has no isolated nodes. Then there exists a chamber $c \in \Delta$ such that $\lambda_i(c) = x$ and $C \subseteq X_i(c)$.*

Proof. Using the absence of isolated nodes, let $j \in I$ be such that $m_{ij} = \infty$. Let $\mathcal{R}$ be an arbitrary residue of type $\{i, j\}$ (which is a tree). The projection $\mathrm{proj}_{\mathcal{R}}(C)$ is a finite set of chambers in $\mathcal{R}$, and can be enclosed by a ball B of finite diameter. Let $\mathcal{P}'$ be any i-panel in $\mathcal{R} \setminus \mathsf{B}$.

Since $\mathcal{R}$ is a tree, it follows that $\operatorname{proj}_{\mathcal{P}'}(\operatorname{proj}_{\mathcal{R}}(C))$ is just a single chamber c'. If $\lambda_i(c') = x$, then we let $c = c'$. If $\lambda_i(c') \neq x$, then we let c be any chamber that is j-adjacent to the unique chamber in $\mathcal{P}'$ with i-colour equal to x. In both cases, $\lambda_i(c) = x$. Let $\mathcal{P}$ be the i-panel containing c.

By construction we have that $\operatorname{proj}_{\mathcal{P}}(\operatorname{proj}_{\mathcal{R}}(C)) = \{c\}$. Hence by Proposition 1.7.41 (iii), it follows that $\operatorname{proj}_{\mathcal{P}}(C) = \{c\}$, and we are done. ■

Lemma 2.5.6. *Let $I = J \sqcup \{k\}$ be a partition of the index set. Let $\mathcal{R}, \mathcal{R}'$ be two distinct J-residues, c, d chambers in $\mathcal{R}$, and c', d' chambers in $\mathcal{R}'$, such that $c \sim_k c'$ and $d \sim_k d'$. Finally, let $i \in J \setminus k^\perp$. Then $\lambda_i(c) = \lambda_i(c') = \lambda_i(d) = \lambda_i(d')$.*

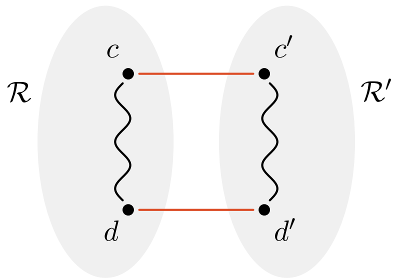

Proof. Let $\mathcal{P} = \mathcal{P}_k(c) = \mathcal{P}_k(c')$ and let $\mathcal{P}' = \mathcal{P}_k(d) = \mathcal{P}_k(d')$. Since residues are convex, it follows that $c = \operatorname{proj}_{\mathcal{P}}(d)$ and $c' = \operatorname{proj}_{\mathcal{P}}(d')$. Hence by Lemma 2.3.4, the panels $\mathcal{P}$ and $\mathcal{P}'$ are parallel. By Proposition 2.3.7 (ii), the chambers $\{c, c', d, d'\}$ are contained in a common residue of type $k \cup k^\perp$. Finally, since $i \notin k \cup k^\perp$, the i-colours of these chambers are identical. ■

In the following lemma, the fact that Δ is right-angled is crucial; the result does not at all hold for, say, general spherical buildings.

Lemma 2.5.7. *Let λ be a colouring of a right-angled building Δ using colour sets X_i. For each $i \in I$, let $Y_i \subseteq X_i$ be a subset of the i-colours such that $|Y_i| \geq 2$. Let $c_0 \in \Delta$ be any chamber, and let Γ be the set of chambers of Δ that are connected to c_0 by a gallery that only takes colours in the restricted sets Y_i. Then Γ is a semiregular subbuilding with the same type as Δ and with parameters $q_i = |Y_i|$.*

Proof. By [Wei03, Proposition 7.18] or [AB08, Theorem 4.66], in order to show that Γ is a building, it suffices to show that Γ is convex. Clearly, Γ is connected, hence consider an arbitrary gallery γ in Γ and a minimal gallery γ' in Δ such that γ and γ' are homotopic — we will demonstrate that γ' is fully contained in Γ as well.

We can reduce γ to γ' by means of only elementary contractions and elementary homotopies on the types, and claim that applying such an operation on the type of a gallery in Γ results in a gallery that is again contained in Γ.

(i) *Elementary contractions.* If a gallery γ in Γ of type $w \cdot ii \cdot w'$ is contracted to a gallery γ' of type $w \cdot w'$ or $w \cdot i \cdot w'$, then the chambers of γ' are a subset of the chambers of γ, and it follows that γ' is again a gallery in Γ.

(ii) *Elementary homotopies.* If a gallery in Γ of type ij with $m_{ij} = 2$ is transformed into a gallery of type ji, then this new gallery is again contained in Γ. Indeed, whenever $c \sim_i d \sim_j c'$ with $c, d, c' \in \Gamma$ and $c \sim_j d' \sim_i c'$ with $d' \in \Delta$, we have $\lambda_i(d') = \lambda_i(c)$ and $\lambda_j(d') = \lambda_j(c')$, so that $d' \in \Gamma$.

This shows that Γ is convex, and hence a subbuilding. The parameters and type of Γ are clear. ■

We can then easily obtain as a corollary …

Lemma 2.5.8. *Any finite set of chambers in a semiregular right-angled building Δ is contained in a finite convex set of chambers.*

Proof. First, we prove that any n chambers $\{c_1, \dots, c_n\}$ are contained in a semiregular *locally finite* subbuilding of Δ. For $n = 1$, this is obvious. For $n \geq 2$, we will make use of a legal colouring λ of Δ. Join c_1 to every chamber in $\{c_2, \dots, c_n\}$ by an arbitrary minimal gallery. Let $Y_i \subseteq X_i$ be the subset of all i-colours that occur as the i-colour of some chamber on one of these newly added minimal galleries, and note that Y_i is finite for every $i \in I$. Let Γ be the set of chambers of Δ that are connected to c_1 by a gallery that only takes colours in the sets $(Y_i)_{i \in I}$.

By Lemma 2.5.7, Γ is a locally finite subbuilding of Δ containing the chambers $\{c_1, \dots, c_n\}$. In Γ these chambers can be enclosed by a finite ball B (of finite radius). Moreover, as B is convex in Γ and Γ is convex in Δ, it follows that B is a finite convex set in Δ containing $\{c_1, \dots, c_n\}$. ■

2.6 Implosions

We have already met a few interesting "retraction-like" maps: the projections $\mathrm{proj}_{\mathcal{R}} \colon \Delta \to \mathcal{R}$ onto residues $\mathcal{R}$, and the retractions $\rho_{c,\mathcal{A}} \colon \Delta \to \mathcal{A}$ onto apartments $\mathcal{A}$ from Definition 1.7.35. In this section we develop a family of maps of a similar nature that we will call *implosions* and that provide a controlled way to "collapse" a coloured building by collapsing the colour sets. To the best of our knowledge, the construction is original.

Definition 2.6.1. Let Δ be a semiregular right-angled building over I and let λ be a legal colouring using colour sets Ω_i for every $i \in I$. Consider an equivalence relation $\equiv_i$ on every set Ω_i and let $I' = \{i \in I \mid \equiv_i \text{ is not the universal relation}\}$. Define a new semiregular right-angled building Δ' over I' with diagram induced by the diagram of Δ, with parameters $q'_i = |\Omega_i/{\equiv_i}|$ (for every $i \in I'$), and with a legal colouring λ' using the quotient $\Omega_i/{\equiv_i}$ as the set of i-colours.

Recall that a map $f \colon X \to Y$ between metric spaces is called *nonexpansive* if it does not increase distances, i.e if $\mathrm{dist}_Y(f(x_1), f(x_2)) \leq \mathrm{dist}_X(x_1, x_2)$ for every pair (x_1, x_2) of points in X.

Proposition 2.6.2. *Let Δ' be an implosion of Δ as in Definition 2.6.1. Let $c_0 \in \Delta$ be any chamber and let $c'_0 \in \Delta'$ be such that $\lambda'_i(c'_0) = [\lambda_i(c_0)]_i$ for every $i \in I'$. Then there exists a nonexpansive epimorphism τ of chamber systems from Δ onto Δ' such that $\lambda'_i(\tau(c)) = [\lambda_i(c)]_i$ for all $c \in \Delta$.*

Proof. We construct τ by induction on the distance from c_0, settling the induction base by declaring $\tau(c_0) = c'_0$. For $c \in \Delta$ such that $\mathrm{dist}(c_0, c) = n + 1$, let d be such that $\mathrm{dist}(c_0, d) = n$ and $d \sim_i c$. If $\lambda_i(c) \equiv_i \lambda_i(d)$ (in particular, if $i \notin I'$), then we set $\tau(c) = \tau(d)$. Otherwise we set $\tau(c)$ to be the unique chamber in Δ' i-adjacent to $\tau(d)$ such that $\lambda'_i(\tau(c)) = [\lambda_i(c)]_i$.

Since $\tau(c)$ *a priori* depends on the choice of d, we need to show that τ is well-defined. In order to do so, suppose that both d_1 and d_2 satisfy $\mathrm{dist}(c_0, d_1) = n = \mathrm{dist}(c_0, d_2)$ and $d_1 \sim_i c \sim_j d_2$. By Lemma 2.2.1, there exists a chamber e such that $\mathrm{dist}(c_0, e) = n - 1$ and $d_1 \sim_j e \sim_i d_2$. Since $\lambda_i(e) = \lambda_i(d_1)$, $\lambda_i(d_2) = \lambda_i(c)$, $\lambda_j(e) = \lambda_j(d_2)$, $\lambda_j(d_1) = \lambda_j(c)$, the images of the paths $e \sim d_1 \sim c$ and $e \sim d_2 \sim c$ end up in the same chamber in Δ', so $\tau(c)$ is indeed well-defined.

This extends τ to the whole of Δ, and by construction, we have $\lambda'_i(\tau(c)) = [\lambda_i(c)]_i$ for all $c \in \Delta$.

It is not hard to see that τ is surjective, since any gallery γ' in Δ' can be "lifted" to a gallery γ in Δ such that $\tau(\gamma) = \gamma'$. Explicitly, let γ' be a gallery

$$d'_0 \sim_{i_1} d'_1 \sim_{i_2} d'_2 \sim_{i_3} \cdots \sim_{i_n} d'_n$$

in Δ'. For every $1 \le k \le n$, let x_k be a representative of the equivalence class $\lambda'_{i_k}(d'_k)$. Let $d_0 \in \Delta$ be any chamber such that $\tau(d_0) = d'_0$, and, for every $1 \le k \le n$, let $d_k \in \Delta$ be the unique chamber i_k-adjacent to d_{k-1} such that the i_k-colour of d_k equals x_k. Note that the i_k-colours of d_{k-1} and d_k cannot be i_k-equivalent, since then d'_{k-1} and d'_k would be the same chamber in Δ'; hence d_k is well-defined. If we then let γ be the gallery

$$d_0 \sim_{i_1} d_1 \sim_{i_2} d_2 \sim_{i_3} \cdots \sim_{i_n} d_n$$

in Δ, we clearly have $\tau(\gamma) = \gamma'$. In particular we see that τ is surjective (using $d'_0 = c'_0$).

It remains to show that $\mathrm{dist}_{\Delta'}(\tau(c_1), \tau(c_2)) \le \mathrm{dist}_{\Delta}(c_1, c_2)$ for all $c_1, c_2 \in \Delta$. It suffices to show that the images of adjacent chambers either are adjacent or coincide, so assume that $c_1 \sim_j c_2$. We use induction on the distance to c_0. If $\mathrm{dist}(c_0, c_1) \ne \mathrm{dist}(c_0, c_2)$, then there is nothing to prove (by the very definition of τ). Suppose then that $\mathrm{dist}(c_0, c_1) = \mathrm{dist}(c_0, c_2) = n + 1$ and let d_1 satisfy $\mathrm{dist}(c_0, d_1) = n$ and $d_1 \sim_i c_1$. There are two possibilities. First, if $i = j$, then also $d_1 \sim_i c_2$ in Δ and either $\tau(c_1) = \tau(c_2)$ or $\tau(c_1) \sim_j \tau(c_2)$ in Δ'. Second, if $i \ne j$, then by Lemma 2.2.2 there exists a chamber d_2 such that $\mathrm{dist}(c_0, d_2) = n$ and $d_1 \sim_j d_2 \sim_i c_2$. Since $\lambda_i(c_1) = \lambda_i(c_2)$, $\lambda_i(d_1) = \lambda_i(d_2)$, $\lambda_j(c_1) = \lambda_j(d_1)$, $\lambda_j(c_2) = \lambda_j(d_2)$, and either $\tau(d_1) = \tau(d_2)$ or $\tau(d_1) \sim_j \tau(d_2)$, we see that either $\tau(c_1) = \tau(c_2)$ or $\tau(c_1) \sim_j \tau(c_2)$ as well. ■

Definition 2.6.3 (implosion). We call the pair (Δ', τ) from Proposition 2.6.2 the *implosion* of Δ with *centre* c_0 (with respect to the relations $\equiv_i$).

Corollary 2.6.4. *Let Δ be a semiregular right-angled building of type M over I, let $J \subseteq I$, and let Γ be the semiregular building of type M_J over J with the same parameters as Δ. Then there is a map $\varphi_J \colon \Delta \to \Gamma$ with the following properties:*

(i) *for every residue $\mathcal{R}$ of type J, the restriction $\varphi_J|_{\mathcal{R}}$ is an isomorphism;*

(ii) *for every residue $\mathcal{R}$ of type $I \setminus J$, the restriction $\varphi_J|_{\mathcal{R}}$ is a constant map.*

Proof. This follows from Proposition 2.6.2 by taking as equivalence relations $\equiv_i$ either the equality relation if $i \in J$ or the universal relation if $i \notin J$. ■

2.7 City products

In this section, we develop a construction for creating new right-angled buildings of a higher rank by glueing together lower rank buildings along another building. Our construction is inspired by the observation that the large-scale geometry of certain right-angled buildings (such as Figure 1.11) resembles that of a tree; the city product structure explains this behaviour in a broad sense. To the best of our knowledge, this product is original.

We start with some combinatorics, the goal of which will become clear later on.

Definition 2.7.1 (weak homotopy). Let $i, j \in I$ with $m_{ij} = 2$ and define the set

$$P(i,j) = \{ w \in \{i,j\}^* \mid w \text{ contains at least one } i \text{ and one } j \}.$$

A *weak homotopy* is a transformation of a word $w_1\, p\, w_2$ into a word $w_1\, p'\, w_2$ where $w_1, w_2 \in I^*$ and $p, p' \in P(i,j)$. Two words w and w' are *weakly homotopic* if w can be transformed into w' by a sequence of weak homotopies.

Definition 2.7.2 (normal form). Let $\prec$ be a total order on I. Endow I^* with the induced lexicographical order. Then every word $w \in I^*$ is homotopic to a unique lexicographically minimal word that we call the *normal form* of w.

Proposition 2.7.3. *Let $\prec$ be a total order on I.*

(i) *If two words are homotopic, then their normal forms are equal.*

(ii) *A word is reduced if and only if its normal form contains no consecutive duplicate letters.*

(iii) *The normal forms of weakly homotopic words are equal up to consecutive duplicate letters.*

Proof. Claim (i) follows immediately from the definitions. For (ii), let $w \simeq w_1\, ii\, w_2$ and assume that the normal form contains no subword ii. Mark the two letters i in $w_1\, ii\, w_2$ and write the normal form as $n_1\, i\, n_2\, i\, n_3$ (where the two letters i are the marked ones). If n_2 is not the empty word, then let k be its first letter; by assumption $k \neq i$. From the homotopy all letters in n_2 are contained in $\{i\} \cup \{i\}^\perp$. It follows that the normal form cannot be lexicographically minimal: if $i \prec k$, then the homotopic word $n_1\, ii\, n_2\, n_3$ is lexicographically smaller, and if $i \succ k$, then $n_1\, n_2\, ii\, n_3$ is smaller. Claim (ii) follows. For claim (iii), it suffices to note that the effect of a weak homotopy of a word on its normal form is that a subword $i^m j^n$ with $m \geq 1$, $n \geq 1$, is replaced by another such word. ∎

Corollary 2.7.4. *If two reduced words $w, w' \in I^*$ are weakly homotopic, then they are homotopic.*

Proof. Letting $\prec$ be any total order, this follows readily from Proposition 2.7.3 (ii) and (iii). ∎

Now let us go back to the building realm and define an operation on the diagrams first.

Definition 2.7.5 (city product). Let M be a diagram of rank n over the index set $\{1, \dots, n\}$, and for every $1 \leq \ell \leq n$, let M_ℓ be a given diagram over I_ℓ. Then we define an new diagram as follows:

(i) the index set is the disjoint union $I = \bigsqcup_{\ell=1}^n I_\ell$;

(ii) for every pair of elements $i \in I_\ell$ and $i' \in I_{\ell'}$ there are two cases. If $\ell = \ell'$ then we set $m_{ii'}$ equal to $m_{ii'}$ (considered in M_ℓ). If $\ell \neq \ell'$ then we set $m_{ii'}$ equal to $m_{\ell\ell'}$ (considered in M).

This defines a diagram over I that we call the *city product* of the diagrams $\{M_1, \dots, M_n\}$ over M and denote by $✠_M(\{M_1, \dots, M_n\})$. Clearly its rank is $\sum_{\ell=1}^n |I_\ell|$.

Note, the special case of a city product over an empty diagram (i.e. $m_{ij} = 2$ for all $1 \leq i \neq j \leq n$) results in nothing more than the disjoint union of the diagrams $M_1, \dots, M_n$ (or in other words, the diagram of the direct product of buildings associated to $M_1, \dots, M_n$). Two more examples are given in Figure 2.7.

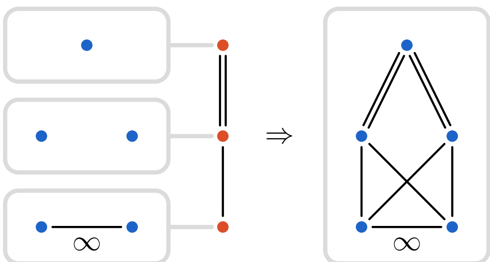

(a) A general example.

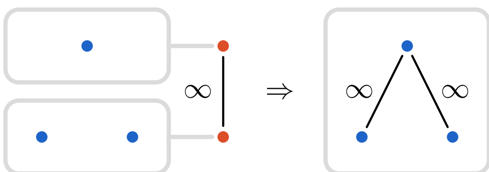

(b) A right-angled example.

Figure 2.7. City products of diagrams.

Lemma 2.7.6. $✠_M(\{M_1, \dots, M_n\})$ *with $n \geq 2$ is irreducible if and only if M is irreducible.*

Proof. This follows immediately from the definition. ■

We can now straightforwardly define city products of buildings.

Definition 2.7.7 (city product). Let M be a right-angled diagram over the index set $\{1, \dots, n\}$, and for every $1 \leq \ell \leq n$, let Δ_ℓ be a given semiregular right-angled building of type M_ℓ over I_ℓ. Then we define the *city product* of the buildings $\{\Delta_1, \dots \Delta_n\}$ over M as follows:

(i) the index set is the disjoint union $I = \bigsqcup_{\ell=1}^{n} I_\ell$;

(ii) the (right-angled) diagram is the city product of diagrams $\maltese_M(\{M_1, \dots, M_n\})$;

(iii) for every $i \in I$, the parameter q_i of the new building is the parameter q_i of Δ_ℓ, where $i \in I_\ell$.

Up to isomorphism, this defines a unique semiregular right-angled building by Theorem 2.1.1, that we denote by $\maltese_M(\{\Delta_1, \dots, \Delta_n\})$. There should be no confusion possible with the city product of diagrams. For ease of notation, for every $i \in I$, we also define $\ell(i)$ to be the number in $\{1, \dots, n\}$ such that $i \in I_{\ell(i)}$.

Note that for every $1 \leq \ell \leq n$, the residues of type $I_\ell \subseteq I$ of the city product $\maltese_M(\{\Delta_1, \dots, \Delta_n\})$ are isomorphic to the original building Δ_ℓ. As a special case of Corollary 2.6.4, we then obtain

Lemma 2.7.8. *Let $\Delta = \maltese_M(\{\Delta_1, \dots, \Delta_n\})$ be a city product and let $1 \leq \ell \leq n$. Then there is a map $\varphi_\ell \colon \Delta \to \Delta_\ell$ with the following properties:*

(i) *for every residue $\mathcal{R}$ of type I_ℓ the restriction $\varphi_\ell|_{\mathcal{R}} \colon \mathcal{R} \to \Delta_\ell$ is an isomorphism;*

(ii) *for every residue $\mathcal{R}$ of type $I \setminus I_\ell$ the restriction $\varphi_\ell|_{\mathcal{R}} \colon \mathcal{R} \to \Delta_\ell$ is a constant map.*

Proof. This follows immediately from Corollary 2.6.4. ■

We can then easily lift colourings of the subbuildings to a colouring of the full city product.

Lemma 2.7.9. *Let $\Delta = \maltese_M(\{\Delta_1, \dots, \Delta_n\})$ be a city product. For every $1 \leq \ell \leq n$, let λ^ℓ be a legal colouring of Δ_ℓ with colour sets Ω_i (where i ranges over I_ℓ). Then the collection of maps*

$$\lambda'_i = \lambda_i^{\ell(i)} \circ \varphi_{\ell(i)}$$

provides a legal colouring of Δ with colour sets Ω_i (where i ranges over I).

Proof. This follows immediately from Lemma 2.7.8 and the definition of legal colourings. ■

The city product construction over a diagram M essentially glues together smaller rank buildings as if they were chambers of a building of type M, hence the fact that the original buildings reemerge locally as residues (Lemma 2.7.8) should not be surprising. However, we can also recover a building of type M at the global scale by relaxing the adjacencies.

Definition 2.7.10 (skeletal building). Let $\Delta = \maltese_M(\{\Delta_1, \dots, \Delta_n\})$ be a city product, where M is a right-angled diagram over $\{1, \dots, n\}$. The *skeletal building* of Δ is the chamber system over the index set $\{1, \dots, n\}$ with the same chamber set as Δ, but with coarser adjacencies: we declare two chambers $c, d \in \Delta$ to be ℓ-adjacent if and only if they lie in the same residue of type I_ℓ.

We will prove in Proposition 2.7.13 that the skeletal building of a city product is, in fact, a building. First we need an auxiliary definition and some combinatorial lemmas, making the bridge between city products and weak homotopies.

Definition 2.7.11 (parkour map). Let Φ be the skeletal building of a city product Δ, so that Φ is a chamber system over $\{1, \ldots, n\}$ and Δ a chamber system over I. The *parkour map* is the map

$$r \colon I^* \to \{1, \ldots, n\}^*$$

that first replaces every letter $i \in I$ by $\ell(i) \in \{1, \ldots, n\}$ and then removes consecutive duplicates. Then when w is the type of a gallery in Δ that visits no chamber twice, $r(w)$ is the type of a gallery in Φ with the same extremities, but cutting short subgalleries in residues of type I_ℓ to a single jump of type ℓ. The maximal subwords of a word $w \in I^*$ with letters in a common subset I_ℓ of indices are called the *blocks* of w. These are precisely the maximal subwords such that the image under r is a single letter.

As an example, consider the index sets

$$I_1 = \{1_a, 1_b, 1_c\}, \quad I_2 = \{2_a, 2_b\}, \quad I_3 = \{3_a, 3_b, 3_c\}, \quad I = I_1 \cup I_2 \cup I_3.$$

Then for the word $w = 2_a\, 2_b\, 3_c\, 1_c\, 1_a\, 1_b\, 1_a\, 3_b$, the image is $r(w) = 2313$. The blocks are the words

$$2_a\, 2_b, \quad 3_c, \quad 1_c\, 1_a\, 1_b\, 1_a, \quad 3_b.$$

The interplay between words in I^* and words in $\{1, \ldots, n\}^*$ is not completely trivial — especially when considering reduced words. As illustrated in Figure 2.8, images of reduced words under the parkour map are not necessarily reduced, nor are images of equivalent words necessarily equivalent. Hence the following slightly technical lemma.

Lemma 2.7.12. *Let $u \in I^*$ and let $r \colon I^* \to \{1, \ldots, n\}^*$ be the parkour map.*

(i) *If $u \simeq u'$, then $r(u)$ and $r(u')$ are weakly homotopic (in the sense of Definition 2.7.1).*

(ii) *If we have a homotopy $r(u) \simeq v$, then there exists $u' \in I^*$ such that $u' \simeq u$ and $r(u') = v'$, where v' is the word obtained from v by removing consecutive duplicate letters.*

$$\begin{tikzcd}
u \arrow[rr, "\simeq"] \arrow[d, "r"] & & u' \arrow[d, "r"] \\
r(u) \arrow[r, "\simeq"] & v \arrow[r, dashed] & v'
\end{tikzcd}$$

(iii) *If u is reduced, then all blocks of u are reduced.*

(iv) *If all blocks of u are reduced and $r(u)$ is reduced, then u is reduced.*

(v) *If $u \simeq u'$ and both $r(u)$ and $r(u')$ are reduced, then $r(u) \simeq r(u')$.*

Proof. Claim (i) follows immediately from the definition.

To be more precise, consider an elementary homotopy $u = u_1\, ij\, u_2 \simeq u_1\, ji\, u_2$. If $\ell(i) = \ell(j)$, then the image under r remains invariant. Henceforth, we can assume that $\ell(i) \neq \ell(j)$. We distinguish three cases for the subword u_1 of u:

[case L.a] u_1 is nonempty and the last letter of u_1 is in $I_{\ell(i)}$,

[case L.b] u_1 is nonempty and the last letter of u_1 is in $I_{\ell(j)}$,

[case L.c] u_1 is the empty word, or the last letter of u_1 is neither in $I_{\ell(i)}$ nor $I_{\ell(i)}$.

Analogously, we distinguish three cases for u_2:

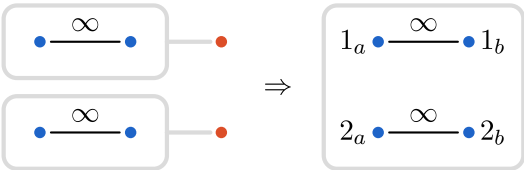

(a) The ambient (reducible) city product.

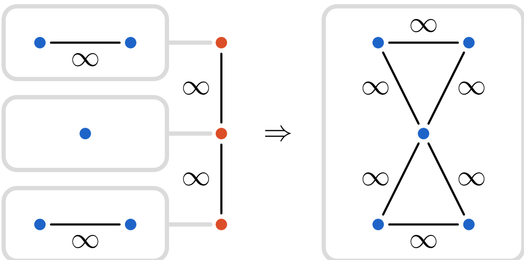

(b) An irreducible city product featuring case (a) as a residue.

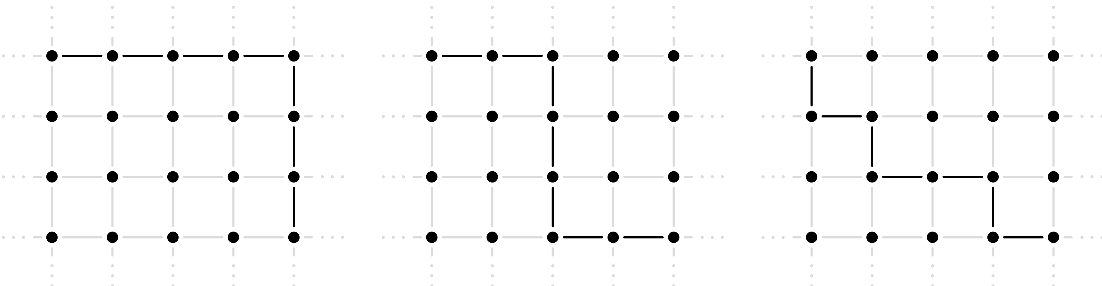

$r(w) = 12,$ $r(w) = 121,$ $r(w) = 212121.$

Figure 2.8. The effect of the parkour map on equivalent types of minimal galleries.

[case R.a] u_2 is nonempty and the first letter of u_2 is in $I_{\ell(i)}$,

[case R.b] u_2 is nonempty and the first letter of u_2 is in $I_{\ell(j)}$,

[case R.c] u_2 is the empty word, or the first letter of u_2 is neither in $I_{\ell(i)}$ nor $I_{\ell(i)}$.

Depending on the nine combinations of possibilities, the elementary homotopy $u_1\, ij\, u_2 \simeq u_1\, ji\, u_2$ transforms the image $r(u)$ by substituting some subword in $\{ij, ji, iji, jij, ijij, jiji\}$ into another such word; we refer to the table in Figure 2.9 for a more detailed overview. In any case the result is weakly homotopic to $r(u)$.

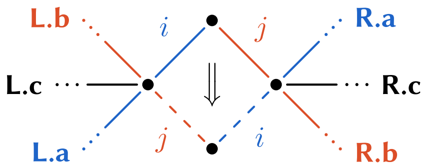

	R.a	**R.b**	**R.c**
L.a	$iji \mapsto iji$	$ij \mapsto ijij$	$ij \mapsto iji$
L.b	$jiji \mapsto ji$	$jij \mapsto jij$	$jij \mapsto ji$
L.c	$iji \mapsto ji$	$ij \mapsto jij$	$ij \mapsto ji$

Figure 2.9. The effect of an elementary homotopy on the image of the parkour map. In the table, we have simply written i and j instead of $\ell(i)$ and $\ell(j)$ for better readability.

For (ii), consider an elementary homotopy of $r(u)$, i.e. let $r(u) = v_1\, \ell_1\ell_2\, v_2 \simeq v_1\, \ell_2\ell_1\, v_2 = v$ with $1 \le \ell_1 \ne \ell_2 \le n$ and such that ℓ_1 and ℓ_2 commute in M. Then we can write $u = u_1\, b_1 b_2\, u_2$ where b_1 and b_2 are the blocks corresponding to ℓ_1 and ℓ_2, respectively. Since b_1 and b_2 have only letters in I_{ℓ_1} and in I_{ℓ_2} which are sets of pairwise commuting generators in the Coxeter system, we have a homotopy $u' = u_1\, b_2 b_1\, u_2 \simeq u_1\, b_1 b_2\, u_2$ that satisfies $r(u') = v'$.

For (iii), simply observe that *any* subword of a reduced word is reduced.

For (iv), assume by means of contraposition that every block of u is reduced while u is not, i.e. there is a homotopy $u \simeq w_1\, ii\, w_2$. Since every elementary homotopy simply swaps two adjacent letters, we can unambiguously define the initial blocks b_1 and b_2 of u that contain the two letters i. Since every block is assumed to be reduced, $b_1 \ne b_2$, hence $u = u_1\, b_1\, u_2\, b_2\, u_3$ such that u_2 is nonempty and $m_{ij} = 2$ for every letter j in u_2. The image then satisfies $r(u) = r(u_1)\, \ell\, r(u_2)\, \ell\, r(u_3)$ where ℓ is the index such that $i \in I_\ell$. By construction, ℓ commutes with every type in $r(u_2)$. Hence $r(u)$ is not reduced.

Finally (v) follows from (i) and Corollary 2.7.4. ∎

Proposition 2.7.13. *Let Φ be the skeletal building of the city product $\maltese_M(\{\Delta_1, \dots, \Delta_n\})$ of right-angled buildings. Then*

(i) *Φ is a right-angled building of type M over $\{1, \dots, n\}$;*

(ii) *Φ is semiregular with parameters $q_\ell = |\Delta_\ell|$ for every $1 \le \ell \le n$;*

(iii) *ℓ-panels of Γ (as sets of chambers) are I_ℓ-residues of Δ and vice versa;*

(iv) *the maps φ_ℓ with $1 \le \ell \le n$ provide a legal colouring of Γ with colour sets Δ_ℓ.*

Proof. The only nontrivial claim is that Φ is indeed a building of type M; the other claims follow immediately from the definitions. Let us first write out the Weyl distance function in Φ and then verify that it satisfies the necessary properties.

Denote by W_Δ be the Weyl group of the building Δ. Recall from Definition 1.7.17 the evaluation morphism $\varsigma_\Delta\colon I^* \to W_\Delta$. On the other hand, we have a Coxeter group W_Φ of type M, together with an evaluation morphism $\varsigma_\Phi\colon \{1,\dots,n\}^* \to W_\Phi$. As a final part of the setup, let $s\colon W_\Delta \to I^*$ be a section of ς_Δ with reduced images such that the word length $|r(s(w))|$ is minimal for every $w \in W_\Delta$. Finally, we can define

$$\delta_\Phi = \varsigma_\Phi \circ r \circ s \circ \delta_\Delta \colon \Phi \times \Phi \to W_\Phi.$$

Notice that the composition $\varsigma_\Phi \circ r \circ s$ is a map $W_\Delta \to W_\Phi$ but by no means a morphism of groups (or even monoids). We bundle all maps in the commutative diagram below.

$$\begin{tikzcd}
& W_\Delta \arrow[d, equal] \arrow[dr, "s"] & \\
\Delta \times \Delta \arrow[r, "\delta_\Delta"] \arrow[d, equal] & W_\Delta \arrow[d, dashed] & I^* \arrow[l, "\varsigma_\Delta"'] \arrow[d, "r"] \\
\Phi \times \Phi \arrow[r, "\delta_\Phi"] & W_\Phi & \{1,\dots,n\}^* \arrow[l, "\varsigma_\Phi"']
\end{tikzcd}$$

Clearly panels of Φ contain at least two chambers, since every such panel of Φ contains a panel of Δ. Now consider a reduced word v in $\{1,\dots,n\}^*$ — we need to demonstrate that $\delta_\Phi(c,d) = \varsigma_\Phi(v)$ if and only if there exists a gallery of type v from c to d in Φ.

First, assume that $\delta_\Phi(c,d) = \varsigma_\Phi(v)$. By definition of δ_Φ this means that the words $(r \circ s \circ \delta_\Delta)(c,d)$ and v are equivalent. Moreover, both words are reduced, and hence homotopic by Theorem 1.7.20. By Lemma 2.7.12 (ii), this homotopy can be realised in I^*, i.e. we can find a word $u \in I^*$ such that $u \simeq (s \circ \delta_\Delta)(c,d)$ and $r(u) = v$. The homotopy $u \simeq (s \circ \delta_\Delta)(c,d)$ yields that u is reduced, hence by the building axioms for Δ, there is a minimal gallery in Δ of type u from c to d. Then $r(u) = v$ is the type of a gallery in Φ from c to d.

Conversely, assume that γ is a gallery of type v from c to d in Φ. We can "lift" γ to a gallery $\overline{\gamma}$ in Δ with the same extremities, by replacing each ℓ-adjacency in γ by a minimal gallery in a residue of type I_ℓ of Δ. Let $\overline{v}$ be the type of $\overline{\gamma}$. Note that $r(\overline{v}) = v$ and that $\overline{v}$ is reduced by Lemma 2.7.12 (iv). Hence, we have $\delta_\Delta(c,d) = \varsigma_\Delta(\overline{v})$, so that $s(\delta_\Delta(c,d))$ and $\overline{v}$ are homotopic by Theorem 1.7.20. Then by Lemma 2.7.12 (v), the images $(r \circ s \circ \delta_\Delta)(c,d)$ and $r(\overline{v}) = v$ are homotopic, so that finally

$$\delta_\Phi(c,d) = (\varsigma_\Phi \circ r \circ s \circ \delta_\Delta)(c,d) = \varsigma_\Phi(v).$$

This concludes our proof that Φ is a right-angled building of type M. ■

As an example, note that the diagram in Figure 2.7 (b) gives rise to the Coxeter system of Figure 1.11. Indeed, the Coxeter complex is a thin building of rank three, but by taking the union of the red and blue adjacencies, we recognise a semiregular tree with parameters $q_1 = 4$ and $q_2 = 2$.

The city product of right-angled diagrams can be interpreted in a purely graph-theoretical way, and occurs in various disguises throughout the literature. Let us present the most common definition.

Definition 2.7.14 (modules). Let G be a simple, undirected graph. A subset X of the vertex set of G is called a *module* of G if it has the property that every vertex $v \notin X$ is either adjacent to all vertices in X or adjacent to no vertex in X. Note that every graph has *trivial modules*: the full set of all vertices, the singletons of single vertices, and the empty set.

Definition 2.7.15 (prime). A simple undirected graph is *prime* if it has no nontrivial modules.

Modules are generalisations of connected components, in the sense that the union X of connected components of a graph can be characterised by the property that every vertex $v \notin X$ is adjacent to no vertex in X. In particular, a prime graph is connected.

Modules were first described by Tibor Gallai in [Gal67]. They have since also been known as homogeneous sets, autonomous sets, partitive sets, or intervals. Modules play a crucial role in László Lovász's celebrated proof of the perfect graph theorem, but are mostly of algorithmic interest since there is an efficient way to compute a "modular decomposition" of a graph — a data structure that encodes all possible ways of decomposing a graph into modules, and serves as a stepping stone for more advances algorithms. The modular decomposition has been studied under various aliases as well, such as the substitution decomposition or prime tree decomposition. We refer to [MS99] for an introduction to the topic.

In the context of right-angled buildings, let us point out that a semiregular right-angled building is a nontrivial city product of lower rank buildings if and only if the underlying graph of its diagram is not prime. It is hence a natural question how "exclusive" prime graphs are. For small n, the prime graphs on n vertices can easily be found by hand. As it turns out, the path graph on four vertices is the unique smallest prime graph. There are four prime graphs on five vertices: the path graph, the cycle graph, and two others informally known as the house graph and the bull graph (Figure 2.10). However, as Table 2.1 clearly shows, the number of isomorphism classes of prime graphs tends to explode about as quickly as the total number of isomorphism classes of (connected) graphs; with increasing n, prime graphs become more and more common.

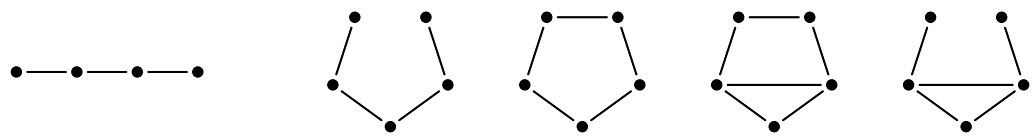

Figure 2.10. All prime diagrams of rank four or five.
For clarity, the labels ∞ on the edges are omitted.

n	total number of … diagrams on n nodes		
	unrestricted	irreducible	prime
1	1	1	0
2	2	1	0
3	4	2	0
4	11	6	1
5	34	21	4
6	156	112	26
7	1 044	853	260
8	12 346	11 117	4 670
9	274 668	261 080	145 870
10	12 005 168	11 716 571	8 110 356
11	1 018 997 864	1 006 700 565	804 203 096

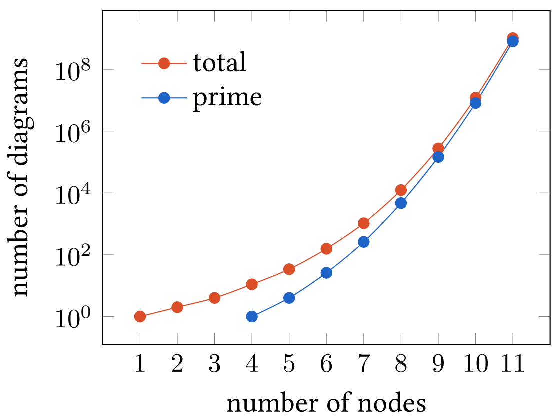

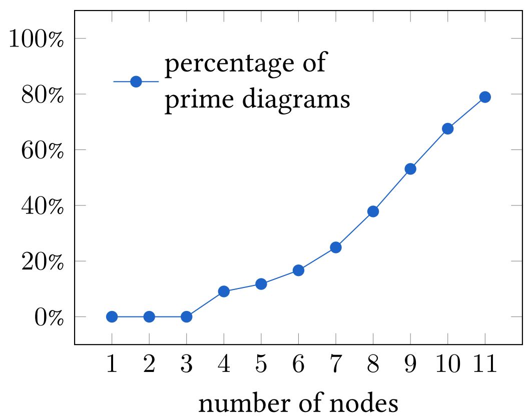

Table 2.1. A comparison of the number of unrestricted, irreducible, and prime diagrams. The first column is [OEISa], the second is [OEISb], the third is [OEISc].

3 Universal groups

> *You can never know everything, and part of what you know is always wrong. Perhaps even the most important part. A portion of wisdom lies in knowing that. A portion of courage lies in going on anyway.*
>
> — Robert Jordan, *Winter's Heart*

Now that we have built up a sufficiently large background, let us finally define the central objects of this thesis: universal groups over right-angled buildings.

3.1 Definition

A colouring of a building is useful for keeping track of the local behaviour of its automorphisms, in the following sense.

Definition 3.1.1 (local action). Let λ be a legal colouring of a semiregular right-angled building Δ with colour set Ω_i for every $i \in I$. Consider an automorphism $g \in \mathrm{Aut}(\Delta)$ and an arbitrary i-panel $\mathcal{P}$. Then we define the *local action of* g *at* $\mathcal{P}$ as the map

$$\sigma_\lambda(g, \mathcal{P}) = \lambda_i\big|_{g\mathcal{P}} \circ g\big|_{\mathcal{P}} \circ \lambda_i\big|_{\mathcal{P}}^{-1},$$

which is a permutation of Ω_i by definition of λ.

In other words, the local action $\sigma_\lambda(g, \mathcal{P})$ is the map that makes the following diagram commute.

$$\begin{tikzcd} \mathcal{P} \arrow[r, "g"] \arrow[d, "\lambda_i"] & g\mathcal{P} \arrow[d, "\lambda_i"] \\ \Omega_i \arrow[r, "\sigma_\lambda(g,\mathcal{P})"] & \Omega_i \end{tikzcd}$$

When the colouring λ is clear from the context, we will usually omit the explicit reference to λ and simply write $\sigma(g, \mathcal{P})$.

Lemma 3.1.2. *Let $g, h \in \mathrm{Aut}(\Delta)$ and let $\mathcal{P}$ be any panel. Then the local actions satisfy*

$$\sigma_\lambda(gh, \mathcal{P}) = \sigma_\lambda(g, h\mathcal{P}) \cdot \sigma_\lambda(h, \mathcal{P}) \qquad \textit{and} \qquad \sigma_\lambda(g, \mathcal{P})^{-1} = \sigma_\lambda(g^{-1}, g\mathcal{P}).$$

Proof. This follows immediately from the definition, or from one look at the diagrams below.

$$\begin{tikzcd} \mathcal{P} \arrow[r, "h"] \arrow[rr, bend left, "gh"] \arrow[d, "\lambda_i"] & h\mathcal{P} \arrow[r, "g"] \arrow[d, "\lambda_i"] & gh\mathcal{P} \arrow[d, "\lambda_i"] \\ \Omega_i \arrow[r, "\sigma_\lambda(h,\mathcal{P})"] \arrow[rr, bend right, "\sigma_\lambda(gh,\mathcal{P})"'] & \Omega_i \arrow[r, "\sigma_\lambda(g,h\mathcal{P})"] & \Omega_i \end{tikzcd} \qquad \begin{tikzcd} \mathcal{P} \arrow[r, bend left, "g"] \arrow[d, "\lambda_i"] & g\mathcal{P} \arrow[l, bend left, "g^{-1}"] \arrow[d, "\lambda_i"] \\ \Omega_i \arrow[r, bend left, "\sigma_\lambda(g,\mathcal{P})"] & \Omega_i \arrow[l, bend left, "\sigma_\lambda(g^{-1},g\mathcal{P})"] \end{tikzcd}$$

■

Proposition 3.1.3. *Let g be an automorphism of Δ. If $\mathcal{P}$ and $\mathcal{P}'$ are two parallel panels in Δ, then the local actions $\sigma_\lambda(g, \mathcal{P})$ and $\sigma_\lambda(g, \mathcal{P}')$ are identical.*

Proof. By Proposition 2.3.7 (i) and (ii), $\mathcal{P}$ and $\mathcal{P}'$ are panels of the same type i in a common residue of type $i \cup i^\perp$. This residue is isomorphic to the direct product $\mathcal{P}_0 \times \mathcal{R}_0$ of a building $\mathcal{P}_0$ of type i and $\mathcal{R}_0$ of type $i^\perp$. Thus, for every chamber $c \in \mathcal{P}$, a minimal gallery from c to $\mathrm{proj}_{\mathcal{P}'}(c)$ is unique and is contained in a residue of type $i^\perp$. In particular $\lambda_i(c) = \lambda_i(\mathrm{proj}_{\mathcal{P}'}(c))$. Moreover, since g is an automorphism, $g \,.\, \mathrm{proj}_{\mathcal{P}'}(c) = \mathrm{proj}_{g\,.\,\mathcal{P}'}(g\,.\,c)$.

Now consider the colour $\lambda_i(g\,.\,c)$. On the one hand, by definition of local actions,

$$\lambda_i(g\,.\,c) = \sigma_\lambda(g, \mathcal{P}) \,.\, \lambda_i(c).$$

On the other hand, using the projection onto $\mathcal{P}'$, we find that

$$\begin{aligned} \lambda_i(g\,.\,c) &= \lambda_i(\mathrm{proj}_{g\,.\,\mathcal{P}'}(g\,.\,c)) \\ &= \lambda_i(g\,.\,\mathrm{proj}_{\mathcal{P}'}(c)) \\ &= \sigma_\lambda(g, \mathcal{P}') \,.\, \lambda_i(\mathrm{proj}_{\mathcal{P}'}(c)) \\ &= \sigma_\lambda(g, \mathcal{P}') \,.\, \lambda_i(c). \end{aligned}$$

Since this holds for all $c \in \mathcal{P}$, we conclude that $\sigma_\lambda(g, \mathcal{P}) = \sigma_\lambda(g, \mathcal{P}')$.

A commutative diagram is worth a thousand words.

$$\begin{tikzcd}
 & g\mathcal{P} \arrow[rr, "\mathrm{proj}_{g\mathcal{P}'}"] \arrow[dd, "\lambda_i"] & & g\mathcal{P}' \arrow[dd, "\lambda_i"] \\
\mathcal{P} \arrow[ur, "g"] \arrow[rr, "\mathrm{proj}_{\mathcal{P}'}"] \arrow[dd, "\lambda_i", "x"'] & & \mathcal{P}' \arrow[ur, "g"] \arrow[dd, "\lambda_i"] & \\
 & \Omega_i \arrow[rr, equal] & & \Omega_i \\
\Omega_i \arrow[ur, "\sigma_\lambda(g,\mathcal{P})"'] \arrow[rr, equal] & & \Omega_i \arrow[ur, "\sigma_\lambda(g,\mathcal{P}')"'] &
\end{tikzcd}$$

∎

Definition 3.1.4 (universal group). Let $\boldsymbol{F}$ be a collection of permutation groups $F_i \leq \mathrm{Sym}(\Omega_i)$, indexed by $i \in I$. Let Δ be a semiregular right-angled building over I with parameters $q_i = |\Omega_i|$, equipped with a colouring λ using the sets Ω_i as i-colours. Then the *universal group* of $\boldsymbol{F}$ over Δ is by definition the group

$$\mathcal{U}^\lambda_\Delta(\boldsymbol{F}) = \big\{ g \in \mathrm{Aut}(\Delta) \;\big|\; \sigma_\lambda(g, \mathcal{P}) \in F_i \text{ for every } i \in I \text{ and every } \mathcal{P} \in \mathrm{Res}_i(\Delta) \big\}.$$

In words, $\mathcal{U}^\lambda_\Delta(\boldsymbol{F})$ is the group of automorphisms that locally act like permutations in F_i. We hence call the group F_i the *local groups*. We do not put any restrictions on the local groups: they are not required to be transitive, nor of finite degree. Note that $\mathcal{U}^\lambda_\Delta(\boldsymbol{F})$ is indeed a subgroup of $\mathrm{Aut}(\Delta)$ by Lemma 3.1.2.

Regarding the name, what is universal about these universal groups? Recall from Burger—Mozes's Proposition 1.5.7 in the tree setting that the universal groups over trees have a maximality property with respect to their local actions. Silva established in [DMdSS18, Proposition 3.7 (iv)] an analogue for right-angled buildings: *if the local groups are transitive and finite permutation groups, then any* *closed chamber-transitive subgroup $H \leq \mathrm{Aut}(\Delta)$ for which the local action on each i-panel is permutationally isomorphic to the group F_i for every $i \in I$, is in fact conjugate to a subgroup of $\mathcal{U}(\boldsymbol{F})$ in* $\mathrm{Aut}(\Delta)$. Although these results assume the local actions to be transitive, we still keep the name.

We note that we can recover the original Burger–Mozes groups: for a finite permutation group F the group $\mathcal{U}(F)$ in the sense of Definition 1.5.4 is easily shown to be isomorphic to the group $\mathcal{U}(\boldsymbol{F})$ over a tree in the sense of Definition 3.1.4, where the local data is given by the groups F and $\mathbb{Z}/2\mathbb{Z}$ (the cyclic group of order two simulating the edge inversions).

It is worth pointing out that the universal group construction is functorial, in the following sense.

Lemma 3.1.5. *For every $i \in I$, let $F_i \leq F_i' \leq \mathrm{Sym}(\Omega_i)$, and let $\boldsymbol{F}$ and $\boldsymbol{F}'$ be as in Definition 3.1.4. Then there is a natural inclusion $\mathcal{U}_\Delta^\lambda(\boldsymbol{F}) \leq \mathcal{U}_\Delta^\lambda(\boldsymbol{F}')$.*

Proof. This follows immediately from the definition. ■

As usual, when the building Δ or the colouring λ is clear from context, we will frequently simplify the notation to $\mathcal{U}(\boldsymbol{F})$, or occasionally even to $\mathcal{U}$. In any case, the choice of legal colouring λ is not essential for the structure of the universal group, as the following lemma shows.

Lemma 3.1.6. *Up to conjugagy, the subgroup $\mathcal{U}_\Delta^\lambda(\boldsymbol{F})$ of* $\mathrm{Aut}(\Delta)$ *is independent of the choice of* λ.

Proof. Let λ and λ' be distinct legal colourings. By Proposition 2.5.4 there is an automorphism g such that $\lambda' \circ g = \lambda$. A quick calculation then yields that

$$\begin{aligned}\sigma_\lambda(h, \mathcal{P}) &= \lambda_i\big|_{h\mathcal{P}} \circ h\big|_{\mathcal{P}} \circ \lambda_i\big|_{\mathcal{P}}^{-1} \\ &= \lambda_i'\big|_{gh\mathcal{P}} \circ g\big|_{h\mathcal{P}} \circ h\big|_{\mathcal{P}} \circ g\big|_{\mathcal{P}}^{-1} \circ \lambda_i'\big|_{g\mathcal{P}}^{-1} \\ &= \lambda_i'\big|_{gh\mathcal{P}} \circ \left(ghg^{-1}\right)\big|_{g\mathcal{P}} \circ \lambda_i'\big|_{g\mathcal{P}}^{-1} \\ &= \sigma_{\lambda'}({}^g h, g\mathcal{P})\end{aligned}$$

for every automorphism h and every panel $\mathcal{P}$, or more clearly in a commutative diagram:

$$\begin{tikzcd}
\mathcal{P} \arrow[rrr, "h"] \arrow[dr, "g"] \arrow[dd, "\lambda_i"'] & & & h\mathcal{P} \arrow[dl, "g"'] \arrow[dd, "\lambda_i"] \\
& g\mathcal{P} \arrow[r, "{}^g h"] \arrow[dl, "\lambda_i'"'] & gh\mathcal{P} \arrow[dr, "\lambda_i'"] & \\
\Omega_i \arrow[rrr, "\sigma_\lambda(h, \mathcal{P})", "\sigma_{\lambda'}({}^g h, g\mathcal{P})"'] & & & \Omega_i
\end{tikzcd}$$

Consequently, $h \in \mathrm{Aut}(\Delta)$ is an element of $\mathcal{U}^\lambda(\boldsymbol{F})$ if and only if ${}^g h$ is an element of $\mathcal{U}^{\lambda'}(\boldsymbol{F})$. ■

Definition 3.1.7 (panel group). For every panel $\mathcal{P}$ of type i, the *panel group* $\mathcal{U}|_{\mathcal{P}}$ is the subgroup of $\mathrm{Sym}(\mathcal{P})$ induced by the action of the panel stabiliser $\mathcal{U}_{\{\mathcal{P}\}}$ on the chambers in $\mathcal{P}$.

We conclude this introductory section with a definition that will turn out to be convenient when speaking about orbits; see, in particular, Proposition 3.1 below.

Definition 3.1.8 (harmony). Let J and K be disjoint subsets of I. Then two residues $\mathcal{R}$ and $\mathcal{R}'$ of type J are called *K-harmonious* if, for every $k \in K$, the (well-defined) colours $\lambda_k(\mathcal{R})$ and $\lambda_k(\mathcal{R}')$ lie in the same orbit of the local group F_k.

When $K = I \setminus J$, we abbreviate K-harmony to harmony. In particular for $J = \emptyset$, two chambers c and c' are harmonious if their colours $\lambda_i(c)$ and $\lambda_i(c')$ lie in the same F_i-orbit for every $i \in I$.

The final proposition of this introductory section is a rather technical one that allows us to extend partial automorphisms on certain subsets of the building to full automorphisms that are, intuitively speaking, "as close to universal group elements as possible". The proof idea is completely similar to the proof of Proposition 2.5.4, but requires the more general Lemma 2.3.18 since we are extending from a convex set instead of merely a chamber.

Proposition 3.1.9. *Let C be a convex, panel-closed subset of Δ. Let $\boldsymbol{F}$ be a collection of local groups for Δ. Let $g \in \mathrm{Aut}(\Delta)$ be an automorphism mapping chambers in C to harmonious chambers (with respect to $\boldsymbol{F}$). Then there exists an automorphism $h \in \mathrm{Aut}(\Delta)$ with the following properties:*

(i) $g|_C = h|_C$;

(ii) *for every i-panel $\mathcal{P}$, either $\mathcal{P}$ is parallel to a panel contained in C, or $\sigma_\lambda(h, \mathcal{P}) \in F_i$ (or both).*

Note that, in particular, h maps chambers to harmonious chambers.

Proof. First we define for every $n \in \mathbb{N}$ the sets

$$\mathsf{B}_n(C) = \{c \in \Delta \mid \mathrm{dist}(c, C) \leq n\},$$
$$\mathsf{S}_n(C) = \{c \in \Delta \mid \mathrm{dist}(c, C) = n\}.$$

Just like in Proposition 2.5.4, we will inductively construct a sequence of elements g_n (with $n \in \mathbb{N}$) such that $g_n|_C = h|_C$ for all n, such that property (ii) holds for all panels $\mathcal{P}$ contained in $\mathsf{B}_n(C)$, and such that g_n and g_m agree on the ball $\mathsf{B}_m(C)$ whenever $m < n$. Note that g_n maps chambers in $\mathsf{B}_n(C)$ to harmonious chambers. For $n = 0$, take $g_0 = g$.

Now assume that $n \geq 1$, and that we have constructed g_n with all the required properties. In order to define g_{n+1} we will construct an automorphism h_n that stabilises $\mathsf{B}_n(g_n . C)$ pointwise and that fixes the mismatching local actions at $\mathsf{S}_{n+1}(g_n . C)$ — we can then set $g_{n+1} = h_n \circ g_n$. Already let $C' = g_n . C = g . C$.

For convenience, we will henceforth call a panel *n-defect* if it intersects both $\mathsf{S}_n(C)$ and $\mathsf{S}_{n+1}(C)$ and it does not satisfy property (ii) with respect to g_n. Explicitly, if an i-panel $\mathcal{P}$ is n-defect, the local action $\sigma_\lambda(g_n, \mathcal{P}) \notin F_i$ and $\mathcal{P}$ is not parallel to any panel contained in C. Observe, if property (ii) holds for any panel, then for every parallel panel as well, since parallelism is transitive and local actions on parallel panels agree. This implies an n-defect panel is not parallel to any panel in $\mathsf{B}_n(C)$.

Let $\mathcal{P}$ be an n-defect i-panel. By Lemma 2.3.18, $\mathcal{P}$ intersects $\mathsf{S}_n(C)$ in a single chamber c. Denote $c' = g_n . c$ and $\mathcal{P}' = g_n . \mathcal{P}$. By the induction hypothesis, c and c' are harmonious. Hence we can find a permutation $f_\mathcal{P} \in F_i$ such that $f_\mathcal{P} . \lambda_i(c) = \lambda_i(c')$, and consecutively a permutation $\pi_\mathcal{P}$ of the chambers in $\mathcal{P}'$ that makes the diagram below commute.

$$\begin{tikzcd}
\mathcal{P} \arrow[r, "g_n"] \arrow[d, "\lambda_i"'] & \mathcal{P}' \arrow[r, "\pi_\mathcal{P}"] & \mathcal{P}' \arrow[d, "\lambda_i"] \\
\Omega_i \arrow[rr, "f_\mathcal{P}"] & & \Omega_i
\end{tikzcd}$$

By Proposition 2.4.9, $\pi_\mathcal{P}$ extends to an automorphism $\widetilde{\pi}_\mathcal{P}$ that fixes all chambers whose projection onto $\mathcal{P}$ is fixed by $\pi_\mathcal{P}$. Note that $\pi_\mathcal{P}$ fixes c' by construction, since both the local action of g_n at $\mathcal{P}$ and the target local action $f(\lambda_i(c), \lambda_i(c'))$ map the colour $\lambda_i(c)$ to the same image $\lambda_i(c')$. Consider $d \in \mathsf{B}_{n+1}(C') \setminus \mathcal{P}'$ and denote $d' = \mathrm{proj}_{\mathcal{P}'}(d)$. If $\mathrm{dist}(d', C') = n$, or in other words if $c' = d'$, then d' is fixed by $\pi_\mathcal{P}$ and d is fixed by $\widetilde{\pi}_\mathcal{P}$. Otherwise $\mathrm{dist}(d', C') = n + 1$. Then, Lemma 2.3.18 yields that $\mathcal{P}'$ is parallel to either the i-panel containing d, or an i-panel contained in $\mathsf{B}_n(C')$. Note that we have excluded the latter possibility by assuming that $\mathcal{P}$ violates property (ii).

We have thus constructed, for every n-defect panel $\mathcal{P}$, an automorphism $\widetilde{\pi}_{\mathcal{P}}$ of the building with the property that all chambers of $\mathsf{B}_{n+1}(C')$ that are moved by $\widetilde{\pi}_{\mathcal{P}}$ are contained in $\mathsf{S}_{n+1}(C') \cap \mathcal{R}$, with $\mathcal{R}$ the set of all chambers in a panel parallel to $\mathcal{P}'$ (in other words, the residue of type $i \cup i^{\perp}$ containing $\mathcal{P}'$).

Next we claim that no chamber in $\mathsf{S}_{n+1}(C')$ lies in more than one n-defect panel. Indeed, suppose $c_1 \sim_i d \sim_j c_2$ with $d \in \mathsf{S}_{n+1}(C')$ and $c_1, c_2 \in \mathsf{B}_n(C')$. As there is a unique chamber in C' nearest to d, we can apply Lemma 2.2.1 to find that $\mathcal{P}_i(d)$ and $\mathcal{P}_j(d)$ both are parallel to panels in $\mathsf{B}_n(C')$, so that $\mathcal{P}_i(d)$ and $\mathcal{P}_j(d)$ cannot be n-defect.

We now define
$$h_n = \prod_{\mathcal{P}} \widetilde{\pi}_{\mathcal{P}} \in \mathrm{Aut}(\Delta),$$
where the product runs over any set of arbitrarily chosen representatives of all equivalence classes of parallel n-defect panels. Note that h_n leaves invariant the set $\mathsf{B}_n(C')$, since all factors do. Hence the automorphism $g_{n+1} = h_n \circ g_n$ satisfies $g_{n+1}|_C = g_n|_C = g|_C$ in particular.

Moreover, let $\mathcal{P}$ be any panel contained in $\mathsf{B}_{n+1}(C)$. If $\mathcal{P} \subseteq \mathsf{B}_n(C)$, then property (ii) holds by the induction hypothesis. If $\mathcal{P} \subseteq \mathsf{B}_{n+1}(C)$, then $\mathcal{P}$ is parallel to some panel contained in $\mathsf{B}_n(C)$ and again property (ii) holds by induction. It hence only remains to verify that property (ii) is valid for an n-defect panel $\mathcal{P}$ with respect to g_{n+1}. Let $\widetilde{\mathcal{P}}$ be the representative of $\mathcal{P}$ in the class of parallel n-defect panels. Then
$$\sigma_\lambda(g_{n+1}, \mathcal{P}) = \sigma_\lambda(g_{n+1}, \widetilde{\mathcal{P}}) = \sigma_\lambda(h_n, g_n . \widetilde{\mathcal{P}}) \circ \sigma_\lambda(g_n, \widetilde{\mathcal{P}}) = \pi_{\widetilde{\mathcal{P}}} \circ \sigma_\lambda(g_n, \widetilde{\mathcal{P}}) = f_{\widetilde{\mathcal{P}}} \in f_i.$$

Finally, by construction, g_{n+1} agrees with g_n on the ball $\mathsf{B}_n(C)$. The sequence of automorphisms $g_0, g_1, g_2, \ldots$ thus obtained, converges to an automorphism satisfying the desired properties. ∎

The reason for the assumption that C is panel-closed in the above proposition should be evident: if a partial automorphism is only defined on a nontrivial subset of a panel, we would need to extend it to the full panel first. This would require us to assume the partial local actions to be extendable to a full permutation in the local group in the first place — panel-closure is a sufficient assumption to get rid of this additional technicality.

3.2 Permutational properties

As a first result, we motivate why we will regularly restrict our attention to irreducible buildings.

Lemma 3.2.1. *Let Δ be a reducible right-angled building Δ over I. Let $J_1, \ldots, J_m$ be the connected components of (the underlying graph of) the diagram of Δ. Then the universal group $\mathcal{U}_\Delta(\boldsymbol{F})$ splits as a direct product*
$$\mathcal{U}_\Delta(\boldsymbol{F}) \cong \mathcal{U}_{\mathcal{R}_1}(\boldsymbol{F}|_{J_1}) \times \cdots \times \mathcal{U}_{\mathcal{R}_m}(\boldsymbol{F}|_{J_m}),$$
where $\mathcal{R}_\ell$ is a residue of type J_ℓ for every $1 \leq \ell \leq m$.

Proof. Since Δ is isomorphic to the direct product $\mathcal{R}_1 \times \cdots \times \mathcal{R}_m$ and has automorphism group $\mathrm{Aut}(\Delta) \cong \mathrm{Aut}(\mathcal{R}_1) \times \cdots \times \mathrm{Aut}(\mathcal{R}_m)$, this follows immediately from the definition. ∎

Next, we can quite easily describe the orbits of $\mathcal{U}_\Delta(\boldsymbol{F})$ on chambers and residues of Δ.

Proposition 3.2.2. *Two residues lie in the same orbit of $\mathcal{U}_\Delta(\boldsymbol{F})$ if and only if they are of the same type and harmonious. In particular, two chambers c and c' lie in the same orbit of $\mathcal{U}_\Delta(\boldsymbol{F})$ if and only if their colours $\lambda_i(c)$ and $\lambda_i(c')$ lie in the same F_i-orbit for every $i \in I$.*

Proof. First, suppose that $g\,.\,\mathcal{R} = \mathcal{R}'$ for some $g \in \mathcal{U}^{\lambda}_{\Delta}(\boldsymbol{F})$ and fix $i \notin J$. Let $c \in \mathcal{R}$ be an arbitrary chamber and let $\mathcal{P} = \mathcal{P}_i(c)$. By definition of a legal colouring, λ_i is constant on $\mathcal{R}$ and $\mathcal{R}'$. Now let

$$f_i = \sigma_\lambda(g, \mathcal{P}) \in F_i,$$

and it readily follows that

$$f_i\,.\,\lambda_i(\mathcal{R}) = f_i\,.\,\lambda_i(c) = \left(\lambda_i\big|_{g\mathcal{P}} \circ g\big|_{\mathcal{P}} \circ \lambda_i\big|_{\mathcal{P}}^{-1} \circ \lambda_i\big|_{\mathcal{P}}\right)(c) = \lambda_i(g\,.\,c) = \lambda_i(g\,.\,\mathcal{R}).$$

Hence, for every $i \notin J$, the i-colours of $\mathcal{R}$ and $\mathcal{R}'$ lie in the same F_i-orbit.

Conversely, consider two harmonious residues $\mathcal{R}$ and $\mathcal{R}'$ of type J, and let $c \in \mathcal{R}$ and $c' \in \mathcal{R}'$ be two chambers with identical J-colours. For every $i \in I \setminus J$, let $f_i \in F_i$ be a permutation such that $f_i\,.\,\lambda_i(\mathcal{R}) = \lambda_i(\mathcal{R}')$. For every $i \in J$, let f_i be the identity permutation. Now define a “recolouring map”

$$\phi\colon \prod_{i\in I} \Omega_i \to \prod_{i\in I} \Omega_i \colon (\Omega_i)_{i\in I} \mapsto (f_i\,.\,\Omega_i)_{i\in I}$$

From Definition 2.5.1 it is clear that $\phi \circ \lambda$ is again a legal colouring of Δ. By construction, we have that $(\phi \circ \lambda)(c) = \lambda(c')$. Proposition 2.5.4 then provides an automorphism $g \in \mathrm{Aut}(\Delta)$ such that $g\,.\,c = c'$ and $\lambda \circ g = \phi \circ \lambda$. Now let $\mathcal{P}$ be any panel of type i and verify that

$$\sigma_\lambda(g, \mathcal{P}) = (\lambda_i \circ g)\big|_{\mathcal{P}} \circ \lambda_i\big|_{\mathcal{P}}^{-1} = (f_i \circ \lambda_i)\big|_{\mathcal{P}} \circ \lambda_i\big|_{\mathcal{P}}^{-1} = f_i.$$

In particular this means that $g \in \mathcal{U}^{\lambda}_{\Delta}(\boldsymbol{F})$, and since $g\,.\,\mathcal{R} = \mathcal{R}'$ we are done. ∎

From this we can immediately deduce the cardinalities of the orbit spaces.

Corollary 3.2.3. *Let $J \subseteq I$. The action of $\mathcal{U}(\boldsymbol{F})$ on the set of J-residues has finitely many orbits if and only if the local groups corresponding to $I \setminus J$ have finite orbits. More precisely,*

$$\left|\mathrm{Res}_J(\Delta)/\,\mathcal{U}(\boldsymbol{F})\right| = \prod_{i\notin J} \left|\Omega_i/F_i\right|.$$

In particular, $\mathcal{U}(\boldsymbol{F})$ acts transitively on the chambers of Δ if and only if all local groups are transitive.

We can now characterise when a universal group is trivial.

Corollary 3.2.4. *The universal group $\mathcal{U}_\Delta(\boldsymbol{F})$ is nontrivial if and only if the diagram of Δ has at least one edge or at least one local group is nontrivial.*

Proof. If the diagram has at least one edge, then Δ has a residue of rank two isomorphic to a tree. Such a residue contains an infinitude of identically coloured chambers. Hence by Proposition 3.2.2 $\mathcal{U}(\boldsymbol{F})$ is nontrivial. If on the other hand the diagram is the union of isolated nodes, then Δ is isomorphic to a direct product of complete graphs, and $\mathcal{U}(\boldsymbol{F})$ is simply the direct product of the local groups. The result follows. ∎

Corollary 3.2.5. *The action of $\mathcal{U}(\boldsymbol{F})$ on Δ is cobounded. More precisely, if Δ has rank n, every ball in Δ of radius n contains a representative chamber of every $\mathcal{U}(\boldsymbol{F})$-orbit.*

Proof. Let $\mathbf{Id}$ be the trivial subgroups of $\mathrm{Sym}(\Omega_i)$, indexed by $i \in I$. We claim that the action of $\mathcal{U}(\mathbf{Id})$ is cobounded. Indeed, let c and c' be any two chambers; then by Lemma 2.5.3, there exists a chamber d in $\mathsf{B}_n(c')$ such that $\lambda(c) = \lambda(d)$, and by Proposition 3.2.2, c and d lie in the same orbit of $\mathcal{U}(\mathbf{Id})$. Since $\mathcal{U}(\mathbf{Id}) \leq \mathcal{U}(\boldsymbol{F})$, it follows that the action of $\mathcal{U}(\boldsymbol{F})$ is cobounded as well. ∎

Corollary 3.2.6. *Let Δ be a right-angled building such that the diagram does not have isolated nodes. Then the action of $\mathcal{U}(\boldsymbol{F})$ on Δ is combinatorially dense.*

Proof. This follows immediately from Corollary 3.2.5 and Proposition 2.4.8. ■

Now, we come to the principal reason why right-angled buildings provide a suitable generalisation of the Burger–Mozes universal groups: the property that any permutation of the chambers in any panel occurs as the local action of some automorphism.

Lemma 3.2.7. *Let f be a permutation in F_i and let $\mathcal{P}$ be an i-panel. Then there exists an automorphism $g \in \mathcal{U}(\boldsymbol{F})$ with the following properties:*

(i) *g stabilises $\mathcal{P}$;*

(ii) *the local action is equal to f at every i-panel;*

(iii) *the local action is trivial at every other panel.*

Proof. Define a new colouring λ', setting $\lambda'_i = f^{-1} \circ \lambda_i$ and leaving $\lambda'_j = \lambda_j$ unchanged for $j \neq i$. Pick any colour $x \in \Omega_i$, let $c \in \mathcal{P}$ be the chamber with $\lambda_i(c) = x$, and let $c' \in \mathcal{P}$ be the chamber with $\lambda_i(c') = f.x$. Now apply Proposition 2.5.4 to find an automorphism g such that $g.c = c'$ and $\lambda' \circ g = \lambda$. In particular g stabilises $\mathcal{P}$. Moreover, similarly to the proof of Proposition 3.2.2, we see that the local actions are either equal to f (on panels of type i) or trivial (on panels of type $j \neq i$). Thus g satisfies all required properties. ■

We can slightly modify the obtained automorphism g from Lemma 3.2.7, using the same technique from [Cap14a, Proposition 4.2].

Lemma 3.2.8. *Let f be a permutation in F_i and let $\mathcal{P}$ be an i-panel. Then there exists an automorphism $h \in \mathcal{U}(\boldsymbol{F})$ with the following properties:*

(i) *h stabilises $\mathcal{P}$;*

(ii) *the local action of h at $\mathcal{P}$ is equal to f;*

(iii) *h fixes all chambers c with the property that $\lambda_i(\mathrm{proj}_{\mathcal{P}}(c))$ is fixed by f.*

Proof. Let g be the automorphism obtained from Lemma 3.2.7. We shall modify g along the wings of chambers in $\mathcal{P}$ corresponding to fixed points of f in order to satisfy property (iii). Define

$$h \colon \Delta \to \Delta \colon c \mapsto \begin{cases} c & \text{if } f \text{ fixes } \lambda_i(\mathrm{proj}_{\mathcal{P}}(c)); \\ g.c & \text{otherwise.} \end{cases}$$

Note that $g|_{\mathcal{P}} = h|_{\mathcal{P}}$ hence properties (i) and (ii) are automatic. The restriction of h to a wing of $\mathcal{P}$ is either the identity or coincides with the restriction of g. In addition, h satisfies property (iii) by construction. It remains to check that h is an automorphism of Δ and that all local actions of h are in the prescribed local groups.

In order to show that h preserves the Weyl distance, let c and d be any two chambers and denote by c' and d' their projections on $\mathcal{P}$. If $c' = d'$, then either $(hc, hd) = (c, d)$ or $(hc, hd) = (gc, gd)$; in both cases we indeed have $\delta(hc, hd) = \delta(c, d)$. Now assume that $c' \neq d'$. Let $\mathcal{R}$ be the residue of type $i \cup i^{\perp}$ containing $\mathcal{P}$ and let c'' and d'' be the projections of c and d on $\mathcal{R}$, respectively. Then c and c'' lie in a common wing of $\mathcal{P}$. As wings are convex by Lemma 2.3.10, it follows that h preserves the Weyl distance from c to c''. Similarly, $\delta(hd'', hd) = \delta(d'', d)$. Finally, since the restriction of h to $\mathcal{R}$ coincides with the restriction of g, we also have that h preserves the Weyl distance from c'' to d'' as well. Our claim that $\delta(hc, hd) = \delta(c, d)$ now follows from Lemma 2.3.11.

To conclude the proof, let $\mathcal{P}'$ be any panel. If $\mathcal{P}'$ is contained in a single wing of $\mathcal{P}$, then the local action $\sigma_\lambda(h, \mathcal{P}')$ is either the identity or $\sigma_\lambda(g, \mathcal{P}')$. Otherwise $\mathcal{P}'$ intersects at least two wings of $\mathcal{P}$ nontrivially. Then by Lemma 2.3.4, $\mathcal{P}'$ and $\mathcal{P}$ are parallel, and $\sigma_\lambda(h, \mathcal{P}') = f$ by Proposition 3.1.3. In conclusion, every local action of h is either trivial or equal to f. ■

Proposition 3.2.9. *Let $\mathcal{P}$ be a panel of type i. Then the groups $\mathcal{U}|_{\mathcal{P}} \leq \mathrm{Sym}(\mathcal{P})$ and $F_i \leq \mathrm{Sym}(\Omega_i)$ are permutationally isomorphic.*

Proof. Consider the local action morphism

$$\sigma_\lambda(\bullet, \mathcal{P})\colon\ \mathcal{U}_{\{\mathcal{P}\}} \to \mathrm{Sym}(\Omega_i)\colon g \mapsto \sigma_\lambda(g, \mathcal{P}).$$

By Lemma 3.2.7, the image of this morphism is the full local group F_i; the kernel is the pointwise stabiliser $\mathcal{U}_{(\mathcal{P})}$. Note that $\mathcal{U}|_{\mathcal{P}}$ and $\mathcal{U}_{\{\mathcal{P}\}} \,/\, \mathcal{U}_{(\mathcal{P})}$ are isomorphic (as abstract groups). The universal property of quotients then yields a natural isomorphism $\sigma|_{\mathcal{P}}$ from $\mathcal{U}|_{\mathcal{P}}$ to F_i.

$$\begin{tikzcd}
\mathcal{U}_{\{\mathcal{P}\}} \arrow[r] \arrow[drr, two heads, "\sigma_\lambda(\bullet,\mathcal{P})"'] & \mathcal{U}_{\{\mathcal{P}\}} \,/\, \mathcal{U}_{(\mathcal{P})} \arrow[r, "\cong"] & \mathcal{U}|_{\mathcal{P}} \arrow[d, "\sigma|_{\mathcal{P}}"] \\
 & & F_i
\end{tikzcd}$$

Moreover, for every $g \in \mathcal{U}|_{\mathcal{P}}$ and every $c \in \mathcal{P}$, we have that

$$\lambda_i(g\,.\,c) = \sigma\big|_{\mathcal{P}}(g)\,.\,\lambda_i(c),$$

so that the pair $(\sigma|_{\mathcal{P}}, \lambda_i)$ is an isomorphism of permutation groups, as required. ■

Next, we present a strengthening of one property mentioned in Theorem 2.1.1, namely the fact that automorphisms of residues can be extended to automorphisms of the full right-angled building: we can do so while keeping the local actions under control.

Proposition 3.2.10. *Let $J \subseteq I$ and let $\mathcal{R}$ be a residue of type J. Then every automorphism of $\mathcal{R}$ in $\mathcal{U}_{\mathcal{R}}(\boldsymbol{F}|_J)$ extends to an automorphism of Δ in $\mathcal{U}_\Delta(\boldsymbol{F})$.*

Proof. This is a corollary of Proposition 3.1.9, taking $\mathcal{R}$ to be the convex, panel-closed set. ■

Corollary 3.2.11. *Let $J \subseteq I$ and let $\mathcal{R}$ be a residue of type J. Then $\mathcal{U}_\Delta(\boldsymbol{F})_{\{\mathcal{R}\}}|_{\mathcal{R}} = \mathcal{U}_{\mathcal{R}}(\boldsymbol{F}|_J)$.*

Proof. On the one hand, the restriction of an automorphism in $\mathcal{U}_\Delta(\boldsymbol{F})_{\{\mathcal{R}\}}$ to $\mathcal{R}$ clearly is an automorphism in $\mathcal{U}_{\mathcal{R}}(\boldsymbol{F}|_J)$. On the other hand, by Proposition 3.2.10, every automorphism in $\mathcal{U}_{\mathcal{R}}(\boldsymbol{F}|_J)$ occurs as the restriction of an automorphism in $\mathcal{U}_\Delta(\boldsymbol{F})_{\{\mathcal{R}\}}$ to $\mathcal{R}$. ■

Let us now, as an example, calculate the universal group over a tree in a very specific case.

Proposition 3.2.12. *Set $I = \{\circ, \bullet\}$ with $m_{\circ\bullet} = \infty$ and let $q_\circ \geq 2$ and $q_\bullet \geq 2$ be integers. Let $F_\circ$ be the trivial permutation group of degree $q_\circ$ and let $F_\bullet$ be a regular permutation group of degree $q_\bullet$. Then $\mathcal{U}(\boldsymbol{F})$ is isomorphic to the free product of $q_\circ$ copies of $F_\bullet$.*

Proof. Let $\mathcal{P}_\circ$ be a panel of type $\circ$. Consider the partition of the tree

$$\Delta = \bigsqcup_{c \in \mathcal{P}_\circ} X_\circ(c)$$

into wings with base chambers in $\mathcal{P}_\circ$. For every chamber $c \in \mathcal{P}_\circ$ and every permutation $f \in F_\bullet$, let $\widetilde{f}$ be the automorphism from Lemma 3.2.8, stabilising $\mathcal{P}_\bullet(c)$ with local action f on $\bullet$-panels and the identity on $\circ$-panels. Note that the set

$$G_c = \left\{ \widetilde{f} \;\middle|\; f \in F_\bullet \right\}$$

is a subgroup of the stabiliser of $\mathcal{P}_\bullet(c)$ in $\mathcal{U}(\boldsymbol{F})$, canonically isomorphic to the local group $F_\bullet$. Let $d \neq d'$ be two different chambers in $\mathcal{P}_\circ$ and let g be any nontrivial automorphism in G_d. Since g has no fixed points in $\mathcal{P}_\bullet(d)$, we have that $g\,.\,X_\circ(d') \subseteq X_\circ(d)$. By the ping-pong lemma, it follows that the subgroups G_c generate a free product (where c ranges over $\mathcal{P}_\circ$).

Since $F_\bullet$ is transitive, Proposition 3.2.2 yields that $\mathcal{U}(\boldsymbol{F})$ acts transitively on the panels of type $\circ$. Let $\mathcal{P}$ be any such panel; we use induction on $\operatorname{dist}(\mathcal{P}, \mathcal{P}_\circ)$ to prove that $\langle G_c \mid c \in \mathcal{P}_\circ \rangle$ acts transitively on $\operatorname{Res}_\circ(\Delta)$ as well. For $\mathcal{P} = \mathcal{P}_\circ$ there is nothing to show. Now assume that $\operatorname{dist}(\mathcal{P}, \mathcal{P}_\circ) = n + 1$, and let $\mathcal{P}'$ be the panel such that $\operatorname{dist}(\mathcal{P}, \mathcal{P}') = n$ and containing some chamber d adjacent to $\mathcal{P}_\circ$. Then by local transitivity there exists $f \in G_d$ such that $f\,.\,\mathcal{P}_\circ = \mathcal{P}'$. Moreover, note that

$$\operatorname{dist}(f^{-1}\,.\,\mathcal{P}, \mathcal{P}_\circ) = \operatorname{dist}(\mathcal{P}, f\,.\,\mathcal{P}_\circ) = \operatorname{dist}(\mathcal{P}, \mathcal{P}') = n,$$

so that by the induction hypothesis, there exists some $g \in \langle G_c \mid c \in \mathcal{P}_\circ \rangle$ such that $g\,.\,\mathcal{P}_\circ = f^{-1}\,.\,\mathcal{P}$. Then $fg \in \langle G_c \mid c \in \mathcal{P}_\circ \rangle$ satisfies $fg\,.\,\mathcal{P}_\circ = \mathcal{P}$ and our claim follows.

Since the stabilisers of $\mathcal{P}_\circ$ in $\langle G_c \mid c \in \mathcal{P}_\circ \rangle$ and in $\mathcal{U}(\boldsymbol{F})$ agree, it follows that $\mathcal{U}(\boldsymbol{F})$ is generated by the subgroups G_c (with c ranging over $\mathcal{P}_\circ$), and we already established that these subgroups generate a free product with factors isomorphic to $F_\bullet$. ■

An action on the chamber set of Δ by automorphisms induces an action on the set of residues of a fixed type. The following proposition allows us to recover the original action in most situations.

Proposition 3.2.13. *Let Δ be a right-angled building of type I and let $J \subseteq I$. Assume that all local groups are nontrivial. Then the following are equivalent:*

(i) *$\mathcal{U}(\boldsymbol{F})$ acts faithfully on the set of J-residues;*

(ii) *J does not contain a connected component of (the underlying graph of) the diagram of Δ.*

Proof. First observe that if $\mathcal{U}$ acts faithfully on $\operatorname{Res}_J(\Delta)$, then it acts faithfully on $\operatorname{Res}_K(\Delta)$ as well, for every subset $K \subseteq J$. Hence, it is sufficient to show that the *minimal* subsets $J \subseteq I$ (with respect to inclusion) such that the action on $\operatorname{Res}_J(\Delta)$ is not faithful, are in one-to-one correspondence with the connected components of the diagram.

Assume that J is a connected component. If $J = I$, then the action on $\operatorname{Res}_J(\Delta)$ is the trivial action on a singleton and is not faithful. If $J \neq I$, then we have a direct product structure $\Delta \cong \mathcal{R} \times \mathcal{R}'$ where $\mathcal{R}$ is a residue of type J and $\mathcal{R}'$ a residue of type $I \setminus J$. By Lemma 3.2.1,

$$\mathcal{U}_\Delta(\boldsymbol{F}) \cong \mathcal{U}_{\mathcal{R}}\big(\boldsymbol{F}\big|_J\big) \times \mathcal{U}_{\mathcal{R}'}\big(\boldsymbol{F}\big|_{I \setminus J}\big).$$

Since the first factor is nontrivial and stabilises every J-residue, we see that $\mathcal{U}_\Delta(\boldsymbol{F})$ does not act faithfully on $\operatorname{Res}_J(\Delta)$.

Conversely, let J be a minimal subset of I such that the action on $\operatorname{Res}_J(\Delta)$ is not faithful. Then we certainly have $J \neq \emptyset$ (because the action on the chambers is faithful). If $J = I$, then the diagram is connected (because of the previous paragraph and the minimality assumption). Henceforth we can assume that there exist types $j \in J$ and $k \notin J$.

Let g be a nontrivial automorphism in $\mathcal{U}$ stabilising all residues of type J. Then by the minimality assumption, there exists some residue $\mathcal{R}$ of type $J \setminus \{j\}$ such that $g \,.\, \mathcal{R} \neq \mathcal{R}$. Let $\mathcal{P}$ be a k-panel intersecting $\mathcal{R}$ in a chamber c. As c and $g \,.\, c$ are contained in a common residue of type J, but not of type $J \setminus \{j\}$, the type of a minimal gallery from c to $g \,.\, c$ has to contain the element j. Moreover for every $d \in \mathcal{P}$, we have $\mathrm{proj}_{\mathcal{P}}(g \,.\, d) = d$, since d and $g \,.\, d$ are contained in a common J-residue as well. Hence $\mathcal{P}$ and $g \,.\, \mathcal{P}$ are parallel, and by Proposition 2.3.7 (ii), we conclude that $m_{jk} = 2$.

Since $j \in J$ and $k \notin J$ were arbitrary, this means that J is a union of connected components of the diagram. By minimality, J is a single connected component. ■

It is worth making the following corollary explicit.

Corollary 3.2.14. *Let Δ be an irreducible right-angled building over I and let $J \subsetneq I$. Then the action of $\mathcal{U}_\Delta(\boldsymbol{F})$ on $\mathrm{Res}_J(\Delta)$ is faithful.*

Proof. This follows immediately from Proposition 3.2.13. ■

Next, we characterise when the universal group acts primitively on $\mathrm{Res}_J(\Delta)$. Our result, Theorem 3.2.15, has been proved by Simon Smith in the rank two case in [Smi17, Theorem 26 (ii)]. Smith remarks that in this setting, there is a surprising similarity with wreath products of permutation groups; the similarity weakens in higher rank.

Theorem 3.2.15 (primitivity). *Let $J \subseteq I$. Then the action of $\mathcal{U}(\boldsymbol{F})$ on the set $\mathrm{Res}_J(\Delta)$ of residues of type J is primitive if and only if all of the following conditions hold:*

(i) *$|I \setminus J| = 1$, so that $I = J \sqcup \{k\}$ for some $k \in I$;*

(ii) *F_k is primitive and nonregular;*

(iii) *F_i is transitive for all $i \in I \setminus k^\perp$.*

Proof. We proceed in five steps.

Step 1. Condition (i) *is necessary.*

Assume that $J \subseteq J' \subseteq I$. Then the residues of type J' induce a partition of the residues of type J, which is preserved by the action of $\mathcal{U}$. In other words, an intermediate set $J \subseteq J' \subseteq I$ determines a block system of imprimitivity for the action on $\mathrm{Res}_J(\Delta)$. The singleton partition corresponds to $J = J'$, the trivial partition in one block to $J' = I$. Hence $\mathcal{U}$ can only act primitively if $|I \setminus J| \leq 1$.

In what follows, we may assume that $I = J \sqcup \{k\}$. Define the graph Γ with vertex set $\mathrm{Res}_J(\Delta)$, and where two J-residues are adjacent if and only if they contain two k-adjacent chambers. Note that $\mathcal{U}$ acts on Γ in a natural way by graph automorphisms. We will simply call J-residues *adjacent* when they are adjacent in Γ. Note that λ_k induces a well-defined vertex colouring on Γ.

Step 2. Condition (ii) (a) *is necessary: if F_k is imprimitive, then $\mathcal{U}$ acts imprimitively on* $\mathrm{Res}_J(\Delta)$.

Let $\approx$ be a nontrivial equivalence relation on Ω_k that is invariant under F_k. Let $\mathcal{P}$ be any k-panel. Then $\approx$ naturally lifts to an equivalence relation on the chambers of $\mathcal{P}$, by declaring $c_1 \approx c_2$ in $\mathcal{P}$ if and only if $\lambda_k(c_1) \approx \lambda_k(c_2)$ in Ω_k.

If $q_k = 2$, then as an imprimitive subgroup of $\mathrm{Sym}(2)$, the group F_k is trivial. By Proposition 3.2.2, the action on $\mathrm{Res}_J(\Delta)$ is intransitive and certainly imprimitive. If $q_k \geq 3$, then let $c_1 \approx c_2 \not\approx c_3$ be three chambers in $\mathcal{P}$, and for $\ell \in \{1, 2, 3\}$, let $\mathcal{R}_\ell$ be the residue of type J containing c_ℓ. We refer to Figure 3.1a.

Since $\{\mathcal{R}_1, \mathcal{R}_2\}$ is an edge in Γ and $\mathcal{U}$ acts by graph automorphisms, the orbital graph with respect to the orbital $\mathcal{U}\,.\,(\mathcal{R}_1, \mathcal{R}_2)$ is a subgraph of Γ. Since the equivalence relation is invariant under $\mathcal{U}$, it follows that all colours in a connected component of the orbital graph are contained in a single equivalence class of Ω_k. In particular, $\mathcal{R}_1$ and $\mathcal{R}_3$ belong to different connected components of Γ. The imprimitivity of $\mathcal{U}$ on $\mathrm{Res}_J(\Delta)$ now follows from Higman's theorem (Theorem 1.1.7).

Step 3. Condition (ii) (b) *is necessary: if F_k is regular, then $\mathcal{U}$ acts imprimitively on* $\mathrm{Res}_J(\Delta)$.

Assume that F_k acts regularly on the set of k-colours. Let $\mathcal{R}_0$ be a J-residue and let $g \in \mathcal{U}$. Define the k-colours $y = \lambda_k(\mathcal{R}_0)$ and $z = \lambda_k(g\,.\,\mathcal{R}_0)$. By regularity, there exists a *unique* element $f \in F_k$ such that $f\,.\,y = z$. We claim that $\lambda_k(g\,.\,\mathcal{R}) = f\,.\,\lambda_k(\mathcal{R})$ for each J-residue $\mathcal{R}$, using induction on the distance n to $\mathcal{R}_0$ in the graph Γ. For the case $n = 0$, there is nothing to show.

Next, consider a residue $\mathcal{R}$ at distance $n+1$ to $\mathcal{R}_0$ (in the graph Γ) and let $\mathcal{R}'$ be a neighbour of $\mathcal{R}$ at distance n. Let $x = \lambda_k(\mathcal{R})$ and $x' = \lambda_k(\mathcal{R}')$, and note that $\lambda_k(g\,.\,\mathcal{R}') = f\,.\,x'$ by the induction hypothesis. Let $\mathcal{P}$ be a k-panel containing k-adjacent chambers of $\mathcal{R}$ and $\mathcal{R}'$. The local action of g at $\mathcal{P}$ is a permutation of F_k taking x' to $f\,.\,x'$ — by regularity, this local action has to be equal to f and we may conclude that $\lambda_k(g\,.\,\mathcal{R}) = f\,.\,\lambda_k(\mathcal{R})$. Our claim follows by induction.

Since y was essentially arbitrary, we find that the partition of $\mathrm{Res}_J(\Delta)$ into sets

$$B(x) = \{\mathcal{R} \in \mathrm{Res}_J(\Delta) \mid \lambda_k(\mathcal{R}) = x\}$$

of identically coloured residues, is in fact a nontrivial block system of imprimitivity.

Step 4. Condition (iii) *is necessary.*

Suppose that F_i acts intransitively on Ω_i for some $i \in I$ with $m_{ik} = \infty$. By Lemma 2.5.6 we obtain a natural well-defined edge colouring of Γ with colours in Ω_i induced by the k-panels. Let $\mathcal{R}_1$ be any residue of type J. Pick two i-adjacent chambers $c \sim_i d$ in $\mathcal{R}_1$ such that $\lambda_i(c)$ and $\lambda_i(d)$ lie in distinct F_i-orbits. Let $\mathcal{R}_2$ and $\mathcal{R}_3$ be two J-residues containing a chamber k-adjacent to c and d, respectively. We refer to Figure 3.1b.

Consider the orbital graph with respect to the orbital $\mathcal{U}\,.\,(\mathcal{R}_1, \mathcal{R}_2)$, which is again a subgraph of Γ. The action of $\mathcal{U}$ on Γ preserves the orbits of the i-colours of edges under F_i. Hence the i-colours in a connected component of the orbital graph are contained in a single F_i-orbit — in particular it follows that $\mathcal{R}_3$ is not contained in the same component as $\mathcal{R}_1$ and $\mathcal{R}_2$. The imprimitivity of $\mathcal{U}$ now again follows from Higman's theorem (Theorem 1.1.7).

Step 5. Assume that conditions (i), (ii), (iii) *hold. Then $\mathcal{U}$ acts primitively on* $\mathrm{Res}_J(\Delta)$.

Let $\approx$ be a nontrivial $\mathcal{U}$-invariant equivalence relation on the set $\mathrm{Res}_J(\Delta)$ of J-residues; we will show that this relation is universal. Consider two equivalent residues $\mathcal{R}_0 \approx \mathcal{R}$, consider a shortest path from $\mathcal{R}_0$ to $\mathcal{R}$ in Γ and let $\mathcal{R}'$ be the J-residue adjacent to $\mathcal{R}$ on this shortest path. Let $c \in \mathcal{R}$ and $c' \in \mathcal{R}'$ be two k-adjacent chambers, and let $\mathcal{P}$ be the k-panel containing c and c'.

By Lemma 1.1.8, since F_k is primitive and nonregular, there is a permutation $f \in F_k$ fixing $\lambda_k(c')$ but not $\lambda_k(c)$. By Lemma 3.2.8, f extends to an automorphism $g \in \mathcal{U}$ fixing c' but not fixing c, and fixing all chambers d such that $\mathrm{proj}_{\mathcal{P}}(d) = c'$. In particular, g fixes c_0. Moreover, $g\,.\,c$ is k-adjacent to $g\,.\,c' = c'$. It follows from $\mathcal{U}$-invariance that $\mathcal{R} \approx \mathcal{R}_0 = g\,.\,\mathcal{R}_0 \approx g\,.\,\mathcal{R}$, so we have constructed two equivalent J-residues *that are adjacent in* Γ. For an illustration we refer to Figure 3.1c.

We now claim that all J-residues containing some chamber in $\mathcal{P}$ are equivalent. Indeed, consider the induced equivalence relation on the k-colours Ω_k where we define two colours to be equivalent if the J-residues of the corresponding chambers in $\mathcal{P}$ are equivalent. This equivalence relation on Ω_k is invariant under F_k by Lemma 3.2.7, hence universal by primitivity of F_k. Hence our claim follows, and in particular, we find that $\mathcal{R} \approx \mathcal{R}' \approx g \,.\, \mathcal{R}$.

Next, we claim that all J-residues adjacent to $\mathcal{R}$ are equivalent to $\mathcal{R}$. Let $\mathcal{R}''$ be such a J-residue adjacent to $\mathcal{R}$ and let $d \in \mathcal{R}$ be k-adjacent to a chamber in $\mathcal{R}''$. Let γ be a minimal gallery joining c and d in $\mathcal{R}$; we show that $\mathcal{R} \approx \mathcal{R}''$ by means of induction on the gallery length n of γ. The case $n = 0$ is precisely the conclusion of the previous paragraph. For a gallery γ of length $n+1$, let d' be the chamber on γ such that $\operatorname{dist}(c, d') = n$ and let $d' \sim_i d$. There are two options, illustrated in Figures 3.1d and 3.1e.

(a) If $m_{ik} = 2$, then all J-residues containing a chamber k-adjacent to d, contain a chamber k-adjacent to d' as well. Hence there is nothing to prove.

(b) If $m_{ik} = \infty$, then by transitivity of F_i there exists some permutation $f \in F_i$ mapping $\lambda_i(d')$ to $\lambda_i(d)$. Extend this permutation by Lemma 3.2.7 to an element $g \in \mathcal{U}$ mapping d' to d. By the induction hypothesis, all J-residues containing a chamber k-adjacent to d' are equivalent to $\mathcal{R}$, hence from $\mathcal{U}$-invariance, all J-residues containing a chamber k-adjacent to $g \,.\, d' = d$ are equivalent to $g \,.\, \mathcal{R} = \mathcal{R}$ as well.

We conclude by induction that all J-residues adjacent to $\mathcal{R}$ are in fact equivalent to $\mathcal{R}$. Repeating this argument and using the fact that Γ is connected, it follows that all J-residues are equivalent, so that $\approx$ is the universal relation. Hence $\mathcal{U}$ acts primitively on $\operatorname{Res}_J(\Delta)$.

This completes the proof. ∎

The following few properties all involve the subgroup $\mathcal{U}(\boldsymbol{F})^+$ generated by the chamber stabilisers,

$$\mathcal{U}_\Delta(\boldsymbol{F})^+ = \big\langle \mathcal{U}(\boldsymbol{F})_c \mid c \in \Delta \big\rangle.$$

Clearly this is a normal subgroup of the universal group, and hence a major obstruction for $\mathcal{U}(\boldsymbol{F})$ to be simple. Remark that $\mathcal{U}(\boldsymbol{F})^+$ is trivial if and only if every local group acts freely; it will take some more work to characterise when $\mathcal{U}(\boldsymbol{F})^+ = \mathcal{U}(\boldsymbol{F})$.

Lemma 3.2.16. *Let $\mathcal{P}$ be an i-panel. Then the local actions of $\mathcal{U}(\boldsymbol{F})^+$ on $\mathcal{P}$ are contained in $(F_i)^+$.*

Proof. Let $g \in \mathcal{U}(\boldsymbol{F})$ be an automorphism that stabilises some chamber $c \in \Delta$. It suffices to show that the local action $\sigma(g, \mathcal{P})$ is a permutation in $(F_i)^+$. We shall use induction on $\operatorname{dist}(c, \mathcal{P})$. First, if $c \in \mathcal{P}$, then $\sigma(g, \mathcal{P})$ indeed stabilises the colour $\lambda_i(c)$. If otherwise $\operatorname{dist}(c, \mathcal{P}) = n+1$, then let $d = \operatorname{proj}_{\mathcal{P}}(c)$, let $d' \sim_j d$ (for some $j \neq i$) be such that $\operatorname{dist}(c, d') = n$, and let $\mathcal{P}'$ be the i-panel containing d'. Then we know that $\lambda_i(d') = \lambda_i(d)$ and that

$$\sigma(g, \mathcal{P}') \,.\, \lambda_i(d') = \lambda_i(g \,.\, d') = \lambda_i(g \,.\, d) = \sigma(g, \mathcal{P}) \,.\, \lambda_i(d).$$

This implies that $\sigma(g, \mathcal{P}) = \sigma(g, \mathcal{P}') \cdot f$ for some permutation $f \in F_i$ stabilising the colour $\lambda_i(d)$. Since $\sigma(g, \mathcal{P}') \in (F_i)^+$ by the induction hypothesis, our conclusion follows. ∎

Proposition 3.2.17. *Let $\mathbf{Id}, \boldsymbol{F}^+, \boldsymbol{F}$ be the collections of subgroups $\{1\} \leq (F_i)^+ \leq F_i$ of the symmetric group $\operatorname{Sym}(\Omega_i)$, indexed by $i \in I$. Then $\mathcal{U}(\boldsymbol{F}^+) = \langle \mathcal{U}(\mathbf{Id}), \mathcal{U}(\boldsymbol{F})^+ \rangle$.*

Proof. The claim that $\mathcal{U}(\boldsymbol{F}^+)$ contains both $\mathcal{U}(\mathbf{Id})$ and $\mathcal{U}(\boldsymbol{F})^+$ as subgroups follows immediately from Lemmas 3.1.5 and 3.2.16, respectively. Conversely, let g be an automorphism in $\mathcal{U}(\boldsymbol{F}^+)$. Let c be an arbitrary chamber, let $i \in I$, and let $f_i = \sigma_\lambda(g, \mathcal{P}_i(c)) \in (F_i)^+$. Now use Proposition 3.2.9 to construct an automorphism $h_i \in \mathcal{U}(\boldsymbol{F})^+$ stabilising $\mathcal{P}_i(c)$ with local action f_i. By construction,

$$\lambda_i(h_i \,.\, c) = \sigma_\lambda(h_i, \mathcal{P}_i(c)) \,.\, \lambda_i(c) = f_i \,.\, \lambda_i(c) = \sigma_\lambda(g, \mathcal{P}_i(c)) = \lambda_i(g \,.\, c),$$

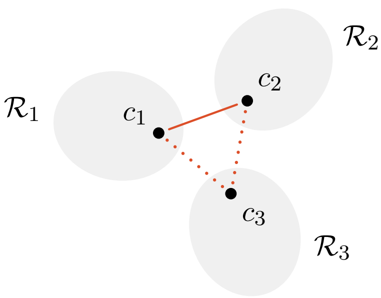

(a) The configuration in step 2.

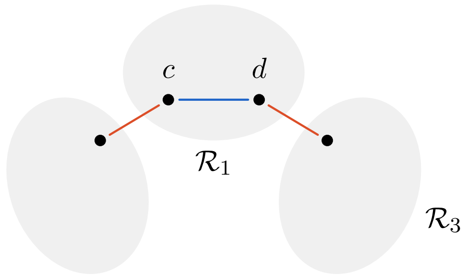

(b) The configuration in step 4.

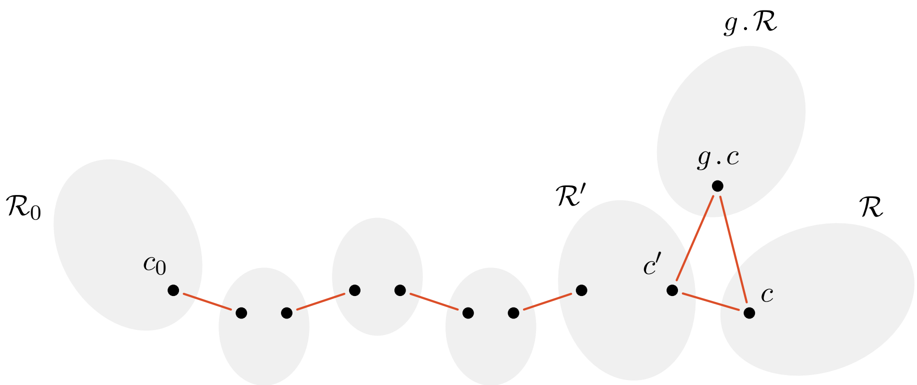

(c) The configuration in step 5.

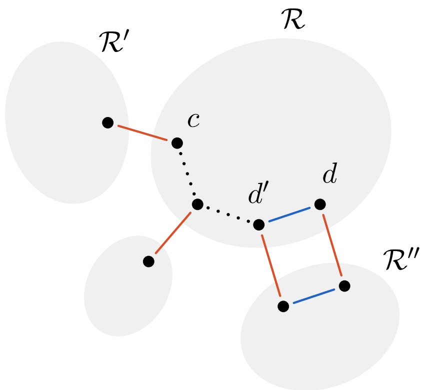

(d) The configuration in step 5, first case.

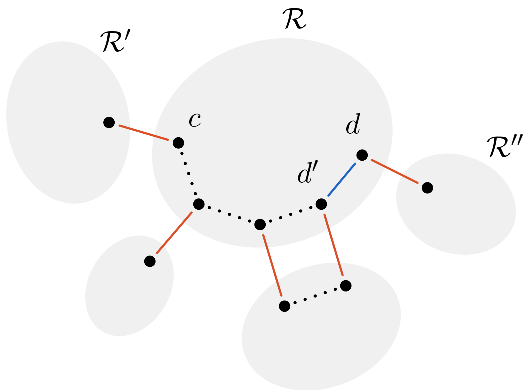

(e) The configuration in step 5, second case.

Figure 3.1. The configurations in the proof of Theorem 3.2.15.

while $\lambda_j(h_i\,.\,c) = \lambda_j(c)$ for all $j \neq i$. Repeating this with $h_i\,.\,c$ for every type in I, we eventually obtain an automorphism

$$h = \prod_{i \in I} h_i \in \mathcal{U}(\boldsymbol{F})^+$$

such that $\lambda_i(h\,.\,c) = \lambda_i(g\,.\,c)$ for every $i \in I$. By Proposition 3.2.2, we can find an automorphism $h' \in \mathcal{U}(\mathbf{Id})$ such that $h'\,.\,(h\,.\,c) = g\,.\,c$. It follows that $(h' \cdot h)^{-1} \cdot g \in \mathcal{U}(\boldsymbol{F})^+$ and hence

$$g \in \mathcal{U}(\mathbf{Id}) \cdot \mathcal{U}(\boldsymbol{F})^+. \qquad \blacksquare$$

The following proposition relates $\mathcal{U}(\boldsymbol{F})^+$ to the notion of implosions from Section 2.6.

Proposition 3.2.18. *For every $i \in I$, define an equivalence relation $\equiv_i$ on Ω_i by declaring i-colours to be equivalent if and only if they are contained in the same orbit of the local group F_i. Let (τ, Δ') be an implosion of Δ with respect to $\equiv_i$. Then the group $\mathcal{U}_\Delta(\boldsymbol{F})^+$ stabilises the fibres of τ.*

In other words, the following diagram commutes.

$$\begin{tikzcd} \Delta \arrow[r, "\mathcal{U}_\Delta(\boldsymbol{F})^+"] \arrow[d, "\tau"] & \Delta \arrow[d, "\tau"] \\ \Delta' \arrow[r, equal] & \Delta' \end{tikzcd}$$

Proof. Note that the implosion map τ identifies harmonious chambers in panels.

Let $g \in \mathcal{U}(\boldsymbol{F})$ stabilise some chamber $c \in \Delta$. Consider any chamber d together with its image $g\,.\,d$. We show that $\tau(d) = \tau(g\,.\,d)$ by induction on $\operatorname{dist}(c, d)$. If $c = d$, then there is nothing to show, since $g\,.\,c = c$. Otherwise if $\operatorname{dist}(c, d) = n + 1$, let d' be a chamber such that $d' \sim_i d$ (for some i) and $\operatorname{dist}(c, d') = n$. We distinguish between two cases.

(a) If $\lambda_i(d) \equiv_i \lambda_i(d')$, or in other words if d and d' are harmonious, then we have $\tau(d) = \tau(d')$ by construction of τ. Moreover, $\tau(d') = \tau(g\,.\,d')$ by the induction hypothesis. Finally, $g\,.\,d'$ and $g\,.\,d$ are again harmonious, hence $\tau(g\,.\,d') = \tau(g\,.\,d)$ by construction. Putting everything together we conclude that $\tau(d) = \tau(g\,.\,d)$.

(b) If $\lambda_i(d') \not\equiv_i \lambda_i(d)$, then by construction,

  - $\tau(d)$ is the unique chamber in Δ' that is i-adjacent to $\tau(d')$ and that has i-colour equal to the orbit $F_i\,.\,\lambda_i(d)$, and
  - $\tau(g\,.\,d)$ is the unique chamber in Δ' that is i-adjacent to $\tau(g\,.\,d')$ and that has i-colour equal to the orbit of $F_i\,.\,\lambda_i(g\,.\,d)$.

  Since $\tau(d') = \tau(g\,.\,d')$ by the induction hypothesis and since $\lambda_i(d)$ and $\lambda_i(g\,.\,d)$ clearly lie in the same orbit of F_i we conclude again that $\tau(d) = \tau(g\,.\,d)$.

By induction, the fibres of τ are stabilised by the stabiliser $\mathcal{U}_c$ and hence by $\mathcal{U}^+$ as well. $\blacksquare$

Remark 3.2.19. Note that the equivalence relations $\equiv_i$ uniquely determine the building Δ' but not the map τ from Proposition 2.6.2. For uniqueness of τ we additionally need to fix a centre in Δ and a compatible image in Δ'. Up to this choice, we then have a well-defined canonical morphism

$$\psi\colon\ \operatorname{Aut}(\Delta) \to \operatorname{Aut}(\Delta')$$

defined by $\psi(g) \,.\, \tau(d) = \tau(g \,.\, d)$ for every $g \in \mathrm{Aut}(\Delta)$ and $d \in \Delta$.

$$\begin{tikzcd} \Delta \arrow[r, "\mathrm{Aut}(\Delta)"] \arrow[d, "\tau"] \arrow[dr, phantom, "\Downarrow\psi"] & \Delta \arrow[d, "\tau"] \\ \Delta' \arrow[r, "\mathrm{Aut}(\Delta')"'] & \Delta' \end{tikzcd}$$

Observe that the kernel of ψ contains precisely those automorphisms that stabilise all fibres of τ. In other words, Proposition 3.2.18 can be rephrased as the statement that $\mathcal{U}_\Delta(\boldsymbol{F})^+ \leq \ker(\psi)$.

On the other hand, note that in general $\mathcal{U}_\Delta(\boldsymbol{F})^+ \neq \ker(\psi)$. Indeed, assume as an extreme case that the local groups are regular. Then by Lemma 3.2.16 we have that $\mathcal{U}_\Delta(\boldsymbol{F})^+$ is trivial, whereas the implosion map collapses Δ to a single chamber. Thus in this case the kernel of ψ is the full automorphism group of Δ.

In the next result we characterise when the universal group is generated by its chamber stabilisers. As it turns out, this not only depends on the local groups, but on the combinatorial structure of the diagram as well. Recall that a *vertex cover* of a graph is some set of vertices that includes at least one endpoint of every edge. Hence, in our setting, a vertex cover of a right-angled diagram over I is a subset $J \subseteq I$ such that whenever $m_{ij} = \infty$, we have $i \in J$ or $j \in J$ (or both).

Theorem 3.2.20. *Assume that the diagram of Δ has no isolated nodes. The following are equivalent.*

(i) *$\mathcal{U}(\boldsymbol{F})^+ = \mathcal{U}(\boldsymbol{F})$;*

(ii) *$\mathcal{U}(\boldsymbol{F})^+$ has finite index in $\mathcal{U}(\boldsymbol{F})$;*

(iii) *the local groups are generated by point stabilisers for every $i \in I$, and are transitive for every i in some vertex cover of (the underlying graph of) the diagram of Δ.*

Proof. The implication (i) $\Rightarrow$ (ii) is of course trivial. For (ii) $\Rightarrow$ (iii), we proceed in two steps.

For transitivity, we use the implosion map $\tau\colon \Delta \to \Delta'$ from Proposition 3.2.18. Suppose by means of contraposition that the indices of the *transitive* local groups do *not* define a vertex cover, i.e. there are intransitive local groups F_i and F_j such that $m_{ij} = \infty$. Then in particular Δ' is not spherical.

Let $\blacktriangleleft$ and $\triangleleft$ be i-colours in different F_i-orbits and $\blacktriangleright$ and $\triangleright$ j-colours in different F_j-orbits. Consider an apartment $\mathcal{A}$ of a residue of type $\{i, j\}$ with colours only in $\{\blacktriangleleft, \triangleleft\}$ and $\{\blacktriangleright, \triangleright\}$, so that $\mathcal{A}$ looks as follows.

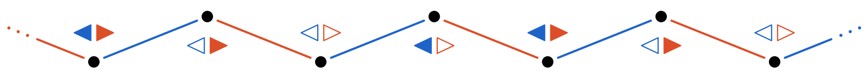

By Proposition 3.2.2, two chambers $c, d \in \mathcal{A}$ lie in the same $\mathcal{U}$-orbit if and only if $4 \mid \mathrm{dist}(c, d)$. Meanwhile τ maps adjacent chambers in $\mathcal{A}$ to *distinct* adjacent chambers in a residue of type $\{i, j\}$ in Δ'. By Proposition 3.2.18 it then follows that no two chambers of $\mathcal{A}$ lie in the same $\mathcal{U}^+$-orbit. We found a $\mathcal{U}$-orbit that is the union of infinitely many $\mathcal{U}^+$-orbits — hence $\mathcal{U}^+$ cannot have finite index in $\mathcal{U}$.

For the claim that all local groups are generated by point stabilisers, fix an index $i \in I$ and assume, again by means of contraposition, that F_i is not generated by point stabilisers. In other words, the F_i-orbits and of $(F_i)^+$ do not agree. Explicitly, there is some orbit $X \in \Omega_i / F_i$ that is the union of at least two $(F_i)^+$-orbits.

Define the new local data $\boldsymbol{F'}$ by replacing F_j with the trivial permutation group of degree q_j for every $i \neq j$, keeping F_i unaltered. Let $H = \mathcal{U}(\boldsymbol{F'}) \leq \mathcal{U}(\boldsymbol{F})$ and set $H^+ = \mathcal{U}(\boldsymbol{F})^+ \cap H$. Since

$$\left[H : H^+\right] \leq \left[\mathcal{U}(\boldsymbol{F}) : \mathcal{U}(\boldsymbol{F})^+\right],$$

it suffices to show that H^+ has infinite index in H.

Consider the induced action of H on the i-tree-wall tree $\Gamma = \Gamma_i$. Since the diagram has no isolated nodes, Γ is an infinite tree. Interpret Γ as a rank two right-angled building over $\{\bullet, \circ\}$. Then, up to a potential nontrivial kernel, H acts on Γ as a universal group where the local groups are $F_\bullet = F_i$ and $F_\circ$ is trivial. Note, if $g \in H$ stabilises a chamber of Δ, then it stabilises a chamber of Γ as well. The nontrivial local actions of H^+ on Γ_i are hence contained in $(F_i)^+$ by Lemma 3.2.16.

Define an equivalence relation $\equiv_\bullet$ by declaring every two $\bullet$-colours (i.e. i-colours) in the same $(F_i)^+$-orbit to be equivalent. Let $\equiv_\circ$ be the identity relation. Since $(F_i)^+$ is assumed to have at least two orbits, the corresponding implosion Γ' of Γ is again an infinite tree. The induced action of H^+ on Γ' is trivial (since all local actions are trivial). The action of H on Γ' however has infinite orbits (since H acts coboundedly by Corollary 3.2.5 and the implosion map is nonexpansive). Hence H^+ cannot have finite index in H.

Finally, for the implication (iii) $\Rightarrow$ (i), it is sufficient to show that $\mathcal{U}$ and $\mathcal{U}^+$ have identical orbits. Let $c, d \in \Delta$ be two harmonious chambers. We construct an automorphism in $\mathcal{U}^+$ taking c to d.

First assume that c and d are contained in some residue $\mathcal{R}$ of *spherical* type $J \subseteq I$. Recall that

$$\mathcal{R} \cong \prod_{j \in J} \mathcal{P}_j(c).$$

For every $j \in J$ there exists some permutation $f_j \in F_j$ such that $f_j \,.\, \lambda_j(c) = \lambda_j(d)$. By assumption f_j can be written as some product of permutations in F_j, each of which fixes some j-colour. Extend these permutations using Lemma 3.2.7 to automorphisms in $\mathcal{U}$, each of which fixes some chamber j-adjacent to c, and take their product over all $j \in J$ (in arbitrary order). We find an automorphism in $\mathcal{U}^+$ that stabilises $\mathcal{R}$ and maps c to a chamber with J-colours equal to those of d, i.e. to d itself.

For the general case, we use induction on the distance $n = \operatorname{dist}(c, d)$. The case $n = 0$ is trivial, and the previous paragraph settles the case $n = 1$. For $n > 1$, let γ be a minimal gallery joining c to d. We distinguish between two cases, depending on the number of pairs of consecutive harmonious chambers on γ.

(a) If there are no such pairs, then let J_γ be the set of all types occurring in γ. For every $j \in J_\gamma$ the local group F_j has at least two orbits. By the vertex cover assumption, any two elements of J_γ commute in the Weyl group of Δ. In other words, the residue of type J_γ containing γ is spherical, and the conclusion follows from the previous paragraph.

(b) If there exists at least one such pair (c', d'), then let $g \in \mathcal{U}^+$ be such that $g \,.\, c' = d'$. Note that

$$\operatorname{dist}(c, g^{-1} \,.\, d) \leq \operatorname{dist}(c, c') + \operatorname{dist}(c', g^{-1} \,.\, d) \leq \operatorname{dist}(c, c') + \operatorname{dist}(d', d) = n - 1.$$

By the induction hypothesis, there exists an element $h \in \mathcal{U}^+$ such that $h \,.\, c = g^{-1} \,.\, d$. Then $gh \in \mathcal{U}^+$ satisfies $gh \,.\, c = d$.

This finishes the proof. ∎

3.3 Topological properties

In order to make $\mathcal{U}_\Delta(\boldsymbol{F})$ into a *topological* group, we will endow it with the permutation topology from Section 1.2.3. Recall that this topology is obtained by taking the pointwise stabilisers of finite sets of chambers as an identity neighbourhood basis.

First, let us characterise when the universal group is unexciting from a topological point of view.

Proposition 3.3.1 (discreteness). *Assume that the diagram has no isolated nodes. Then the following are equivalent:*

(i) *the local group F_i acts freely on Ω_i for every $i \in I$;*

(ii) *the universal group $\mathcal{U}(\boldsymbol{F})$ acts freely on Δ;*

(iii) *the universal group $\mathcal{U}(\boldsymbol{F})$ is a discrete topological group.*

Proof. For the implication (i) $\Rightarrow$ (ii), suppose all of the local groups act freely and that $g \in \mathcal{U}(\boldsymbol{F})$ fixes some chamber c. Then for every $i \in I$, the local action of g at the i-panel containing c is a permutation in F_i with a fixed point, hence is the identity. It follows that every panel containing c is fixed by g. Because a building is connected, g is the identity.

The implication (ii) $\Rightarrow$ (iii) is an immediate consequence of Corollary 1.2.26.

We prove the final implication (iii) $\Rightarrow$ (i) by contraposition. Assume some local group F_i does not act freely on Ω_i and let f be a nontrivial permutation in F_i fixing some i-colour x. Let C be any finite set of chambers. By Proposition 2.5.5, there is an i-panel $\mathcal{P}$ such that $\mathrm{proj}_{\mathcal{P}}(C) = \{c\}$ with $\lambda_i(c) = x$. By Lemma 3.2.8, the permutation f of the colours of $\mathcal{P}$ extends to a nontrivial element of the universal group that stabilises C pointwise. In conclusion, no pointwise stabiliser of a finite set of chambers is trivial. ∎

Of course, the local permutation groups can be endowed with the permutation topology as well, and it should not come as a surprise that there are strong connections between topological properties of the local groups and of the full universal group.

The proof of the following lemma is completely similar to the proof of [Smi17, Lemma 7] for trees.

Lemma 3.3.2. *Let $\mathcal{P}$ be a panel of type i. Then the map*

$$\sigma_\lambda(\bullet, \mathcal{P})\colon \operatorname{Aut}(\Delta) \to \operatorname{Sym}(\Omega_i)\colon g \mapsto \sigma_\lambda(g, \mathcal{P})$$

is continuous w.r.t. the permutation topologies on $\operatorname{Aut}(\Delta)$ *and* $\operatorname{Sym}(\Omega_i)$.

Proof. Abbreviate $\sigma_\lambda(\bullet, \mathcal{P})$ to σ for readability. By Lemma 3.2.7, σ is surjective. Let $V \subseteq \operatorname{Sym}(\Omega_i)$ be an open subset and consider any g in the preimage $\sigma^{-1}(V) \subseteq \operatorname{Aut}(\Delta)$. Then $\sigma(g)$ is contained in some open neighbourhood in V, i.e. in the coset $\sigma(g) \cdot (\operatorname{Sym}\Omega_i)_{(X)}$ of the pointwise stabiliser of some finite colour set $X \subseteq \Omega_i$. Consider the finite set $C = \{c \in \mathcal{P} \mid \lambda_i(c) \in X\}$. Note that for all $h \in (\operatorname{Aut}\Delta)_{(C)}$ we have

$$\sigma_\lambda(gh, \mathcal{P}) = \sigma_\lambda(g, h\,.\,\mathcal{P}) \cdot \sigma_\lambda(h, \mathcal{P})$$

where h stabilises the panel $\mathcal{P}$ and $\sigma_\lambda(h, \mathcal{P})$ fixes the set X. It follows that

$$\sigma\big(g \cdot (\operatorname{Aut}\Delta)_{(C)}\big) \subseteq \sigma(g) \cdot (\operatorname{Sym}\Omega_i)_{(X)} \subseteq V.$$

Hence, g is contained in the open neighbourhood

$$g \cdot (\operatorname{Aut}\Delta)_{(C)} \subseteq \sigma^{-1}(V).$$

Since g was arbitrary, the preimage $\sigma^{-1}(V)$ is open. ∎

Proposition 3.3.3. *The following are equivalent:*

(i) *the local group F_i is closed in $\mathrm{Sym}(\Omega_i)$ for every $i \in I$;*

(ii) *the universal group $\mathcal{U}(\boldsymbol{F})$ is closed in $\mathrm{Aut}(\Delta)$.*

Proof. First we establish the implication (i) ⇒ (ii). By Lemma 3.3.2 we have for every panel $\mathcal{P}$ of Δ a continuous map

$$\sigma_{\mathcal{P}} \colon \mathrm{Aut}(\Delta) \to \mathrm{Sym}(\Omega_i) \colon g \mapsto \sigma_\lambda(g, \mathcal{P})$$

(where i is the type of $\mathcal{P}$). It now suffices to observe that

$$\mathcal{U}(\boldsymbol{F}) = \bigcap_{\mathcal{P}} \sigma_{\mathcal{P}}^{-1}(F_i).$$

We prove (ii) ⇒ (i) by contraposition, so assume that some local group F_i is not closed. Recall that the permutation topology agrees with the topology of pointwise convergence, and let $(f_n)_{n\in\mathbb{I}}$ be a net of permutations in F_i such that

$$f_n \to f \in \mathrm{Sym}(\Omega_i) \setminus F_i.$$

For every $x \in \Omega_i$ let $m(x) \in \mathbb{I}$ be an index such that $f_n . x = f . x$ for every $n > m(x)$.

Now let $\mathcal{P}$ be an i-panel. For every $n \in \mathbb{I}$, let $g_n \in \mathcal{U}(\boldsymbol{F})$ be the automorphism as in Lemma 3.2.7, stabilising $\mathcal{P}$ with local action f_n. Then the net $(g_n)_{n\in\mathbb{I}}$ converges to an automorphism g in $\mathrm{Aut}(\Delta)$. Indeed, for any chamber $c \in \Delta$, define

$$m(c) = \min\bigl\{ m(\lambda_i(d)) \bigm| d \text{ is a chamber on a minimal gallery from } c \text{ to } \mathrm{proj}_{\mathcal{P}}(c) \bigr\}$$

and $g_n . c$ remains constant for all $n > m(c)$. It is clear that g is an automorphism, that g stabilises the panel $\mathcal{P}$ and that the local action $\sigma_\lambda(g, \mathcal{P}) = f$, hence

$$g_n \to g \in \mathrm{Aut}(\Delta) \setminus \mathcal{U}(\boldsymbol{F}).$$

This shows that the universal group is not closed in $\mathrm{Aut}(\Delta)$. ∎

We can characterise when the univeral group is locally compact.

Proposition 3.3.4 (local compactness). *Assume that all F_i are closed in $\mathrm{Sym}(X_i)$. Then the following are equivalent:*

(i) *every suborbit of F_i is finite, for every $i \in I$;*

(ii) *every point stabiliser in F_i is compact, for every $i \in I$;*

(iii) *every chamber stabiliser in $\mathcal{U}(\boldsymbol{F})$ is compact;*

(iv) *the universal group $\mathcal{U}(\boldsymbol{F})$ is a locally compact topological group.*

Proof. Note that $\mathcal{U}(\boldsymbol{F})$ is closed by Proposition 3.3.3. Its chamber stabilisers are open, hence closed. The equivalence (i) ⇔ (ii) is Proposition 1.2.30.

For (ii) ⇒ (iii), take an arbitrary chamber $c \in \Delta$. By Proposition 1.2.30 it suffices to show that all $\mathcal{U}_c$-orbits are finite. Take a second chamber d and let

$$c = c_0 \sim c_1 \sim c_2 \sim \cdots \sim c_{k-1} \sim c_k = d$$

be a minimal gallery from c to d. Then the cardinality of the orbit $\mathcal{U}_c . d$ is at most

$$\bigl|\mathcal{U}_{c_0} . c_1\bigr| \cdot \bigl|\mathcal{U}_{(c_0, c_1)} . c_2\bigr| \cdots \bigl|\mathcal{U}_{(c_0, \ldots, c_{k-1})} . c_k\bigr| \quad \leq \quad \bigl|\mathcal{U}_{c_0} . c_1\bigr| \cdot \bigl|\mathcal{U}_{c_1} . c_2\bigr| \cdots \bigl|\mathcal{U}_{c_{k-1}} . c_k\bigr|.$$

Note that each orbit in the latter product is contained in some panel. The permutation isomorphism from Proposition 3.2.9 yields a bijection to an orbit of a point stabiliser of a local group, i.e.

$$\mathcal{U}_{c_{\ell-1}} . c_\ell \quad \leftrightarrow \quad (F_i)_{(\lambda_i(c_{\ell-1}))} . \lambda_i(c_\ell)$$

for each $0 < \ell \leq k$ (where i is the type such that $c_{\ell-1} \sim_i c_\ell$). Since the local groups are compact, all these orbits are finite, hence so is the cardinality of $\mathcal{U}_c . d$.

The implication (iii) $\Rightarrow$ (iv) is obvious.

We conclude with the implication (iv) $\Rightarrow$ (ii). By local compactness, there exists some finite set C of chambers whose pointwise stabiliser in $\mathcal{U}$ is compact. By Proposition 1.2.30, every $\mathcal{U}_{(C)}$-orbit is finite. Let $i \in I$ and let $x \in \Omega_i$ be any i-colour. Then by Proposition 2.5.5, there is an i-panel $\mathcal{P}$ such that the projection $\operatorname{proj}_{\mathcal{P}}(C)$ is a single chamber with i-colour x. It follows that Ω_i cannot contain infinite $(F_i)_x$-orbits — otherwise by Lemma 3.2.7, the corresponding chambers in $\mathcal{P}$ would lie in an infinite $\mathcal{U}_{(C)}$-orbit. Thus, every point stabiliser of every local group is compact. ■

Proposition 3.3.5 (compact generation of panel stabilisers). *Assume $\mathcal{U}(\boldsymbol{F})$ is locally compact. Then, for every $i \in I$, the following are equivalent:*

(i) *the local group F_i is compactly generated;*

(ii) *the setwise stabiliser in $\mathcal{U}(\boldsymbol{F})$ of an i-panel is compactly generated.*

Proof. By Proposition 3.2.9 and Lemma 3.3.2, we have an isomorphism

$$\mathcal{U}_{\{\mathcal{P}\}} \,/\, \mathcal{U}_{(\mathcal{P})} \cong F_i$$

of topological groups. Moreover, chamber stabilisers in $\mathcal{U}$ are compact by Proposition 3.3.4, hence $\mathcal{U}_{(\mathcal{P})}$ is compact as well. The result now follows from the observation that any Hausdorff quotient of a compactly generated group is itself compactly generated, and conversely, that any extension of a compactly generated group by a compact group is itself compactly generated. ■

Characterising when the full universal group $\mathcal{U}(\boldsymbol{F})$ is compactly generated, turns out to be quite a bit harder. We still make no assumptions regarding finiteness or transitivity of the local groups, though we do assume $\mathcal{U}(\boldsymbol{F})$ to be closed and locally compact. We are grateful to Pierre-Emmanuel Caprace for discussing the problem with us and for helping establish the rank two case.

First, a necessary condition.

Theorem 3.3.6. *Assume that Δ is irreducible and $\mathcal{U}_\Delta(\boldsymbol{F})$ is closed, locally compact, and compactly generated. Then every local group F_i has only finitely many orbits.*

Proof. Pick any chamber $c \in \Delta$. By Proposition 3.3.4, the chamber stabiliser $\mathcal{U}_c$ is a compact open subgroup. Let $T = \{t_1, \dots, t_n\}$ be a finite set as in Lemma 1.2.22, so that T together with $\mathcal{U}_c$ is a good generating set. By Lemma 2.5.8, the set $\{c, t_1 . c, \dots, t_n . c\}$ is contained in a finite convex subset $B \subseteq \Delta$.

First, we claim that the set $\mathcal{U} . B$ is connected. Let $g \in \mathcal{U}$ and $d \in B$ be arbitrary; we show that $g . d$ is connected to c by some gallery contained in $\mathcal{U} . B$. Let $g = t_k \cdots t_1 \cdot s$ with $t_1, \dots, t_k \in T$ and $s \in \mathcal{U}_c$. We use induction on k. For $k = 0$, we simply have that $g \in \mathcal{U}_c$ so that $g . d$ is connected to $g . c = c$ by a gallery in $g . B$. For $k \geq 1$, write $g = t_k \cdot g'$. By the induction hypothesis, $g' . d$ is connected to c by a gallery in $\mathcal{U} . B$. Consequently $g . d$ is connected to $t_k . c$ by a gallery in $\mathcal{U} . B$ and we can concatenate this gallery with one from $t_k . c$ to c in B. We conclude that the set $\mathcal{U} . B$ is indeed connected.

Next, we claim that every F_i-orbit has a "representative chamber" in B. More precisely, we claim that for each $i \in I$, the map

$$\widetilde{\lambda}_i \colon B \to \Omega_i/F_i \colon b \mapsto F_i \,.\, \lambda_i(b)$$

is surjective. Suppose by means of contradiction that some F_j-orbit Y is not in the image of $\widetilde{\lambda}_j$ for some $j \in I$. Call a chamber *neglected* if its j-colour is contained in the orbit Y. By assumption, no chamber in B is neglected, and Proposition 3.2.2 implies that no chamber in $\mathcal{U} \,.\, B$ is neglected. Let Γ be the set of all chambers of Δ connected to $\mathcal{U} \,.\, B$ by some gallery that does not pass through any neglected chamber. Then by Lemma 2.5.7, Γ is a subbuilding of Δ and hence convex. Moreover Γ is $\mathcal{U}$-invariant by construction. This contradicts the minimality of the action (Corollary 3.2.6).

In conclusion, for every $i \in I$ we have a surjective map $\widetilde{\lambda}_i$ from a finite set B onto the set of orbits of i-colours. Hence the local groups F_i have finitely many orbits. ■

Together with an additional assumption on the local groups, having finitely many orbits is a sufficient condition for the universal group to be compactly generated.

Theorem 3.3.7. *Assume that $\mathcal{U}(\boldsymbol{F})$ is closed and locally compact. Moreover, assume that every local group F_i is compactly generated and has finitely many orbits. Then $\mathcal{U}(\boldsymbol{F})$ is compactly generated.*

Proof. We construct an explicit compact generating set Q. For every $i \in I$, choose a transversal Υ_i for the action of $F_i \leq \operatorname{Sym}(\Omega_i)$. Every Υ_i is finite by assumption. Let $c \in \Delta$ be some chamber such that $\lambda_i(c) \in \Upsilon_i$ for all $i \in I$. Let $n = |I|$ be the rank of Δ and define the ball

$$\widetilde{B} = \{d \in \Delta \mid \operatorname{dist}(c, d) \leq n\}$$

and the *finite* subset

$$B = \{d \in \widetilde{B} \mid \lambda_i(d) \in \Upsilon_i \text{ for all } i \in I\}.$$

Note that every possible colour combination with colours in Υ_i occurs inside B. Hence $\mathcal{U} \,.\, B = \Delta$ by Proposition 3.2.2. Next, define

$$\widetilde{D} = \{d \in \Delta \mid \operatorname{dist}(c, d) = n + 1 \text{ and } d \text{ is adjacent to a chamber in } B\}$$

and the *finite* subset

$$D = \{d \in \widetilde{D} \mid \lambda_i(d) \in \Upsilon_i \text{ for all } i \in I\}.$$

For each of the finitely many pairs $(b, d) \in B \times (B \cup D)$ of identically coloured chambers, pick an element $t_{(b,d)} \in \mathcal{U}$ such that $t \,.\, b = d$. Define the set $T = \{t_{(b,d)} \mid (b, d) \in B \times (B \cup D)\}$ and note that $D \subseteq T \,.\, B$.

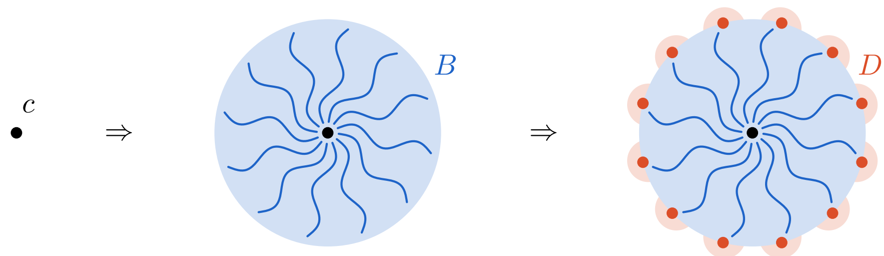

Figure 3.2. An impression of the construction in the proof of Theorem 3.3.7.

The local groups are assumed to be compactly generated, hence by Proposition 3.3.5, the (setwise) panel stabilisers are compactly generated as well. Let $S_{\mathcal{P}}$ be a compact generating set for $\mathcal{U}_{\{\mathcal{P}\}}$ for every panel $\mathcal{P}$. Define the set S as the union of all $S_{\mathcal{P}}$ over all panels $\mathcal{P}$ containing some chamber in B. As a finite union of compact sets, S is compact.

Finally, define $Q = S \cup T \cup \mathcal{U}_c$ and observe that Q is compact. We proceed in four steps to show that $\mathcal{U}(\boldsymbol{F}) = \langle Q \rangle = \langle S, T, \mathcal{U}_c \rangle$.

Step 1. For each chamber $d \in \widetilde{B}$, we have $d \in \langle S \rangle \,.\, B$.

We show this by induction on $\operatorname{dist}(c, d)$. If $d = c$ there is nothing to show. Otherwise let $d' \sim_i d$ be such that $\operatorname{dist}(c, d') = \operatorname{dist}(c, d) - 1$. Since $d' \in \widetilde{B}$, we can write $d' = h \,.\, b'$ for some $h \in \langle S \rangle$ and $b' \in B$. Now let $\mathcal{P}$ be the i-panel that contains b' and let b be the unique chamber of $\mathcal{P}$ such that $\lambda_i(b)$ is the representative of the orbit of $\lambda_i(h^{-1} \,.\, d)$ in the transversal Υ_i. Then there is some $s \in \mathcal{U}_{\{\mathcal{P}\}}$ such that $s \,.\, b = h^{-1} \,.\, d$, and we conclude that indeed $d = hs \,.\, b \in \langle S \rangle \,.\, B$.

Step 2. For each chamber $d \in \widetilde{D}$, we have $d \in \langle S, T \rangle \,.\, B$.

By definition of $\widetilde{D}$, the chamber d is i-adjacent to some chamber in B (for some $i \in I$). Let $\mathcal{P}$ be the i-panel containing d. Let d' be the unique chamber in $\mathcal{P}$ such that $\lambda_i(d')$ is the representative of the orbit of $\lambda_i(d)$ in Υ_i. There are two possible cases. If $d' \in B$, then we can immediately set $b = d'$ and $t = 1$. Otherwise $d' \in D$, and we can find some $b \in B$ and $t \in T$ sending b to d'. Moreover, there exists some $s \in \mathcal{U}_{\{\mathcal{P}\}}$ sending d' to d, and we conclude that $d = st \,.\, b \in \langle S, T \rangle \,.\, B$.

Step 3. For each chamber $d \in \Delta$, we have $d \in \langle S, T \rangle \,.\, B$.

We again use induction on $\operatorname{dist}(c, d)$. If $d = c$ there is nothing to show. Otherwise let $d' \sim_i d$ be such that $\operatorname{dist}(c, d') = \operatorname{dist}(c, d) - 1$. Then by induction, $d' = h \,.\, b$ for some $h \in \langle S, T \rangle$ and $b \in B$. Notice that $\operatorname{dist}(d, h \,.\, c) \leq n + 1$. There are three cases to consider.

(i) If $d \in h \,.\, B$, then we are done.

(ii) If $d \notin h \,.\, B$ but $d \in h \,.\, \widetilde{B}$, then by Step 1, we indeed have that $d \in h\langle S \rangle \,.\, B$.

(iii) If $d \notin h \,.\, \widetilde{B}$, then $\operatorname{dist}(d, h \,.\, c) = n + 1$. As $h^{-1} \,.\, d$ is adjacent to $b \in B$, we have $h^{-1} \,.\, d \in \widetilde{D}$. By Step 2, we indeed have that $d \in h\langle S, T \rangle \,.\, B$.

In any case we find that $d \in \langle S, T \rangle \,.\, B$ and we may conclude that $\langle S, T \rangle \,.\, B = \Delta$ by induction.

Step 4. The universal group $\mathcal{U}(\boldsymbol{F})$ is generated by Q.

Let $g \in \mathcal{U}(\boldsymbol{F})$ be arbitrary. By Step 3, we have $g \,.\, c = h \,.\, b$ for some $h \in \langle S, T \rangle$ and $b \in B$. As b and c lie in the same $\mathcal{U}$-orbit, we must have $\lambda(b) = \lambda(c)$, by construction of B. Hence there exists some element $t \in T$ that takes c to b. Putting everything together, we find that $ht \,.\, c = h \,.\, b = g \,.\, c$, from which it follows that $g \in ht \cdot \mathcal{U}_c$. In particular, $g \in \langle Q \rangle$. ■

We obtain as a corollary:

Corollary 3.3.8. *Assume that $\mathcal{U}(\boldsymbol{F})$ is closed and locally compact. Moreover assume that every local group F_i is compactly generated for each $i \in I$. Then the following are equivalent:*

(i) *the local group F_i has finitely many orbits, for every $i \in I$;*

(ii) *the universal group $\mathcal{U}(\boldsymbol{F})$ has finitely many orbits on Δ;*

(iii) *the universal group $\mathcal{U}(\boldsymbol{F})$ is compactly generated.*

Proof. This follows immediately combining Proposition 3.2.2 and Theorems 3.3.6 and 3.3.7. ■

The natural question is now whether the converse to Theorem 3.3.7 holds as well. We established half of the converse in Theorem 3.3.6, but we have not yet found a complete and general proof. The missing piece of the puzzle is the following seemingly innocuous conjecture.

Conjecture 3.3.9. *Assume that $\mathcal{U}(\boldsymbol{F})$ is closed, locally compact, and compactly generated. Then every local group F_i is compactly generated.*

If true, this would yield a precise characterisation: a closed, locally compact universal group $\mathcal{U}(\boldsymbol{F})$ is compactly generated if and only if all local groups are compactly generated and have finitely many orbits.

As motivation, we present some affirmative results in specific special cases. First we assume some additional information about the subgroups $(F_i)^+$ of the local groups.

Theorem 3.3.10. *Assume that $\mathcal{U}(\boldsymbol{F})$ is closed, locally compact, and compactly generated. Moreover, assume that every F_i^+ is compactly generated. Then the local groups F_i are compactly generated.*

Proof. Fix $i \in I$. Using Proposition 1.2.16, we can write F_i as the directed union

$$F_i = \overrightarrow{\bigcup_{\ell \in \mathbb{I}}} H_\ell$$

over the natural directed system of compactly generated open subgroups of F_i. Take any chamber $c \in \Delta$, let $x = \lambda_i(c)$, and let $(F_i)_x$ be the stabiliser of x, which is compact by Proposition 3.3.4. We can hence assume that 0 is the least element of $\mathbb{I}$ and that $H_0 = (F_i)_x$.

For every $\ell \in \mathbb{I}$, we define the local data $\boldsymbol{F}_\ell$ by slightly modifying $\boldsymbol{F}$, replacing the local group F_i by H_ℓ. Define the family of open subgroups

$$K_\ell = \big\langle \mathcal{U}(\boldsymbol{F}_\ell), \mathcal{U}(\boldsymbol{F})_c \big\rangle \leq \mathcal{U}(\boldsymbol{F}).$$

Then $\mathcal{U}(\boldsymbol{F})$ is the directed union

$$\mathcal{U}(\boldsymbol{F}) = \overrightarrow{\bigcup_{\ell \in \mathbb{I}}} K_\ell.$$

By assumption, $\mathcal{U}(\boldsymbol{F})$ is compactly generated, say $\mathcal{U}(\boldsymbol{F}) = \langle Q \rangle$ (with Q a compact set). The family $\{K_\ell \mid \ell \in \mathbb{I}\}$ defines an open cover of Q. By compactness, it follows that $Q \subseteq K_\ell$ for some ℓ, i.e. the subgroups $\mathcal{U}(\boldsymbol{F}_\ell)$ and $\mathcal{U}(\boldsymbol{F})_c$ generate the full group $\mathcal{U}(\boldsymbol{F})$.

Now consider the local actions of $\mathcal{U}(\boldsymbol{F})$ at the i-panel $\mathcal{P}$ containing c. On the one hand, the local actions are given by F_i (by Proposition 3.2.9). On the other hand, the local actions are generated by the local actions of $\mathcal{U}(\boldsymbol{F}_\ell)$ and $\mathcal{U}(\boldsymbol{F})_c$ — in other words, by the subgroups H_ℓ (by construction) and $(F_i)^+$ (by Lemma 3.2.16). Hence $F_i = \langle H_\ell, (F_i)^+ \rangle$ is compactly generated. ■

Remark 3.3.11. Derek Holt kindly provided an example of a group F acting on a set X which, when endowed with the permutation topology, is totally disconnected and locally compact, has finitely many orbits, and is generated by its point stabilisers, but is not compactly generated. We reproduce his example from [Hol19].

Define F as the semidirect product $A \rtimes \langle t \rangle$ of any infinite abelian group A of odd finite exponent, with $t^2 = 1$ and $tat = a^{-1}$ for all $a \in A$. Let F act by left translation on the set $F/\langle t \rangle$ of left cosets of $\langle t \rangle$ in F. This action is transitive, faithful, every point stabiliser $F_{a\langle t \rangle} = \{1, a^2 t\}$ has order two, and finally $F = F^+$. However, F is not generated by *finitely many* point stabilisers.

Hence, a permutation group F with finitely many orbits and whose point stabilisers all have finite orbits, need not be generated by finitely many point stabilisers — not even in the restrictive setting where F is generated by all point stabilisers. The question remains whether or not such a group F can occur as a local group of a compactly generated universal group.

In the realm of rank two right-angled buildings however, i.e. for trees, we do not need additional assumptions. The following result appeared in [Cas20, Proposition 4.1] with a technical proof using group cohomology; we give an argument due to Colin Reid ([Rei19]) involving Bass–Serre theory.

Proposition 3.3.12. *Let G be a compactly generated t.d.l.c. group. Let G act on a tree T without edge inversions, such that the quotient graph $G\backslash T$ is finite and edge stabilisers are compact open subgroups. Then vertex stabilisers are compactly generated.*

Proof. In the quotient graph $G\backslash T$, let $\{v_1, \dots, v_k\}$ be a set of representatives of the vertices, let E be a set of (directed) representatives of the edges, and let $E' \subseteq E$ be a subset of edges representing a spanning tree. From Bass–Serre theory, we can write G in the form

$$G \cong \frac{G_{v_1} * \cdots * G_{v_k} * F(E)}{\left\langle\!\left\langle \bar{e} \cdot \alpha_e(g) \cdot e = \alpha_{\bar{e}}(g) \ \text{ for } g \in G_e, \quad e \cdot \bar{e} \ \text{ for } e \in E, \quad e \ \text{ for } e \in E' \right\rangle\!\right\rangle}.$$

Here $F(E)$ is the free group on E, the bar notation $\bar{e}$ denotes edge inversion, and $\alpha_e \colon G_e \to G_{o(e)}$ is the natural embedding of an edge group into the vertex group of the origin of the edge.

Fix a vertex $v_1 = v$ and note that the subgroup $H_0 = \langle \alpha_e(G_e) \mid e \in E, o(e) = v \rangle \leq G_v$ is open and compactly generated. Use Proposition 1.2.16 to write G_v as a directed union

$$G_v = \overrightarrow{\bigcup_{\ell \in \mathbb{I}}} H_\ell$$

over the natural directed system of compactly generated open subgroups, where we can assume 0 to be the least element of $\mathbb{I}$. For each $\ell \in \mathbb{I}$, define $K_\ell = \langle H_\ell, G_{v_2}, \dots, G_{v_k}, F(E) \rangle \leq G$. Then G is the directed union

$$G = \overrightarrow{\bigcup_{\ell \in \mathbb{I}}} K_\ell.$$

By assumption, G is compactly generated, say $G = \langle Q \rangle$. The family $\{K_\ell \mid \ell \in \mathbb{I}\}$ defines an open cover of Q. By compactness, it follows that $Q \subseteq K_\ell$ for some $\ell \in \mathbb{I}$. More explicitly every $g \in G_v$ is a product of elements in $H_\ell \cup G_{v_2} \cup \cdots \cup G_{v_k} \cup F(E)$. By the normal form theorem for graphs of groups ([Hig76]), one can conclude that $g \in H_\ell$, i.e. that $G_v = H_\ell$ is compactly generated. ■

Corollary 3.3.13. *Assume that Δ is a semiregular tree and that $\mathcal{U}_\Delta(\boldsymbol{F})$ is closed, locally compact, and compactly generated. Then the two local groups are compactly generated as well.*

Proof. Let $G = \mathcal{U}(\boldsymbol{F})$ and $T = \Delta$. Note that the quotient graph $G\backslash T$ is finite by Theorem 3.3.6 and Proposition 3.2.2, and that chamber stabilisers are compact by Proposition 3.3.4. Proposition 3.3.12 then yields that the panel stabilisers are compactly generated, and the result follows from Proposition 3.2.9. ■

Remark 3.3.14. The similarity between the proofs of Theorem 3.3.10 and Proposition 3.3.12 is certainly striking. Hence, one way to settle Conjecture 3.3.9 might be to find a suitable generalisation of the necessary Bass–Serre theory and normal forms, e.g. using the theory of complexes of groups by Martin Bridson and André Haefliger ([BH99]).

3.4 Simplicity

In this section, we develop an analogue for right-angled buildings of Tits's simplicity criterion for groups of tree automorphisms (Theorem 1.4.17). We will establish a general result for groups acting on right-angled buildings, although the conditions we eventually need to impose will be motivated by the universal groups in particular.

Throughout this section, G refers to a fixed group acting on Δ by automorphisms.

First we define some special subgroups of G, with support either contained in some wing or in the complement of some wing.

Definition 3.4.1. For every chamber $c \in \Delta$ and $i \in I$, we define the subgroups

$$V_i(c) = \{g \in G \mid g \,.\, d = d \text{ for all chambers } d \notin X_i(c)\},$$
$$W_i(c) = \{g \in G \mid g \,.\, d = d \text{ for all chambers } d \in X_i(c)\}.$$

Note that $V_i(c)$ fixes the full i-tree-wall containing c, and that $V_i(c)$ equals the intersection of all $W_i(d)$ with $d \in \mathcal{P}_i(c) \setminus \{c\}$.

Definition 3.4.2 (independence property). Let $\mathcal{P}$ be an i-panel and let $\mathcal{T}$ be the corresponding i-tree-wall. For every $c \in \mathcal{P}$ we define the morphism

$$\varphi_c \colon G_{(\mathcal{T})} \to V_i(c), \qquad \text{with } \varphi_c(g) \colon \Delta \to \Delta \colon d \mapsto \begin{cases} g \,.\, d & \text{if } d \in X_i(c); \\ d & \text{otherwise.} \end{cases}$$

The same technique from the proof of Lemma 3.2.8 shows that $\varphi_c(g)$ is indeed an automorphism of Δ, which is clearly contained in $V_i(c)$. We then have an induced embedding

$$G_{(\mathcal{T})} \hookrightarrow \prod_{c \in \mathcal{P}} V_i(c) \hookrightarrow \mathrm{Aut}(\Delta)_{(\mathcal{T})}. \tag{$*$}$$

We say that G satisfies the *independence property* when for every $\mathcal{P}$ the first monomorphism above is in fact an isomorphism.

Proposition 3.4.3. *$\mathcal{U}(\boldsymbol{F})$ satisfies the independence property.*

Proof. Let $g \in \prod_{c \in \mathcal{P}} V_i(c)$ and identify g with its image in $\mathrm{Aut}(\Delta)_{(\mathcal{T})}$. Every panel of Δ is either contained in $\mathcal{T}$ or in an i-wing with base chamber in $\mathcal{T}$. Thus g is an automorphism of Δ with the property that all local actions are either trivial or equal to a local action of some element in $\mathcal{U}(\boldsymbol{F})$. In other words, $g \in \mathcal{U}(\boldsymbol{F})$. ∎

Following the technique of [LB16, Lemma 4.4 and Theorem 4.5], when the action of G is combinatorially dense, we could proceed to show that any nontrivial normal subgroup $N \trianglelefteq G$ contains the derived subgroup of $V_i(c)$ for every $i \in I$ and $c \in \Delta$. If G additionally satisfies the independence property, we readily obtain that N contains in fact the derived subgroups of tree-wall fixators, as

$$\Big[G_{(\mathcal{T})}, G_{(\mathcal{T})}\Big] = \Big[\prod_{c \in \mathcal{P}} V_i(c), \prod_{c \in \mathcal{P}} V_i(c)\Big] = \prod_{c \in \mathcal{P}} \Big[V_i(c), V_i(c)\Big] \leq N.$$

However, in Proposition 3.4.7 we will establish a stronger result, namely that N contains every full tree-wall fixator (instead of merely the derived subgroup).

The following lemma is a straightforward adaptation from [DMdSS18, Lemma 3.17] that holds in a more general setting than only for universal groups.

Lemma 3.4.4. *Let $i \in I$ and let c and d be two chambers in a common i-panel $\mathcal{P}$ of Δ. Let $g \in G$ and assume that $\mathcal{P}$ and $g.\mathcal{P}$ are not parallel, that $\operatorname{proj}_{\mathcal{P}}(g.c) = d$, and $\operatorname{proj}_{g.\mathcal{P}}(d) = g.c$. Consider an automorphism*

$$b \in \prod_{e \in \mathcal{P}\setminus\{c,d\}} V_i(e).$$

Then there exists an automorphism $h \in G$ such that $b = [h, g]$.

Proof. Denote $V_0 = \prod_{e \in \mathcal{P}\setminus\{c,d\}} V_i(e) \leq G$. For every $n \geq 0$, let

$$\mathcal{P}_n = g^n.\mathcal{P}, \qquad c_n = g^n.c, \qquad d_n = g^n.d, \qquad V_n = g^n V_0 g^{-n} \leq G.$$

Then for every $n \geq 0$, the support of V_n is contained in $\bigcup_{e \in \mathcal{P}_n\setminus\{c_n,d_n\}} X_i(e)$. Since by assumption $\operatorname{proj}_{\mathcal{P}}(g.c) = d$ and $\operatorname{proj}_{g.\mathcal{P}}(d) = g.c$, it follows that for every chamber $e \in \mathcal{P}_n \setminus \{c_n, d_n\}$ and every $m > n$, we have $e \in X_i(c_m)$ while $c_m \notin X_i(e)$. Hence $X_i(e) \subseteq X_i(c_m)$ by Lemma 2.3.12. Similarly for every chamber $e' \in \mathcal{P}_m \setminus \{c_m, d_m\}$ we have $X_i(e') \subseteq X_i(d_n)$. Consequently the sets

$$\bigcup_{e \in \mathcal{P}_n\setminus\{c_n,d_n\}} X_i(e) \qquad \text{and} \qquad \bigcup_{e \in \mathcal{P}_n\setminus\{c_m,d_m\}} X_i(e)$$

are disjoint. In other words, the subgroups V_m and V_n have disjoint support when $m \neq n$. Hence the product

$$V = \prod_{n\geq 0} V_n \tag{$*$}$$

can be identified with a subgroup of G, e.g. by [Cap14a, Lemma 5.3]. Note that V fixes the chambers c_m and d_m for every $m \geq 0$, since every factor V_n does.

Now, let us turn our attention to the automorphism $b \in V_0$. For every $n \geq 0$, let $h_n = g^n \cdot b \cdot g^{-n}$ and notice that $h_n \in V_n$. The tuple $(h_n)_{n\geq 0} \in V$ can be identified with an automorphism $h \in G$. The commutator $[h, g]$ then fixes the chambers c_m and d_m for every $m \geq 0$. Furthermore, denoting by y_n the component in V_n of an element $y \in V$ according to the decomposition in $(*)$, we obtain that $[h, g]_n = h_n \cdot (gh^{-1}g^{-1})_n$ for every $n \geq 0$. Hence $[h, g]_0 = b$, while

$$[h, g]_n = h_n \cdot (g\, h^{-1}\, g^{-1})_n = h_n \cdot g\, h_{n-1}^{-1}\, g^{-1} = h_n \cdot h_n^{-1} = \mathrm{id},$$

since $g\, V_{n-1}\, g^{-1} = V_n$. In conclusion, $[h, g] = b$, which proves the lemma. ■

Proposition 3.4.5. *Let Δ be an irreducible right-angled building and assume that the diagram has no isolated nodes. Let $N \trianglelefteq G$ be a nontrivial normal subgroup and assume that the action of G on Δ is combinatorially dense. Then the action of N is combinatorially dense as well.*

Proof. This follows immediately from Propositions 1.4.14 and 2.4.5. ■

Proposition 3.4.6. *Let Δ be a right-angled building and assume that the diagram has no isolated nodes. Let $i \in I$ and let e_1, e_2 be adjacent edges of the i-tree-wall tree Γ_i. Let $N \trianglelefteq G$ be a nontrivial normal subgroup and assume that the action of G on Δ is combinatorially dense. Then N contains an i-hyperbolic automorphism, the axis of which contains both e_1 and e_2.*

Proof. By Proposition 3.4.5, the action of N on Δ is combinatorially dense. Proposition 2.4.5 then yields that the induced action on Γ_i is geometrically dense, so we can apply Proposition 1.4.13. ■

Proposition 3.4.7. *Let Δ be a thick irreducible right-angled building. Let G act on Δ by automorphisms. Assume that G satisfies the independence property and the action is combinatorially dense. Let $N \trianglelefteq G$ be a nontrivial normal subgroup. Then N contains the fixator $G_{(\mathcal{T})}$ of every i-tree-wall $\mathcal{T}$.*

Proof. Let $\mathcal{T}$ be any arbitrary fixed i-tree-wall and let $\mathcal{P}$ be an i-panel of $\mathcal{T}$. In the i-tree-wall tree, $\mathcal{T}$ corresponds to a vertex v, and the edges incident to v correspond to the chambers in $\mathcal{P}$. Let c, d be two distinct chambers of $\mathcal{P}$. By Proposition 3.4.6, there exists some i-hyperbolic element $g \in G$ such that the axis of g in Γ_i contains the two edges corresponding to c and d. We can assume that d is the chamber pointing towards the attracting end of g in $\partial\Gamma_i$ and c towards the repelling end, replacing g by g^{-1} if necessary.

Note that $g\,.\,\mathcal{T} \neq \mathcal{T}$ since g is an i-hyperbolic automorphism. Moreover, by construction, we have $\operatorname{proj}_{\mathcal{P}}(g\,.\,\mathcal{P}) = d$ and $\operatorname{proj}_{\mathcal{P}}(g\,.\,\mathcal{P}) = g\,.\,c$. Hence we can apply Lemma 3.4.4 to obtain that

$$\prod_{e \in \mathcal{P} \setminus \{c,d\}} V_i(e) \subseteq N. \tag{$*$}$$

Since Δ is assumed to be thick, there is at least one such chamber $e \in \mathcal{P} \setminus \{c, d\}$. Now recall that c and d were chosen arbitrarily in $\mathcal{P}$ — we can thus repeat the proof with different pairs $\{c, d\}$ and multiply the resulting equations $(*)$ to obtain that

$$G_{(\mathcal{T})} = \prod_{e \in \mathcal{P}} V_i(e) \subseteq N.$$

As an aside, we remark that we can already apply Proposition 3.4.7 to obtain that the nondiscrete universal groups are monolithic.

Proposition 3.4.8. *Assume not all local groups F_i are free. Then $\mathcal{U}(\boldsymbol{F})$ is monolithic; the monolith is the subgroup generated by all tree-wall fixators and is simple.*

Proof. Write $M = \langle \mathcal{U}_{(\mathcal{T})} \rangle$, where $\mathcal{T}$ ranges over all tree-walls. Note that M is nontrivial, since not all local groups act freely. Thanks to Proposition 3.4.3 and Corollary 3.2.6, we may apply Proposition 3.4.7 to obtain that M is contained in every nontrivial normal subgroup of $\mathcal{U}$. In other words, $\mathcal{U}$ is monolithic with monolith M.

For the simplicity of M, let $N \trianglelefteq M$ be the intersection of all nontrivial normal subgroups of M. A characteristic subgroup of M, it follows that N is a normal subgroup of $\mathcal{U}$. By Proposition 3.4.7 again, N contains the monolith M, so that in fact $N = M$. This shows that M is simple.

Now finally, our simplicity criterion.

Theorem 3.4.9. *Let Δ be a thick and irreducible right-angled building of rank at least two. Let G be a closed subgroup of $\operatorname{Aut}(\Delta)$. Assume that G satisfies the independency property and that the action is hereditarily combinatorially dense. If G^+ is nontrivial, then it is simple group.*

Proof. We use induction on the rank of Δ. In the rank two case, the fact that G^+ is simple follows from Tits's Theorem 1.4.17. For higher rank, let $i \in I$ be such that the diagram remains irreducible upon removal of i (note that such i definitely exists: any leaf of a spanning tree of the underlying graph does the trick). Let $\mathcal{R}$ be any residue of type $I \setminus \{i\}$ and let c be any chamber in $\mathcal{R}$.

For every group $H \leq \operatorname{Aut}(\Delta)$ — not necessarily stabilising $\mathcal{R}$ — we define

$$H\big|_{\mathcal{R}} = \left\{ h\big|_{\mathcal{R}} \;\middle|\; h \in H_{\{\mathcal{R}\}} \right\} \leq \operatorname{Aut}(\mathcal{R}).$$

Also define the restriction morphism

$$\rho\colon G_{\{\mathcal{R}\}} \to G\big|_{\mathcal{R}} \colon g \mapsto g\big|_{\mathcal{R}}$$

with kernel $G_{(\mathcal{R})}$.

Let $N \trianglelefteq G^+ \trianglelefteq G$ be a series of normal subgroups. Restricting to the stabiliser of $\mathcal{R}$, we find that also $N|_{\mathcal{R}} \trianglelefteq G^+|_{\mathcal{R}} \trianglelefteq G|_{\mathcal{R}}$ are normal subgroups. By the induction hypothesis and Proposition 3.4.5, the subgroup $(G|_{\mathcal{R}})^+$ generated by stabilisers of chambers in $\mathcal{R}$ is either trivial or simple. In any case it follows that $(G|_{\mathcal{R}})^+ \leq N|_{\mathcal{R}} \trianglelefteq G^+|_{\mathcal{R}}$. The subgroups stabilising the chamber c in $\mathcal{R}$ then satisfy

$$G_c|_{\mathcal{R}} = \left(G|_{\mathcal{R}}\right)_c = \left(G|_{\mathcal{R}}\right)_c^+ \leq \left(N|_{\mathcal{R}}\right)_c \leq \left(G^+|_{\mathcal{R}}\right)_c = G_c^+|_{\mathcal{R}} = G_c|_{\mathcal{R}}$$

from which we conclude that $N_c|_{\mathcal{R}} = G_c|_{\mathcal{R}}$.

On the other hand, pick $j \in I \setminus \{i\}$ such that $m_{ij} = \infty$. Then the j-tree-wall $\mathcal{T}$ of Δ containing c is completely contained in $\mathcal{R}$, i.e. $G_{(\mathcal{R})} \leq G_{(\mathcal{T})}$. Moreover, by Proposition 3.4.7, we have $G_{(\mathcal{T})} \leq N$. Consequently $N_{(\mathcal{R})} = G_{(\mathcal{R})}$.

In other words, the restrictions of the morphism ρ to the chamber stabilisers N_c and G_c have both identical kernels $N_{(\mathcal{R})} = G_{(\mathcal{R})}$ and identical images $N_c|_{\mathcal{R}} = G_c|_{\mathcal{R}}$. It follows that in fact $N_c = G_c$.

$$\begin{tikzcd}
0 \arrow[r] & N_{(\mathcal{R})} \arrow[r] \arrow[d, equal] & N_c \arrow[r, "\rho"] \arrow[d, dashed] & N_c|_{\mathcal{R}} \arrow[r] \arrow[d, equal] & 0 \\
0 \arrow[r] & G_{(\mathcal{R})} \arrow[r] & G_c \arrow[r, "\rho"] & G_c|_{\mathcal{R}} \arrow[r] & 0
\end{tikzcd}$$

Since c and $\mathcal{R}$ were arbitrary, we have that $G^+ = \langle G_c \mid c \in \Delta \rangle = \langle N_c \mid c \in \Delta \rangle \leq N$. Since N was an arbitrary normal subgroup of G^+, we conclude that G^+ is simple (or trivial). ■

Corollary 3.4.10. *Let M be an irreducible diagram over I. Let $\boldsymbol{F}$ be a collection of closed permutation groups of degree at least three, at least one of which does not acts freely. Then $\mathcal{U}(\boldsymbol{F})^+$ is simple.*

Proof. We know that $\mathcal{U}(\boldsymbol{F})$ is closed by Proposition 3.3.3, satisfies the independence property by Proposition 3.4.3, and is hereditarily combinatorially dense by Corollaries 3.2.6 and 3.2.11. Hence we can apply Theorem 3.4.9. Finally, $\mathcal{U}(\boldsymbol{F})^+$ cannot be trivial unless all local groups act freely. ■

Corollary 3.4.11. *Let M be an irreducible diagram over I. Let $\boldsymbol{F}$ be a collection of closed permutation groups of degree at least three, indexed by I. Assume that not all local actions are free. Then $\mathcal{U}(\boldsymbol{F})$ is simple if and only if F_i is generated by point stabilisers for every $i \in I$ and transitive for every i in some vertex cover of M.*

Proof. By Proposition 3.3.1, $\mathcal{U}(\boldsymbol{F})^+$ is a nontrivial normal subgroup of $\mathcal{U}(\boldsymbol{F})$. Hence this follows immediately from Corollary 3.4.10 and Theorem 3.2.20. ■

3.5 City products

Earlier in Lemma 3.2.1, we already observed that the universal group construction behaves nicely with respect to disjoint unions of diagrams. This operation on diagrams is a special case of the city product from Section 2.7 (over a diagram with only isolated nodes). In this section, we generalise the special case of Lemma 3.2.1 to arbitrary city products.

We start with carefully defining all buildings involved. For each $\ell \in \{1, \dots, n\}$, consider

- a diagram M_ℓ over an index set I_ℓ,
- a semiregular right-angled building Δ_ℓ of type M_ℓ,
- a colouring λ_ℓ of Δ_ℓ with colour sets Ω_i (indexed by $i \in I_\ell$),
- a collection $\boldsymbol{F}_\ell$ of permutation groups $F_i \leq \mathrm{Sym}(\Omega_i)$ (indexed by $i \in I_\ell$).

Then for each $\ell \in \{1, \dots, n\}$, we have a universal group $\mathcal{U}_\ell = \mathcal{U}(\boldsymbol{F}_\ell)$ over Δ_ℓ. Next, consider

- a right-angled diagram M over $\{1, \dots, n\}$,
- the city product $\Delta = \maltese_M(\{\Delta_1, \dots, \Delta_n\})$ of type $\maltese_M(\{M_1, \dots, M_n\})$ over $I = \bigsqcup_{\ell=1}^n I_\ell$,
- the "unified" local data $\boldsymbol{F}$ over I, defined by $(\boldsymbol{F})_i = (\boldsymbol{F}_{\ell(i)})_i$ for every $i \in I$.

Finally, consider

- the skeletal building Φ of the city product Δ,
- the local data $\boldsymbol{F'}$ over $\{1, \dots, n\}$, defined by $(\boldsymbol{F'})_\ell = \mathcal{U}(\boldsymbol{F}_\ell)$ for every $\ell \in \{1, \dots, n\}$.

Theorem 3.5.1. *With notation from above, we have*

$$\mathcal{U}_\Delta(\boldsymbol{F}) \cong \mathcal{U}_\Phi(\boldsymbol{F'}).$$

In colloquial terms, the universal group over a city product of buildings is isomorphic to the universal group over the skeletal building of the universal groups over the factor buildings.

Proof. Equip Δ with the colouring λ' from Lemma 2.7.9, assigning colours in the sets Ω_i (indexed by $i \in I$). Also equip its skeletal building Φ with the colouring φ from Proposition 2.7.13, assigning colours in the sets Δ_ℓ (indexed by $\ell \in \{1, \dots, n\}$).

First, every automorphism of Δ induces an automorphism of its skeletal building, hence we have a natural monomorphism

$$\iota \colon \mathrm{Aut}(\Delta) \hookrightarrow \mathrm{Aut}(\Phi).$$

Let $g \in \mathcal{U}_\Delta(\boldsymbol{F}) \leq \mathrm{Aut}(\Delta)$, let $\mathcal{R}$ be any panel of Φ of type ℓ, and consider the local action of $\iota(g)$ as an automorphism of Φ at the panel $\mathcal{R}$. For readability, we will identify g with its image $\iota(g)$. We can also identify $\mathcal{R}$ with a residue of Δ of type I_ℓ (which is isomorphic to Δ_ℓ). Then the local action

$$\sigma_\varphi(g, \mathcal{R}) = \varphi_\ell\big|_{g\,.\,\mathcal{R}} \circ g\big|_{\mathcal{R}} \circ \varphi_\ell\big|_{\mathcal{R}}^{-1}$$

is the composition of three isomorphisms $\Delta_\ell \to \mathcal{R} \to g\,.\,\mathcal{R} \to \Delta_\ell$ and is hence an automorphism of Δ_ℓ (instead of merely a permutation). Hence, we can consider the local action of $\sigma_\varphi(g, \mathcal{R})$ at an i-panel $\mathcal{P}$ of Δ_ℓ with $i \in I_\ell$. First, we define the i-panel $\mathcal{P}' = \varphi_\ell|_{\mathcal{R}}^{-1}(\mathcal{P})$ of Δ. Then

$$\begin{aligned}
\sigma_{\lambda_\ell}(\sigma_\varphi(g, \mathcal{R}), \mathcal{P}) &= (\lambda_\ell)_i\big|_{\sigma_\varphi(g,\mathcal{R})\,.\,\mathcal{P}} \circ \big(\varphi_\ell\big|_{g\,.\,\mathcal{R}} \circ g\big|_{\mathcal{R}} \circ \varphi_\ell\big|_{\mathcal{R}}^{-1}\big)\big|_{\mathcal{P}} \circ (\lambda_\ell)_i\big|_{\mathcal{P}}^{-1} \\
&= (\lambda_\ell)_i\big|_{\sigma_\varphi(g,\mathcal{R})\,.\,\mathcal{P}} \circ \varphi_\ell\big|_{g\,.\,\mathcal{P}'} \circ g\big|_{\mathcal{P}'} \circ \varphi_\ell\big|_{\mathcal{P}'}^{-1} \circ (\lambda_\ell)_i\big|_{\mathcal{P}}^{-1} \\
&= \lambda'_i\big|_{g\,.\,\mathcal{P}'} \circ g\big|_{\mathcal{P}'} \circ \lambda'_i\big|_{\mathcal{P}'}^{-1} \\
&= \sigma_{\lambda'}(g, \mathcal{P}'). \qquad (*)
\end{aligned}$$

A commutative diagram helps tremendously.

$$
\begin{tikzcd}
\Delta \arrow[r, "g"] \arrow[d, equal] \arrow[dr, phantom, "\Downarrow \iota"] \arrow[ddd, bend right=60, "\lambda'_i"'] & \Delta \arrow[d, equal] \arrow[ddd, bend left=60, "\lambda'_i"] \\
\Phi \arrow[r, "\iota(g)"'] \arrow[d, "\varphi_\ell"'] & \Phi \arrow[d, "\varphi_\ell"] \\
\Delta_\ell \arrow[r, "\sigma_\varphi(g, \bullet)"] \arrow[d, "(\lambda_\ell)_i"'] & \Delta_\ell \arrow[d, "(\lambda_\ell)_i"] \\
\Omega_i \arrow[r, "\sigma_{\lambda_\ell}(\sigma_\varphi(g, \bullet), \bullet)"] & \Omega_i
\end{tikzcd}
$$

Since $g \in \mathcal{U}_\Delta(\boldsymbol{F})$, the result of Equation $(*)$ is a permutation in F_i — hence we can conclude that $\sigma_\varphi(g, \mathcal{R}) \in \mathcal{U}(\boldsymbol{F}_\ell)$. This shows that the image of $\mathcal{U}_\Delta(\boldsymbol{F})$ under ι is contained in $\mathcal{U}_\Phi(\boldsymbol{F'})$.

Conversely, let $g \in \mathcal{U}_\Phi(\boldsymbol{F'})$. We can identify g (an automorphism of Φ) with a permutation of Δ, and claim that this permutation is type-preserving, i.e. g is in fact an automorphism of Δ. Indeed, let $c \sim_i d$ be i-adjacent chambers in Δ. Let $\ell \in \{1, \dots, n\}$ be such that $i \in I_\ell \subseteq I$ and let $\mathcal{R}$ be the residue of Δ of type I_ℓ containing c and d. Then $\mathcal{R}$ is an ℓ-panel of Φ. The local action $\sigma_\varphi(g, \mathcal{R})$ is an element of $(\boldsymbol{F'})_\ell = \mathcal{U}(\boldsymbol{F}_\ell) \leq \operatorname{Aut}(\Delta_\ell)$. Hence

$$g\big|_{\mathcal{R}} = \varphi_\ell\big|_{g\,.\,\mathcal{R}}^{-1} \circ \sigma_\varphi(g, \mathcal{R}) \circ \varphi_\ell\big|_{\mathcal{R}}$$

is a composition of isomorphisms $\mathcal{R} \to \Delta_\ell \to \Delta_\ell \to g\,.\,\mathcal{R}$, each of which preserves i-adjacency. In particular $g\,.\,c \sim_i g\,.\,d$. Since c and d were arbitrary, we conclude that $g \in \operatorname{Aut}(\Delta)$.

Let $\mathcal{P}'$ be any i-panel in Δ with $i \in I_\ell$, let $\mathcal{R}$ be the I_ℓ-residue of Δ containing $\mathcal{P}'$, let $\mathcal{P} = \varphi_\ell(\mathcal{P}')$ in Δ_ℓ. The reverse calculation of Equation $(*)$ shows that the local action satisfies

$$\sigma_{\lambda'}(g, \mathcal{P}') = \sigma_{\lambda_\ell}(\sigma_\varphi(g, \mathcal{R}), \mathcal{P}).$$

Since $\sigma_\varphi(g, \mathcal{R}) \in \mathcal{U}(\boldsymbol{F}_\ell)$ we have that $\sigma_{\lambda'}(g, \mathcal{P}') \in (\boldsymbol{F}_\ell)_i = F_i$. Hence, $g \in \mathcal{U}_\Delta(\boldsymbol{F})$.

in conclusion, the restriction of ι is an isomorphism $\mathcal{U}_\Delta(\boldsymbol{F}) \to \mathcal{U}_\Phi(\boldsymbol{F}')$. ■

4 Restricted universal groups

> *All you need to know for the moment is that the universe is a lot more complicated than you might think, even if you start from a position of thinking it's pretty damn complicated in the first place.*
>
> — Douglas Adams, *The Hitchhiker's Guide to the Galaxy*

We finish our study of the universal groups and now follow Le Boudec's idea ([LB16]) to allow for singularities in the building automorphisms.

4.1 Definition

Definition 4.1.1. Using the same setup as in Definition 3.1.4, let $\boldsymbol{F}$ be a collection of permutation groups $F_i \leq \mathrm{Sym}(\Omega_i)$, indexed by $i \in I$. Let Δ be a semiregular right-angled building over I with parameters $q_i = |\Omega_i|$, equipped with a colouring λ using the sets Ω_i as i-colours. Then we define

$$\mathcal{G}_\Delta^\lambda(\boldsymbol{F}) = \left\{ g \in \mathrm{Aut}(\Delta) \;\middle|\; \sigma_\lambda(g, \mathcal{P}) \in F_i \text{ for every } i \in I \text{ and } \textit{all but finitely many } i\text{-panels } \mathcal{P} \right\}.$$

We will, again, usually abbreviate $\mathcal{G}_\Delta^\lambda(\boldsymbol{F})$ to $\mathcal{G}(\boldsymbol{F})$ if the context is clear. Note that by Lemma 3.1.2 $\mathcal{G}(\boldsymbol{F})$ is closed under composition and inversion, and is hence a subgroup of $\mathrm{Aut}(\Delta)$. We clearly have an inclusion

$$\mathcal{U}_\Delta^\lambda(\boldsymbol{F}) \leq \mathcal{G}_\Delta^\lambda(\boldsymbol{F}) \leq \mathrm{Aut}(\Delta).$$

Definition 4.1.2 (singularity). For an element $g \in \mathcal{G}(\boldsymbol{F})$, an i-panel $\mathcal{P}$ with $\sigma_\lambda(g, \mathcal{P}) \notin F_i$ will be called a *singularity* of g. Then in other words, $\mathcal{G}(\boldsymbol{F})$ is the group of all building automorphisms with only finitely many singularities. We denote by $S(g)$ the set of all singularities of g.

Lemma 4.1.3. *Let Δ be a right-angled building over I and let $i, j \in I$ be such that $m_{ij} = 2$. If q_j is infinite, then a panel of type i can never be a singularity of an element $g \in \mathcal{G}_\Delta(\boldsymbol{F})$.*

Proof. Since the i-panels in a residue of type $\{i, j\}$ with $m_{ij} = 2$ are parallel, this follows immediately from Proposition 3.1.3. ■

Even though the definition of $\mathcal{G}(\boldsymbol{F})$ may suggest otherwise, the local action at a singularity cannot just be any permutation. Indeed, some panels cannot be a singularity at all! One simple observation is that due to Proposition 3.1.3, the set of panels parallel to any singularity must be finite. However, consider the following.

Definition 4.1.4 (ladder, rungs). A *ladder* is a rank three Coxeter system of type

$$M = \begin{pmatrix} 1 & 2 & 2 \\ 2 & 1 & \infty \\ 2 & \infty & 1 \end{pmatrix}.$$

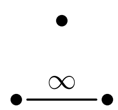

The *rungs* of a ladder are the panels of the type corresponding to the isolated node of the diagram. The origin of the names should be evident from one look at a ladder's Coxeter complex, which we have visualised in Figure 4.1.

When working with right-angled buildings, *ladders* will also refer to residues of the building which have the type of a ladder, and its *rungs* will be the corresponding panels of the building.

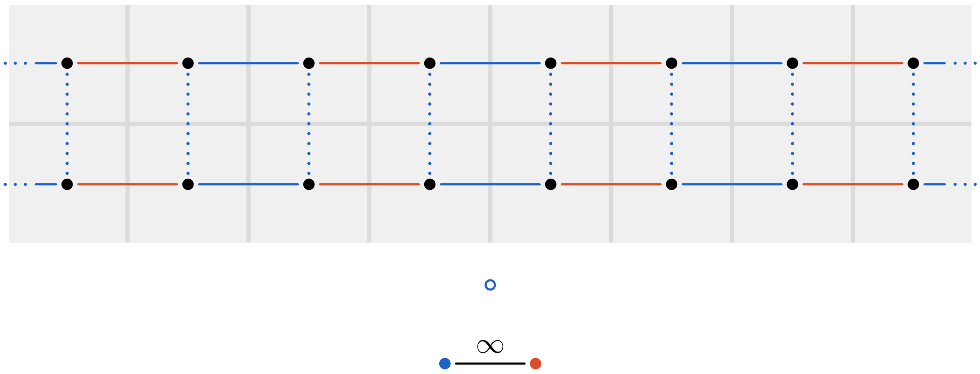

Figure 4.1. The Coxeter complex of a ladder.
Nodes ● and ● correspond to solid edges, node ○ to dotted edges.

Lemma 4.1.5. *Let Δ be a right-angled building over I and let $\mathcal{R}$ be a ladder in Δ. Then a rung of $\mathcal{R}$ can never be a singularity of an element $g \in \mathcal{G}_\Delta(\boldsymbol{F})$.*

Proof. Since the rungs of a ladder are parallel, this follows immediately from Proposition 3.1.3. ∎

Proposition 4.1.6. *Let $\mathcal{P}$ be an i-panel of a right-angled building Δ such that the set of all panels parallel to $\mathcal{P}$ is unbounded. Then $\mathcal{P}$ is the rung of a ladder in Δ.*

Proof. Let $\mathcal{P}'$ be a panel such that $\operatorname{dist}(\mathcal{P}, \mathcal{P}') \geq |I|$. Let $c \in \mathcal{P}$ and let $c' = \operatorname{proj}_{\mathcal{P}'}(c)$. A minimal gallery γ from c to c' has only types in $i^\perp$ (by Proposition 2.3.7) and length $|I|$. By the pigeonhole principle, there is a type $j \in i^\perp$ that occurs at least twice in the type of γ. In between, there must exist a type $k \in i^\perp$ in γ such that $m_{jk} = \infty$, since the type of γ is a reduced word. Hence, we have found three types $\{i, j, k\}$ satisfying $m_{ij} = m_{ik} = 2$ and $m_{jk} = \infty$, i.e. the i-panel $\mathcal{P}$ is the rung of at least one ladder in Δ. ∎

As a corollary to Lemma 4.1.5, depending on the building's diagram, it is possible that allowing for a finite number of singularities in fact does not expand the group. For instance, recall the Bourdon building from Figure 1.12, the diagram of which was a pentagon. Every vertex of the pentagon has an "opposite" edge, together defining a ladder with the original vertex corresponding to the rungs. The Coxeter complex clearly illustrates that every panel is the rung of some ladder in the building. Hence, for the Bourdon building Δ of type $I_{5,2}$ it follows that in fact

$$\mathcal{G}_\Delta^\lambda(\boldsymbol{F}) = \mathcal{U}_\Delta^\lambda(\boldsymbol{F}).$$

Similar to the prime graphs in the city products from Section 2.7, a natural question is again how exceptional this example is. Let us first formalise some definitions.

Definition 4.1.7 (ladderless, ladderfull). Let G be a simple, undirected graph. A vertex v of G will be called a *rung* if there exist two other, adjacent vertices $w_1 \sim w_2$ in G such that v is adjacent to neither w_1 nor w_2. We call the graph G *ladderfull* if every vertex is a rung, and *ladderless* if no vertex is a rung.

Then for any building Δ with a ladderfull diagram, we have $\mathcal{G}^{\lambda}_{\Delta}(\boldsymbol{F}) = \mathcal{U}^{\lambda}_{\Delta}(\boldsymbol{F})$.

Computational results for graphs on at most eleven vertices are presented in Table 4.1. In addition, as an illustration, we include all ten irreducible ladderfull diagrams in Figure 4.2. The percentage of ladderfull graphs versus the total number of simple undirected graphs looks quite interesting, and we are interested in the asymptotic behaviour as n goes to infinity. However, what mostly catches the eye are the numbers of ladderless diagrams — these are precisely the partition numbers from number theory and combinatorics. There is a simple explanation.

Proposition 4.1.8. *A simple undirected graph is ladderless if and only if its complement is a disjoint union of complete graphs.*

Proof. A graph is ladderless if and only if its complement graph has no path graph as an induced subgraph on three vertices. This is then easily shown to imply that every connected component of the complement graph must be a complete graph. ■

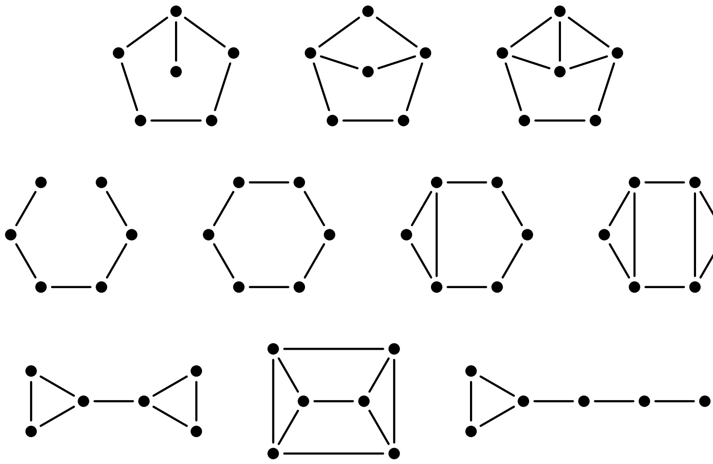

Figure 4.2. All irreducible diagrams of rank six where every panel is the rung of some ladder. For clarity, the labels ∞ on the edges are omitted.

Even for singularities that are not the rung of a ladder, there is some restriction on the local action. Recall the Young overgroups from Definition 1.1.16.

Proposition 4.1.9. *Let $g \in \mathcal{G}(\boldsymbol{F})$ and let $\mathcal{P}$ be a singularity of g of type i. Assume that i is not an isolated node of the diagram. Then at least one of the following holds:*

- *the local action $\sigma_\lambda(g, \mathcal{P})$ is contained in the Young overgroup $\widehat{F_i}$ of the local group, or*
- *F_i has at least two infinite orbits, while F_j has finite degree for all $j \neq i$.*

Proof. Assume that $\sigma = \sigma_\lambda(g, \mathcal{P})$ is not contained in $\widehat{F_i}$ — i.e., there exists a colour $x \in \Omega_i$ such that x and $\sigma.x$ are contained in *different* F_i-orbits. Let $c \in \mathcal{P}$ be such that $\lambda_i(c) = x$, let $j \neq i$, and let d be any chamber j-adjacent to c. Since $\lambda_i(c) = \lambda_i(d)$ and $\lambda_i(g\,.\,c) = \lambda_i(g\,.\,d)$, the i-panel containing d is again a singularity. As there are only finitely many singularities, this implies that the parameter q_j is finite, for every $j \neq i$.

Next, we claim that both orbits $F_i\,.\,x$ and $F_i\,.\,(\sigma.x)$ are infinite. Assume by means of contradiction that $X = F_i\,.\,x$ is finite. Since σ is a permutation, and $x \in X$ with $\sigma.x \notin X$, there exists a colour $y \notin X$ with $\sigma.y \in X$. Let $c' \in \mathcal{P}$ be such that $\lambda_i(c') = y$ and note that $c \neq c'$. Since i is not an isolated node in the diagram, consider $k \in I$ with $m_{ik} = \infty$ and choose $d \sim_k c$ and $d' \sim_k c'$. Then

n	total number of … diagrams on n nodes				
	unrestricted	irreducible	ladderfull	irreducible & ladderfull	ladderless
4	11	6	1	0	5
5	34	21	4	1	7
6	156	112	24	10	11
7	1 044	853	191	132	15
8	12 346	11 117	3 095	2 719	22
9	274 668	261 080	95 208	90 871	30
10	12 005 168	11 716 571	5 561 999	5 452 862	42
11	1 018 997 864	1 006 700 565	592 458 683	586 604 553	56

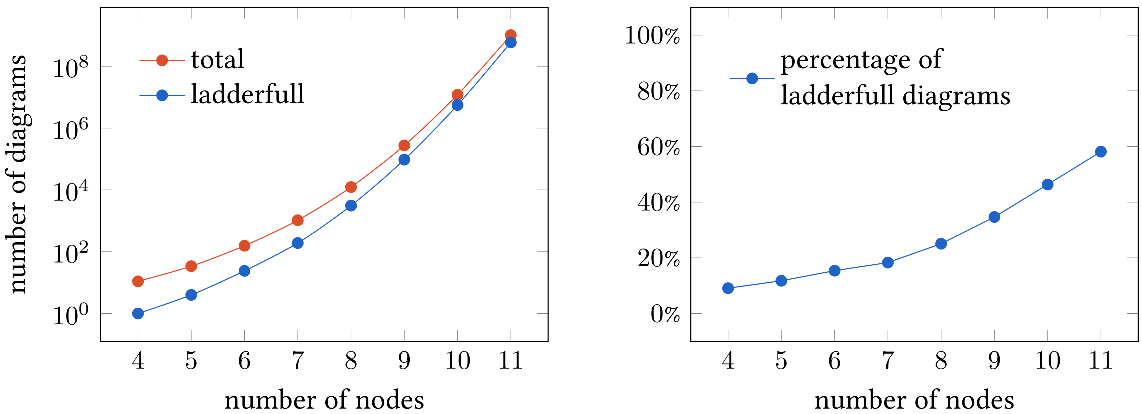

Table 4.1. A comparison of the number of irreducible, ladderfull, and ladderless diagrams. The first column is [OEISa], the second is [OEISb], the fifth is [OEISd].

the local actions at the i-panels containing d and d' are again permutations not contained in $\widehat{F}_i$. We can repeat this construction with those panels to obtain an apartment of type $\{i, k\}$,

$$\cdots \sim_i d \sim_k c \sim_i c' \sim_k d' \sim_i \cdots$$

every chamber of which lies in a singularity of type i. This contradiction shows that the orbit $F_i \,.\, x$ cannot be finite. A similar construction shows that $F_i \,.\, (\sigma . x)$ cannot be finite. ■

It follows from Proposition 4.1.9 the that local actions at singularities are contained in the Young overgroups of the local groups, except for possibly one single type, depending on the finiteness of the parameters. Motivated by this observation and by Lemma 4.1.3, it makes sense to restrict our setting to *locally finite* right-angled buildings for the most interesting results. We then have in particular the following corollary.

Corollary 4.1.10. *Assume that Δ is locally finite. Then we have an inclusion $\mathcal{G}(\boldsymbol{F}) \leq \mathcal{U}(\widehat{\boldsymbol{F}})$, with $\widehat{\boldsymbol{F}}$ the local data obtained from the Young overgroups of the local groups in $\boldsymbol{F}$.*

Proof. This follows immediately from Proposition 4.1.9. ■

From now on, we will always assume that the building Δ is locally finite,
i.e. that the local groups are permutation groups of finite degree.

Definition 4.1.11 (restricted universal group). Let $\boldsymbol{F}$ and $\acute{\boldsymbol{F}}$ be two collections of permutation groups of finite degree, indexed by $i \in I$ and satisfying

$$F_i \leq \acute{F}_i \leq \widehat{F}_i \leq \mathrm{Sym}(\Omega_i)$$

for every $i \in I$. In particular, F_i and $\acute{F}_i$ have identical orbits. Let Δ be a semiregular right-angled building over I with parameters $q_i = |\Omega_i|$, equipped with a colouring λ using the sets Ω_i as i-colours. Then the *restricted universal group* of $\boldsymbol{F}$ and $\acute{\boldsymbol{F}}$ over Δ is by definition the group

$$\mathcal{G}^{\lambda}_{\Delta}(\boldsymbol{F}, \acute{\boldsymbol{F}}) = \mathcal{G}^{\lambda}_{\Delta}(\boldsymbol{F}) \cap \mathcal{U}^{\lambda}_{\Delta}(\acute{\boldsymbol{F}}).$$

In words, $\mathcal{G}^{\lambda}_{\Delta}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ is the group of all automorphisms that locally act like permutations in F_i but with a finite number of exceptions, where the local action still follows a prescribed fashion, namely a permutation in $\acute{F}_i$. We will continue to refer to those exceptions as *singularities*, i.e. a singularity of g is any i-panel $\mathcal{P}$ such that $\sigma_\lambda(g, \mathcal{P}) \in \acute{F}_i \setminus F_i$. The set of all singularities will still be denoted as $S(g)$.

The local groups $\acute{F}_i$ can be chosen between F_i and $\widehat{F}_i$. In the extreme cases, we clearly have

$$\mathcal{G}^{\lambda}_{\Delta}(\boldsymbol{F}, \boldsymbol{F}) = \mathcal{U}^{\lambda}_{\Delta}(\boldsymbol{F}) \qquad \text{and} \qquad \mathcal{G}^{\lambda}_{\Delta}(\boldsymbol{F}, \widehat{\boldsymbol{F}}) = \mathcal{G}^{\lambda}_{\Delta}(\boldsymbol{F}),$$

where the latter equality follows by Corollary 4.1.10.

Remark 4.1.12. Restricted universal groups were first introduced by Adrien Le Boudec in [LB16], in the setting of trees. The name originates from [CRW19], where the authors remark the analogy with *restricted direct products*. There are a few related constructions in the literature, but especially in topological group theory, the restricted direct product of a collection of groups $\{G_i\}_{i \in I}$ with respect to subgroups $\{H_i \leq G_i\}_{i \in I}$ (over any index set I) is by definition the subgroup

$$\left\{ (g_i)_{i \in I} \in \prod_{i \in I} G_i \;\middle|\; g_i \in H_i \text{ for all but finitely many } i \in I \right\}$$

of the standard direct product. Recall also Example 1.2.23 (iv). Just like the restricted direct product imposes extra restraints on all but finitely many factors, automorphisms in $\mathcal{G}^{\lambda}_{\Delta}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ are essentially automorphisms in $\mathcal{U}^{\lambda}_{\Delta}(\acute{\boldsymbol{F}})$ with extra constraints on all but finitely many panels.

4.2 Permutational properties

To start, we mention some analogues of properties of the standard universal groups in Section 3.2. Like before, we will regularly focus on irreducible buildings, although the reduction to lower rank is slightly less "clean" for restricted universal groups.

Lemma 4.2.1. *Let Δ be a reducible right-angled building over I. Let $J_1, \dots, J_m$ be the connected components of (the underlying graph of) the diagram with at least two nodes, and let J be their union. Let $k_1, \dots, k_n$ be the isolated nodes of the diagram, and let K be their union.*

(i) *If $m = 0$ (i.e. all nodes are isolated), then $\mathcal{G}_\Delta(\boldsymbol{F}, \acute{\boldsymbol{F}}) = \mathcal{U}_\Delta(\acute{\boldsymbol{F}}) \cong \acute{F}_{k_1} \times \cdots \times \acute{F}_{k_n}$.*

(ii) *If $m = 1$, then $\mathcal{G}_\Delta(\boldsymbol{F}, \acute{\boldsymbol{F}}) \cong \mathcal{G}_{\mathcal{R}_J}(\boldsymbol{F}|_J, \acute{\boldsymbol{F}}|_J) \times \mathcal{U}_{\mathcal{R}_K}(\boldsymbol{F}|_K) \cong \mathcal{G}_{\mathcal{R}_J}(\boldsymbol{F}|_J, \acute{\boldsymbol{F}}|_J) \times F_{k_1} \times \cdots \times F_{k_n}$, where $\mathcal{R}_J$ and $\mathcal{R}_K$ are residues of types J and K, respectively.*

(iii) *If $m \geq 2$, then $\mathcal{G}_\Delta(\boldsymbol{F}, \acute{\boldsymbol{F}}) = \mathcal{U}_\Delta(\boldsymbol{F}) \cong \mathcal{U}_{\mathcal{R}_1}(\boldsymbol{F}|_{J_1}) \times \cdots \times \mathcal{U}_{\mathcal{R}_m}(\boldsymbol{F}|_{J_m})$.*

In particular, if the diagram is reducible and has no isolated nodes, then $\mathcal{G}_\Delta(\boldsymbol{F}, \acute{\boldsymbol{F}}) = \mathcal{U}_\Delta(\boldsymbol{F})$.

Proof. Case (i) is immediate, since there are only finitely many panels of each type. Case (ii) follows readily from the fact that Δ is isomorphic to the direct product of $\mathcal{R}_J$ with a finite complete graph for every $k \in K$. Case (iii) follows from Lemma 4.1.5, since every panel is the rung of a ladder. ■

The next few propositions follow for free from the inclusions $\mathcal{U}_\Delta(\boldsymbol{F}) \leq \mathcal{G}_\Delta(\boldsymbol{F}, \acute{\boldsymbol{F}}) \leq \mathcal{U}_\Delta(\acute{\boldsymbol{F}})$ and the corresponding properties of the universal groups.

Proposition 4.2.2. *Two residues lie in the same orbit of $\mathcal{G}_\Delta(\boldsymbol{F}, \acute{\boldsymbol{F}})$ if and only if they are of the same type and harmonious. In particular, two chambers c and c' lie in the same orbit of $\mathcal{G}_\Delta(\boldsymbol{F}, \acute{\boldsymbol{F}})$ if and only if their colours $\lambda_i(c)$ and $\lambda_i(c')$ lie in the same F_i-orbit for every $i \in I$.*

Proof. For every $i \in I$, the orbits of the local groups F_i and $\acute{F}_i$ agree. By Proposition 3.2.2, the same is true for the orbits of $\mathcal{U}_\Delta(\boldsymbol{F})$ and $\mathcal{U}_\Delta(\acute{\boldsymbol{F}})$ on Δ. The result follows. ■

The next lemma yields in particular a converse to Lemma 4.1.5: every panel that is *not* the rung of a ladder, is a singularity of some element of $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$.

Lemma 4.2.3. *Let f be a permutation in $\acute{F}_i$ and let $\mathcal{P}$ be an i-panel. Then there exists an automorphism $g \in \mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ with the following properties:*

(i) *g stabilises $\mathcal{P}$;*

(ii) *the local action is equal to f at every panel parallel to $\mathcal{P}$;*

(iii) *the local action is a permutation in F_i at every other panel.*

Proof. First extend the permutation on $\mathcal{P}$ to an automorphism of the building (e.g. in $\mathcal{U}(\acute{\boldsymbol{F}})$, using Lemma 3.2.7). Then apply Proposition 3.1.9 with $\mathcal{P}$ as the convex panel-closed set. ■

Proposition 4.2.4. *The action of $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ on Δ is cobounded.*

Proof. This follows immediately from the fact that $\mathcal{U}(\boldsymbol{F}) \leq \mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ and Corollary 3.2.5. ■

Proposition 4.2.5. *Let Δ be a right-angled building such that the diagram does not have isolated nodes. Then the action of $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ on Δ is combinatorially dense.*

Proof. This follows immediately from the fact that $\mathcal{U}(\boldsymbol{F}) \leq \mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ and Corollary 3.2.6. ■

Definition 4.2.6. For every panel $\mathcal{P}$ we define the subgroup

$$K_{\mathcal{P}} = \big\{ g \in \mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}}) \;\big|\; g \text{ stabilises } \mathcal{P} \text{ and every singularity is parallel to } \mathcal{P} \big\}.$$

For every panel $\mathcal{P}$ and integer $n \geq 0$ we define the subgroup

$$K_{n,\mathcal{P}} = \big\{ g \in \mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}}) \;\big|\; g \text{ stabilises } \mathcal{P} \text{ and } \operatorname{dist}(\mathcal{P}, \mathcal{P}') \leq n \text{ for every singularity } \mathcal{P}' \big\}.$$

Note that both are indeed subgroups by Lemma 3.1.2.

Lemma 4.2.7. *Let $\mathcal{P}$ be an i-panel and $f \in \acute{F}_i$ be a permutation. Then there exists an automorphism $g \in K_{\mathcal{P}}$ such that $\sigma_\lambda(g, \mathcal{P}) = f$.*

Proof. This is a special case of Proposition 3.1.9. ∎

Lemma 4.2.7 yields in particular as a direct consequence that

$$\big[K_{\mathcal{P}} : \mathcal{U}(\boldsymbol{F})_{\{\mathcal{P}\}}\big] = \big[\acute{F}_i : F_i\big].$$

Proposition 4.2.8. *Let $\mathfrak{P}$ be a set of representatives of harmonious panels of Δ. Then*

$$\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}}) = \big\langle \mathcal{U}(\boldsymbol{F}), K_{\mathcal{P}} \mid \mathcal{P} \in \mathfrak{P} \big\rangle.$$

Proof. Let $g \in \mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$. We proceed by induction on the number $|S(g)|$ of singularities. Of course, if $S(g) = \emptyset$ then $g \in \mathcal{U}(\boldsymbol{F})$ and there is nothing to show, so assume that $S(g) \geq 1$ and let $\mathcal{P}$ be a singularity of type i. There is an i-panel $\mathcal{P}_0 \in \mathfrak{P}$ such that $g \,.\, \mathcal{P}$ and $\mathcal{P}_0$ are harmonious. Moreover, there is an automorphism $h \in \mathcal{U}(\boldsymbol{F})$ such that $hg \,.\, \mathcal{P} = \mathcal{P}_0$ by Propositions 3.2.2 and 4.2.2. We can assume that the local action of h at $\mathcal{P}$ is the identity by Lemma 3.2.7. Next let $\sigma = \sigma_\lambda(g, \mathcal{P})$. Using Lemma 4.2.7 we can find an automorphism $h_0 \in K_{\mathcal{P}_0}$ satisfying $\sigma_\lambda(h_0, \mathcal{P}_0) = \sigma^{-1}$. Let us bundle everything in a commutative diagram:

$$\begin{tikzcd}
\mathcal{P} \arrow[rrr, "h'"] \arrow[dr, "h"] \arrow[ddd, "\lambda_i"'] & & & \mathcal{P} \arrow[r, "g"] \arrow[ddd, "\lambda_i"] & g\,.\,\mathcal{P} \arrow[ddd, "\lambda_i"] \\
& \mathcal{P}_0 \arrow[r, "h_0"] \arrow[d, "\lambda_i"'] & \mathcal{P}_0 \arrow[ur, "h^{-1}"'] \arrow[d, "\lambda_i"] & & \\
& \Omega_i \arrow[r, "\sigma^{-1}"] \arrow[dl, equal] & \Omega_i \arrow[dr, equal] & & \\
\Omega_i \arrow[rrr, "\sigma^{-1}"] & & & \Omega_i \arrow[r, "\sigma"] & \Omega_i
\end{tikzcd}$$

Let us also abbreviate by h' the conjugation $h^{-1} \cdot h_0 \cdot h$. Then singularities of h' are parallel to $\mathcal{P}$. Consider the automorphism $g' = g \cdot h'$. By construction, the local action of g' at $\mathcal{P}$ is the identity. Now let $\mathcal{P}'$ be any other panel that is not parallel to $\mathcal{P}$. Then neither is $h' \,.\, \mathcal{P}'$ parallel to $\mathcal{P}$. Since

$$\sigma_\lambda(g', \mathcal{P}') = \sigma_\lambda(g, h' \,.\, \mathcal{P}') \circ \sigma_\lambda(h', \mathcal{P}'),$$

where the second factor is known to be a permutation in F_i, we find that $\mathcal{P}'$ is a singularity of g' if and only if $h' \,.\, \mathcal{P}'$ is a singularity of g. Together with the singularities of g parallel to $\mathcal{P}$, this implies that g' has strictly less singularities than g. The induction hypothesis finishes the proof. ∎

It is worth noting that the set $\mathfrak{P}$ in Proposition 4.2.8 is finite. Indeed, denote the number of orbits of the local groups by $m_i = |\Omega_i/F_i|$. Then by Proposition 4.2.2 the action of $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ on the i-panels has $\prod_{j \neq i} m_j$ orbits. Summing over i,

$$|\mathfrak{P}| = \sum_{i \in I} \prod_{j \neq i} m_j = \sum_{i \in I} \frac{1}{m_i} \cdot \prod_{i \in I} m_i.$$

In the following lemma, with slight abuse of notation, we identify the set $S(g)$ of panels of Δ with the set S of chambers contained in a panel in $S(g)$.

Lemma 4.2.9. *Let $g \in \mathcal{G}(\boldsymbol{F})$ and let S be the set of all chambers contained in some singularity of g. Let $\mathcal{U} = \mathcal{U}(\boldsymbol{F})$. Then the conjugation ${}^g\mathcal{U}_{(S)} = g \cdot \mathcal{U}_{(S)} \cdot g^{-1}$ in $\mathcal{G}(\boldsymbol{F})$ equals $\mathcal{U}_{(g\,.\,S)}$.*

Proof. Since $g\,.\,S(g) = S(g^{-1})$, it suffices by symmetry to show that ${}^g\mathcal{U}_{(S)} \subseteq \mathcal{U}_{(g\,.\,S)}$. Clearly $g\,.\,S$ remains fixed by ${}^g\mathcal{U}_{(S)} = g\,\mathcal{U}_{(S)}\,g^{-1}$ so we only need to check that automorphisms in ${}^g\mathcal{U}_{(S)}$ have no singularities. Let $h \in \mathcal{U}_{(S)}$ and consider an arbitrary panel $\mathcal{P}$. If $\mathcal{P} \in g\,.\,S(g)$, then as observed earlier, $\mathcal{P}$ is fixed by ${}^g h$ and cannot be a singularity of ${}^g h$. Otherwise $\mathcal{P} \notin g\,.\,S(g)$, in which case $g^{-1}\,.\,\mathcal{P}$ is not a singularity of g, and neither is $hg^{-1}\,.\,\mathcal{P}$ since h fixes S. By Lemma 3.1.2,

$$\sigma_\lambda({}^g h, \mathcal{P}) = \sigma_\lambda(g, hg^{-1}\mathcal{P}) \cdot \sigma_\lambda(h, g^{-1}\mathcal{P}) \cdot \sigma_\lambda(g, g^{-1}\mathcal{P})^{-1}$$

is the product of three permutations in F_i (where i is the type of $\mathcal{P}$). Hence $\mathcal{P}$ is not a singularity of ${}^g h$. Since $\mathcal{P}$ was arbitrary, this concludes our proof. ■

Corollary 4.2.10. *$\mathcal{G}(\boldsymbol{F})$ commensurates the compact open subgroups of $\mathcal{U}(\boldsymbol{F})$.*

Proof. Note that in Lemma 4.2.9, the set S is finite. Hence $\mathcal{U}_{(S)}$ and $\mathcal{U}_{(g\,.\,S)}$ are compact open subgroups of $\mathcal{U}(\boldsymbol{F})$. Together with Proposition 1.2.14, the corollary follows. ■

4.3 Topological properties

The topology on $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ is slightly tricky. If we want to lift the topological structure from $\mathcal{U}(\boldsymbol{F})$, we want the embedding $\mathcal{U}(\boldsymbol{F}) \hookrightarrow \mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ to be a continuous open map. We proceed as follows.

Let $\mathfrak{B}$ be the collection of all compact open subgroups of $\mathcal{U}(\boldsymbol{F}) \leq \mathcal{G}(\boldsymbol{F})$. By van Dantzig's theorem (Corollary 1.2.19), $\mathfrak{B}$ is an identity neighbourhood basis for the topology on $\mathcal{U}(\boldsymbol{F})$. Note that $\mathfrak{B}$ is a filter base on $\mathcal{G}(\boldsymbol{F})$ that satisfies all properties in Lemma 1.2.5. Indeed, (i) and (ii) are immediate, and in order to check (iii), take any $U \in \mathfrak{B}$ and $g \in \mathcal{G}(\boldsymbol{F})$. Then by Corollary 4.2.10, the intersection of U and $g \cdot U \cdot g^{-1}$ has finite index in U and is hence again compact and open in $\mathcal{U}(\boldsymbol{F})$. We now obtain a unique, well-defined group topology on $\mathcal{G}(\boldsymbol{F})$ by Lemma 1.2.5.

Definition 4.3.1. We endow $\mathcal{G}(\boldsymbol{F})$ with the group topology described in the above paragraph, and the restricted universal group $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}}) \leq \mathcal{G}(\boldsymbol{F})$ with the subspace topology induced by $\mathcal{G}(\boldsymbol{F})$.

Instead of using the technical Lemma 1.2.5 of Bourbaki (that we included without proof), we refer the reader who prefers a more concrete approach to [GL18, Lemma 8.4].

Note that when $\boldsymbol{F} = \acute{\boldsymbol{F}}$, the topology on $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}}) = \mathcal{U}(\boldsymbol{F})$ agrees with the permutation topology, but in general it does not. Indeed, in our topology on $\mathcal{G}(\boldsymbol{F})$, a subgroup $H \leq \mathcal{G}(\boldsymbol{F})$ is open if and only if H contains a pointwise stabiliser of a finite set of chambers *intersected with* $\mathcal{U}(\boldsymbol{F})$. In other words, the topology on $\mathcal{G}(\boldsymbol{F})$ in Definition 4.3.1 is finer than the permutation topology — in which, for instance, a subgroup $\mathcal{U}(\boldsymbol{F})_c$ would not be open.

Proposition 4.3.2. *Endow $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ with the topology from Definition 4.3.1.*

(i) *The inclusion $\mathcal{U}(\boldsymbol{F}) \hookrightarrow \mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ is continuous and open.*

(ii) *$\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ is a totally disconnected, locally compact group.*

Proof. For $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}}) = \mathcal{G}(\boldsymbol{F})$, claim (i) is a restatement of the definition and (ii) is an immediate consequence. For the general case, both claims follow from the definition of the topology on $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ as the subspace topology induced by $\mathcal{G}(\boldsymbol{F})$. ■

With Proposition 3.3.1, this immediately characterises when the topology on $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ is discrete.

Corollary 4.3.3. *$\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ is discrete if and only if every local group in $\boldsymbol{F}$ is free.*

In the following proposition, we can assume without loss of generality that the local data $\boldsymbol{F}$ and $\acute{\boldsymbol{F}}$ already take into account the rung restriction (Lemma 4.1.5) by asking that $F_i = \acute{F}_i$ for every $i \in I$ that is the type of a rung.

Proposition 4.3.4. *Let Δ be a right-angled building over index set I. Let $\boldsymbol{F}$ and $\boldsymbol{F}'$ be the local data as in Definition 4.1.11, and assume that $F_i = \acute{F}_i$ for every $i \in I$ that is the type of a rung in Δ. Then the closure of the restricted universal group $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ as a subgroup of $\mathrm{Aut}(\Delta)$ is the group $\mathcal{U}(\acute{\boldsymbol{F}})$.*

Proof. Let $g \in \mathcal{U}(\acute{\boldsymbol{F}})$. It suffices to find a sequence of automorphisms $g_n \in \mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$, where $n \in \mathbb{N}$, converging to g in the permutation topology on $\mathrm{Aut}(\Delta)$. Let $c \in \Delta$ be any chamber and, for every $n \in \mathbb{N}$, define the set B_n to be the convex closure of the ball $\mathsf{B}_n(c)$. Note that B_n is panel-closed.

Proposition 3.1.9 yields an automorphism g_n that agrees on B_n with g. Moreover, the only panels where the local action of g_n is not guaranteed to be a permutation in $\boldsymbol{F}$ are the panels parallel to some panel in B_n. But by assumption on the rungs, the parallel classes of singularities are bounded (Proposition 4.1.6), and hence finite. Hence we have $g_n \in \mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$, and this finishes our proof. ■

Corollary 4.3.5. *Let $\boldsymbol{F}$ and $\boldsymbol{F}'$ be the local data as in Definition 4.1.11, and assume that $F_i = \acute{F}_i$ for every $i \in I$ that is the type of a rung in Δ. Then the following are equivalent:*

(i) *$\boldsymbol{F} = \acute{\boldsymbol{F}}$;*

(ii) *$\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}}) = \mathcal{U}(\boldsymbol{F})$;*

(iii) *$\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ is closed in $\mathrm{Aut}(\Delta)$;*

(iv) *chamber stabilisers of $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ are compact;*

(v) *chamber stabilisers of $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ are closed in $\mathrm{Aut}(\Delta)$.*

Proof. Implications (i) $\Rightarrow$ (ii) $\Rightarrow$ (iii) $\Rightarrow$ (iv) $\Rightarrow$ (v) are straightforward. For (iii) $\Rightarrow$ (i), if $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ is closed, then $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}}) = \mathcal{U}(\acute{\boldsymbol{F}})$ by Proposition 4.3.4, and hence $\boldsymbol{F} = \acute{\boldsymbol{F}}$. Finally (v) $\Rightarrow$ (iii) follows from Lemma 1.2.29. ■

Proposition 4.3.6. *$K_{\mathcal{P}}$ and $K_{n,\mathcal{P}}$ are compact open subgroups of $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$, for every $\mathcal{P}$ and n.*

Proof. Since $K_{\mathcal{P}}$ contains the setwise panel stabiliser $\mathcal{U}_{\{\mathcal{P}\}}$ as a subgroup of finite index, it suffices to note that $\mathcal{U}_{\{\mathcal{P}\}}$ is compact and open in $\mathcal{U}(\boldsymbol{F})$ (as it is a union of finitely many cosets of chamber stabilisers), and hence in $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ as well. The same argument works for $K_{n,\mathcal{P}}$ once we note that the set $\{c \in \Delta \mid \mathrm{dist}(c, \mathcal{P}) \leq n\}$ contains only finitely many panels. ■

An immediate corollary is the following fact.

Corollary 4.3.7. *The group $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ is compactly generated.*

Proof. Combining Propositions 4.2.8 and 4.3.6, we see that $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ is generated by a finite number of compactly generated subgroups. ■

From Proposition 4.3.6 also follows that panel stabilisers, although not compact, are still quite close to being compact, in the following precise sense.

Definition 4.3.8 (regionally compact). A topological group is said to be *regionally compact* if it is the increasing union of a family of compact open subgroups.

In the literature, regionally compact groups are also called *locally elliptic*. The authors of [CRW19, Remark 1.0.1] argue why this terminology is best avoided: "locally" is suggestive of a property to be satisfied in some identity neighbourhood basis, which is not the case here.

Proposition 4.3.9. *Let $\mathcal{P}$ be a panel of Δ. Then the stabiliser $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})_{\{\mathcal{P}\}}$ is regionally compact.*

Proof. Since we can write

$$\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})_{\{\mathcal{P}\}} = \overrightarrow{\bigcup_{n\in\mathbb{N}}} K_{n,\mathcal{P}}$$

the result follows immediately from Proposition 4.3.6. ■

4.4 Simplicity

Recall for general groups $G \leq \mathrm{Aut}(\Delta)$ our Theorem 3.4.9, a criterion for G^+ to be simple. One of the assumptions of Theorem 3.4.9 is that G is closed. Unfortunately, by Corollary 4.3.5, this only occurs in degenerate boundary cases in the context of restricted universal groups. Yet we can still show that $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ is *virtually* simple under the same conditions for $\mathcal{U}(\acute{\boldsymbol{F}})$ to be simple. The proof idea is due to Pierre-Emmanuel Caprace, Colin Reid, and Phillip Wesolek ([CRW19]).

Proposition 4.4.1. *$\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ satisfies the independence property.*

Proof. This is almost exactly the same argument as in Proposition 3.4.3, the independence property for universal groups, taking into account the running assumption that the parameters are all finite. Indeed, identify $g \in \prod_{c\in\mathcal{P}} V_i(c)$ with its image in $\mathrm{Aut}(\Delta)_{(\mathcal{T})}$. Every panel of Δ is either contained in $\mathcal{T}$ or in *one of the finitely many* i-wings with base chamber in $\mathcal{T}$. Hence, every nontrivial local action of g agrees with a local action of an element in $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$, and g has in total a finite number of singularities. In other words, $g \in \mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$. ■

We can recycle the argument from Proposition 3.4.8 to show that, in the nondiscrete case, $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ is monolithic.

Proposition 4.4.2. *Assume not all local groups F_i are free. Then $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ is monolithic; the monolith is the subgroup generated by all tree-wall fixators and is simple.*

Proof. Write $M = \langle \mathcal{G}_{(\mathcal{T})} \rangle$, where $\mathcal{T}$ ranges over all tree-walls. Note that M is nontrivial, since not all local groups act freely. Thanks to Propositions 4.2.5 and 4.4.1, we may apply Proposition 3.4.7 to obtain that every normal subgroup of $\mathcal{G}$ contains M. Hence $\mathcal{G}$ is monolithic with monolith M.

For the simplicity of M, let $N \trianglelefteq M$ be the intersection of all nontrivial normal subgroups of M. A characteristic subgroup of M, it follows that N is a normal subgroup of $\mathcal{G}$. By Proposition 3.4.7 again, N contains the monolith M, so that in fact $N = M$. This shows that M is simple. ■

As an interesting corollary of Proposition 4.4.2 and using the running assumption that the building is locally finite, we then obtain the following.

Proposition 4.4.3. *Assume that the diagram of Δ is not ladderfull. Then $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ is topologically simple if and only if it is abstractly simple.*

Proof. If every local group F_i acts freely, then $\mathcal{G}$ is discrete and there is nothing to prove.

Otherwise, $\mathcal{G}$ is monolithic by Proposition 4.4.2 and the monolith is generated by tree-wall fixators. By the assumption on the diagram, there exists at least one index $i \in I$ such that an i-tree-wall $\mathcal{T}$ contains only finitely many chambers. Then the fixator $\mathcal{G}_{(\mathcal{T})}$ — as the intersection of finitely many chamber stabilisers — is open. Consequently, the monolith of $\mathcal{G}$ is open, and so is every nontrivial normal subgroup of $\mathcal{G}$. ∎

We again assume without loss of generality that $F_i = \acute{F}_i$ for every rung type $i \in I$.

Theorem 4.4.4. *Let Δ be a thick irreducible right-angled building over index set I. Let $\boldsymbol{F}$ and $\acute{\boldsymbol{F}}$ be the local data as in Definition 4.1.11. Assume that $F_i = \acute{F}_i$ for every $i \in I$ that is the type of a rung. Moreover assume that not all local groups $\acute{F}_i$ are free.*

Then the restricted universal group $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ is virtually simple if and only if $\acute{F}_i$ is generated by point stabilisers for every $i \in I$ and transitive for every i in some vertex cover of the diagram of Δ.

Proof. Abbreviate $\mathcal{G} = \mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ and $\mathcal{U} = \mathcal{U}(\acute{\boldsymbol{F}})$.

First, suppose that $\mathcal{G}$ has a simple subgroup M of finite index. We can assume that M is a normal subgroup by Lemma 1.1.11. By Proposition 4.3.4, $\mathcal{G}$ is dense in $\mathcal{U}$. Since $\mathcal{U}^+$ is a simple nondiscrete open subgroup of $\mathcal{U}$ by Corollary 3.4.10 and Proposition 3.3.1, it follows that $\mathcal{U}^+$ is contained in $\overline{M}$. Then $N = \mathcal{U}^+ \cap \mathcal{G}$ is a nontrivial normal subgroup of $\mathcal{G}$ and intersects M nontrivially. Since M is simple, we obtain that $M \leq N$, hence N has finite index in $\mathcal{G}$. Taking the closure, it then follows that $\mathcal{U}^+$ has finite index in $\mathcal{U}$, and the characterisation follows from Theorem 3.2.20.

Conversely, suppose the local data $\acute{\boldsymbol{F}}$ satisfies the assumptions postulated. Then by Corollary 3.4.11, $\mathcal{U}$ is simple. Let $M = \langle \mathcal{G}_{(\mathcal{T})} \rangle$ be the open normal subgroup of $\mathcal{G}$ generated by all fixators of tree-walls $\mathcal{T}$ of Δ. By Proposition 4.4.2, M is a simple group, and the monolith of $\mathcal{G}$. We will show M to have finite index in $\mathcal{G}$.

The closure of M is a normal subgroup of $\mathcal{U}(\acute{\boldsymbol{F}})$. Hence, M is dense in $\mathcal{U}$, and consequently in $\mathcal{G}$ as well. This implies that the M-orbits and $\mathcal{G}$-orbits on the building Δ agree: for any $c \in \Delta$ and $g \in \mathcal{G}$ the stabiliser $\mathcal{G}_c$ is open, hence the intersection $M \cap g \cdot \mathcal{G}_c$ is nonempty, and any automorphism $h \in M$ in the intersection satisfies $h \,.\, c = g \,.\, c$. It follows that $\mathcal{G} = \mathcal{G}_c \cdot M$.

Let $\mathcal{P}$ be a panel of Δ. By the previous paragraph, $\mathcal{G} = \mathcal{G}_{\{\mathcal{P}\}} \cdot M$. Together with Proposition 4.3.9, we obtain an increasing union

$$\mathcal{G} = \overrightarrow{\bigcup_{n \in \mathbb{N}}} H_n \cdot M,$$

where the $H_n \leq \mathcal{G}_{\{\mathcal{P}\}}$ are compact open subgroups. On the other hand, $\mathcal{G}$ is compactly generated by Corollary 4.3.7. The family $\{H_n \cdot M\}_{n \in \mathbb{N}}$ defines an open cover of any compact generating set and hence $\mathcal{G} = H_n \cdot M$ for some $n \in \mathbb{N}$.

In conclusion, M is both open and cocompact in $\mathcal{G}$. The coset space $\mathcal{G}/M$ being both compact and discrete, it follows that indeed M has finite index in $\mathcal{G}$, which concludes our proof. ∎

5 Open questions

> *Well. What can we do, except try to do better?*
>
> — Joe Abercrombie, *The Blade Itself*

Mathematical research isn't all sunshine and rainbows. And even with the successful completion of one theorem, five new questions arise. In this concluding chapter, we mention a couple of research problems that we encountered thoughout this thesis and did not have the time or ingenuity for.

5.1 Geometrical vs. combinatorial density

Recall that we introduced three definitions for an action of a group to be *minimal*, depending on whether G acts on a tree (Definition 1.4.1), on a building (Definition 2.4.4), or on a $\mathrm{CAT}(0)$ space (Definition 1.6.7). We did show in Proposition 2.4.5 that minimality on a right-angled building can be characterised in terms of minimality on its tree-wall trees.

However, we also noted in Remark 2.4.6 that the minimality of a group $G \leq \mathrm{Aut}(\Delta)$ acting on a right-angled building Δ is by no means equivalent to minimality of the induced action on its Davis realisation. Hence the following natural question, trying to bridge Definitions 1.6.7 and 2.4.4.

Question 5.1.1. *Let Δ be any right-angled building, let $\mathbb{K}(\Delta)$ be its Davis realisation, and let G act on Δ by automorphisms. Can we characterise when the induced action on $\mathbb{K}(\Delta)$ is minimal in terms of the action on Δ?*

Similarly, we presented in the same Definition 2.4.4 a competing notion for a group action of G on a right-angled building to be *dense*. On the one hand, more useful to our study was the requirement that G leaves no point at infinity of any tree-wall tree invariant — a property we called *combinatorial density*. On the other hand, the more traditional approach would involve the $\mathrm{CAT}(0)$ Davis realisation. By Remark 2.4.6 again, one does not want to require the induced action on the Davis realisation to be geometrically dense in the sense of Definition 1.6.7, but a suitable alternative might be the following.

Definition 5.1.2. Let G be a group acting on a right-angled building Δ by automorphisms. We call the action *geometrically dense* if it is minimal on Δ and if moreover the induced action on the Davis realisation $\mathbb{K}(\Delta)$ has no fixed point at infinity in the boundary $\partial\mathbb{K}(\Delta)$.

Question 5.1.3. *Let Δ be any right-angled building, let $\mathbb{K}(\Delta)$ be its Davis realisation, and let G act on Δ by automorphisms. Is is true that the action of G on Δ is combinatorially dense if and only if it is geometrically dense in the sense of Definition 5.1.2?*

The main hindrance for answering Question 5.1.3 seems to be the following. Any minimal infinite gallery γ in Δ (by which we mean that its finite subgalleries are all minimal) induces a path in the i-tree-wall, that may or may not be infinite, depending on whether γ eventually stays in a single residue of type $\{i\} \cup \{i\}^{\perp}$ or of type $I \setminus \{i\}$, or not. Conversely, it is not too hard to "lift" an infinite path in the i-tree-wall to a gallery in Δ that induces it. However, it is no longer straightforward to try a similar conversion strategy in the CAT(0) realisation. For any point in $\partial\mathbb{K}(\Delta)$, we can pick a representative geodesic in $\mathbb{K}(\Delta)$ and read off a gallery of Δ from the visited Davis chambers and vertices of the simplicial complex, but the converse requires the construction of infinite geodesics in $\mathbb{K}(\Delta)$ from discrete approximations.

To illustrate the difficulties, we mention only one recent result of Timothée Marquis, who studied this correspondence between infinite minimal galleries and geodesic rays in [Mar19].

Definition 5.1.4. Let Δ be a right-angled building.

(i) For any chamber c, denote by $c^{(0)}$ the barycenter of the Davis chamber $\mathbb{K}(c)$ in $\mathbb{K}(\Delta)$.

(ii) For any minimal gallery γ in Δ, denote by $\gamma^{(0)}$ the piecewise geodesic path in $\mathbb{K}(\Delta)$ joining the points $c^{(0)}$ where c runs over the consecutive chambers of γ.

(iii) For any $c \in \Delta$ and $\eta \in \partial\mathbb{K}(\Delta)$, consider the minimal galleries γ starting in c such that $\gamma^{(0)}$ lies at bounded Hausdorff distance from some geodesic ray pointing to η. Define the *geodesic ray bundle* $\operatorname{Geo}(c, \eta)$ as the union of all barycenters $d^{(0)}$ with d on such a minimal gallery γ.

Theorem 5.1.5. *Let Δ be a locally finite, hyperbolic building, let $c_1, c_2 \in \Delta$, and let $\eta \in \partial\mathbb{K}(\Delta)$. Then the symmetric difference of the geodesic ray bundles $\operatorname{Geo}(c_1, \eta)$ and $\operatorname{Geo}(c_2, \eta)$ is finite.*

Proof omitted. This is [Mar19, Theorem A].

Note that hyperbolicity is a nontrivial assumption — a result of Moussong ([Mou88, Theorem 17.1]) implies that a right-angled building is hyperbolic precisely when its diagram does not contain the rank four subdiagram below (the Davis realisation of which is the Euclidean plane).

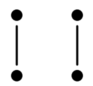

5.2 Compact generation of universal groups

The most jarring open result of this thesis must be our Conjecture 3.3.9, on the characterisation of when the universal group $\mathcal{U}(\boldsymbol{F})$ is compactly generated. Let us repeat the problem.

Question 5.2.1. *Assume that $\mathcal{U}(\boldsymbol{F})$ is closed, locally compact, and compactly generated. Does it then follow that every local group F_i is compactly generated?*

As we noted, Question 5.2.1 is the keystone to finishing the claim that a closed and locally compact universal group is compactly generated if and only if all local groups are compactly generated and have finitely many orbits.

Let us briefly and intuitively sketch one unsuccessful way to tackle the problem, using a variant of our city product construction of Section 2.7. We can glue together right-angled buildings over the index sets $J_1 \sqcup K$ and $J_2 \sqcup K$ along a tree, while amalgamating their common residues of type K. More precisely, assume the index set of a building admits a nontrivial partition $J_1 \sqcup J_2 \sqcup K$, such that $m_{j_1 j_2} = \infty$ for all $j_1 \in J_1$ and $j_2 \in J_2$. Then define the semiregular graph with vertex set

$$\operatorname{Res}_{J_1 \sqcup K}(\Delta) \cup \operatorname{Res}_{J_2 \sqcup K}(\Delta)$$

and declare a residue of type $J_1 \sqcup K$ to be adjacent to a residue of type $J_2 \sqcup K$ if and only if they share a common residue of type K. This defines a semiregular tree that we denote by T.

Next, for every $J \subsetneq I$, an action of $G \leq \mathrm{Aut}(\Delta)$ on the chambers of Δ also induces an action on the set $\mathrm{Res}_J(\Delta)$ of all J-residues. Let us assume that Δ is irreducible, so that the action is faithful by Proposition 3.2.13. We denote the induced permutation group by $G \curvearrowright_J \leq \mathrm{Sym}(\mathrm{Res}_J(\Delta))$. Then we have an abstract isomorphism $G \cong G \curvearrowright_J$ by faithfulness.

Endowing the two groups with the permutation topology, it is not too hard to show that this is, in fact, an isomorphism of topological groups. In particular, $\mathcal{U}_\Delta(\boldsymbol{F})$ is compactly generated if and only if $\mathcal{U}_\Delta(\boldsymbol{F}) \curvearrowright_J$ is.

Going back to the tree T, and using similar techniques as in Section 2.7, we can proceed to show that $\mathcal{U}_\Delta(\boldsymbol{F}) \curvearrowright_K$ is isomorphic to a *subgroup* of

$$\mathcal{U}_T\Big(\mathcal{U}_{\mathcal{R}_1}(\boldsymbol{F}|_{J_1 \cup K}) \curvearrowright_K, \mathcal{U}_{\mathcal{R}_2}(\boldsymbol{F}|_{J_2 \cup K}) \curvearrowright_K\Big), \tag{$*$}$$

where $\mathcal{R}_1$ and $\mathcal{R}_2$ are residues of type $J_1 \cup K$ and type $J_2 \cup K$, respectively. It is only a subgroup, because $\mathcal{U}_\Delta(\boldsymbol{F})$ imposes extra "compatibility conditions" on adjacent residues of type $J_1 \sqcup K$ and $J_2 \sqcup K$ — the local actions have to match on the amalgamated K-residues. In $(*)$, there is no such restriction.

If we would be able to show that the group in $(*)$ is compactly generated from the assumption that its subgroup $\mathcal{U}_\Delta(\boldsymbol{F}) \curvearrowright_K$ is compactly generated, then we could use Corollary 3.3.13 to obtain that $\mathcal{U}_{\mathcal{R}_1}$ and $\mathcal{U}_{\mathcal{R}_2}$ are compactly generated and make an inductive argument work. We do need that Δ is irreducible with index type that admits a partition $J_1 \sqcup J_2 \sqcup K$, but this is not an obstruction: we can consider an edge $\{j_1, j_2\}$ in the diagram and set $J_1 = \{j_1\}$ and $J_2 = \{j_2\}$, and with some more care, ensure that $I \setminus \{j_1\}$ and $I \setminus \{j_2\}$ remain irreducible.

Unfortunately however, it appears that showing $(*)$ to be compactly generated is about just as hard as showing $\mathcal{U}_\mathcal{R}(F|_K)$ to be compactly generated (with $\mathcal{R}$ a residue of type K).

We note that a partition $\{j_1\} \sqcup \{j_2\} \sqcup K$ of the index set of Δ, with $m_{j_1 j_2} = \infty$, corresponds to what Haglund and Paulin call a *scindement* (French) in [HP03].

5.3 Tidiness and the scale function

George Willis introduced in [Wil94] the *scale function* and the concept of *tidy subgroups* for t.d.l.c. groups. The motivation was the following question of Karl Hofmann. As a topological analogue of torsion, call an element $g \in G$ of a locally compact group *periodic* if $\langle g \rangle$ has compact closure in G. The subset $\mathrm{Per}(G)$ of all periodic elements need not be closed if G is connected, but using the scale function, Willis could demonstrate $\mathrm{Per}(G)$ to be closed if G is t.d.l.c.

Since then, the notions have been used to define more structural invariants of the group (such as the space of directions and maximal scale-multiplicative semigroups) with applications in various other areas (such as random walks, ergodic theory, and dynamical systems).

The definitions are rather technical, and not easy to compactly motivate; we refer to the literature for more details and background. In this section G always denotes a locally compact group.

Definition 5.3.1 (scale function, tidy). We define the *scale* of an automorphism $\alpha \in \mathrm{Aut}(G)$ as

$$s(\alpha) = \min_V \big[\alpha(V) : \alpha(V) \cap V\big],$$

taken over all compact open subgroups $V \leq G$. A subgroup attaining the minimum is *tidy* for α.

There is a *tidying procedure* that takes as input any compact open subgroup V and produces a tidy subgroup for an automorphism. With this procedure one can prove the following characterisation of tidiness in terms of structural properties of V only.

Definition 5.3.2. For a compact open subgroup $V \leq G$ and $\alpha \in \mathrm{Aut}(G)$, define

$$V_+ = \bigcap_{n \geq 0} \alpha^n(V), \qquad V_- = \bigcap_{n \leq 0} \alpha^n(V), \qquad V_{++} = \bigcup_{n \geq 0} \alpha^n(V_+), \qquad V_{--} = \bigcup_{n \leq 0} \alpha^n(V_-).$$

Note that V_{++} and V_{--} are again subgroups of G (being increasing unions of subgroups).

Theorem 5.3.3. *A compact open subgroup V is tidy for the automorphism α if and only if V_{--} is closed ("tidy below") and $V = V_+V_- = V_-V_+$ ("tidy above").*

Proof omitted. We refer to [Wil01].

In his recent paper [Byw19], Timothy Bywaters used this tidying procedure to explicitly calculate the scale function on the Le Boudec groups $\mathcal{G}(F, F')$ of Section 1.5.3. Among other things he was also able to describe the space of directions of $\mathcal{G}(F, F')$ in terms of the action on the tree and could construct maximal scale-multiplicative semigroups using this space of directions.

We highlight one example. It is not too hard to show that all elliptic automorphisms $g \in \mathcal{G}(F, F')$ are uniscalar (i.e. $s(g) = 1$). In fact, the tree analogue of some compact open subgroup $K_{n,\mathcal{P}}$ (as in our Definition 4.2.6) is tidy for g — see [Byw19, Proposition 3.1] for the precise statement. The following theorem by Bywaters characterises conversely when the elliptic automorphisms are the only uniscalar elements.

Theorem 5.3.4. *In $\mathcal{G}(F, F')$, the subset of uniscalar elements equals the subset of elliptic elements if and only if F has distinct point stabilisers.*

Proof omitted. This is [Byw19, Corollary 3.23].

Bywaters studied far more and also considered the space of directions of $\mathcal{G}(F, F')$ and associated scale-multiplicative semigroups, but it would lead us too far to introduce these concepts here.

The calculations in [Byw19] quickly grow quite complicated. Moreover, they rely on the notion of a *pando* for a hyperbolic element $g \in \mathcal{G}(F, F')$, which is a specific subtree associated to g.

Question 5.3.5. *Can the results in [Byw19] be generalised to our setting of restricted universal groups over right-angled buildings? In particular, what can we say about the behaviour of the scale function on $\mathcal{G}_\Delta(\boldsymbol{F}, \acute{\boldsymbol{F}})$?*

5.4 Restricted universal groups over locally infinite buildings

Recall from Section 4.1 that we assumed the building Δ to be locally finite, since even a innocuous fundamental result like Corollary 4.1.10 gets far more technical without this assumption. In fact, it is not so clear what the "best" generalisations of Le Boudec's Definition 1.5.10 would be.

A promising candidate might be the following:

$$\mathcal{G}^\lambda_\Delta(\boldsymbol{F}) = \big\{ g \in \mathrm{Aut}(\Delta) \;\big|\; \text{the set of all } i\text{-panels } \mathcal{P} \text{ with } \sigma_\lambda(g, \mathcal{P}) \notin F_i \text{ is bounded in } \Delta \big\}.$$

With this definition, Lemma 4.1.5 would still be valid and the existence of ladders would still impose restrictions. However the independence property Proposition 4.4.1 would not hold anymore: one could construct an automorphism that has in every wing around a central tree-wall $\mathcal{T}$ a bounded set of singularities, but in such a way that the distances between these bounded sets and $\mathcal{T}$ grows arbitrarily large.

In any case, it goes without saying that the topological structure of such a locally infinite variation would again be quite subtle.

Question 5.4.1. *What can we say about the family of groups $\mathcal{G}^{\lambda}_{\Delta}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ in the locally infinite case?*

5.5 Lattices in universal groups

One research question the author would have liked to study is the existence of lattices in universal and restricted universal groups over locally finite right-angled buildings. We refer to Section 1.5.3 for some more motivation.

Question 5.5.1. *What can we say about (cocompact) lattices in $\mathcal{U}^{\lambda}_{\Delta}(\boldsymbol{F})$?*

Question 5.5.2. *What conditions can we impose on the local groups that guarantee that $\mathcal{G}^{\lambda}_{\Delta}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ does not admit lattices? How does the combinatorics of the diagram relate to the existence of lattices? Can we use similar constructions as in [LB16] to construct other new interesting topological groups?*

We mention that Anne Thomas in [Tho06] studied lattices in the full automorphism group $\mathrm{Aut}(\Delta)$ of a right-angled building and showed that they share many properties with tree lattices.

5.6 Spheromorphisms of buildings

A final, very broad research question is whether there is an interesting generalisation of Lederle's results in Section 1.5.4. The hands-on definition of the Neretin group (or a Burger–Mozes variant) involves almost-automorphisms of trees, and it is not immediately clear how a general right-angled building analogue would look like.

However, given a group G acting on a topological space, there is a general framework to rigorously define a group of homeomorphisms that "locally look like" elements of G.

Definition 5.6.1. Let G be a group acting on a topological space X. Then the *topological full group* (with respect to this action) is the group of all homeomorphisms $\varphi\colon X \to X$ such that for every $x \in X$ there is a group element $g \in G$ and an open neighbourhood U where $\varphi|_U = g|_U$.

Lederle established the following lemma; here we reuse our notation $\mathrm{A}G$ from Section 1.5.4 for the group of equivalence classes of G-almost-automorphisms.

Lemma 5.6.2. *Let T be a regular tree and let $G \leq \mathrm{Aut}(T)$. Then the group $\mathrm{A}G$ is isomorphic to the topological full group of G acting on ∂T.*

Proof omitted. This is [Led17, Lemma 2.19].

We could hence quite straightforwardly define the topological full group of $\mathcal{U}_{\Delta}(\boldsymbol{F})$ with its action on the boundary $\partial\mathbb{K}(\Delta)$ of the Davis realisation (equipped with a suitable topology). Let us denote this group by $\mathcal{N}_{\Delta}(\boldsymbol{F})$. We remark again, however, that the interplay between the combinatorics of the building and the geometry of its $\mathrm{CAT}(0)$ realisation is not quite trivial.

Question 5.6.3. *What can we say about $\mathcal{N}_{\Delta}(\boldsymbol{F})$? How does the combinatorics of the diagram relate to $\mathcal{N}_{\Delta}(\boldsymbol{F})$? Can we find conditions on the local data that prohibit $\mathcal{N}_{\Delta}(\boldsymbol{F})$ from admitting (cocompact) lattices, similar to Theorem 1.5.19 (iii)? Can we establish similar simplicity results?*

Conclusions

> *If we do nothing with the knowledge we gain, then we have wasted our study. Books can store information better than we can; what we do that books cannot is interpret. So if one is not going to draw conclusions, then one might as well just leave the information in the texts.*
>
> — Brandon Sanderson, *The Way of Kings*

Recall the general goal that we sketched in the introduction: the goal of the original Burger–Mozes construction was to provide families of topological groups, the global structure of which depends on the local data. In the variations by Smith, Le Boudec, and Lederle, this local-to-global principle remained a central anchor point. The results of Silva showed that right-angled buildings allow for a fruitful generalisation, in which the global structure of the group not only depends on local data, but also on the diagram of the building. Silva assumed the local groups to be finite and transitive, while we explicitly did not.

In this concluding section, we present a brief summary of our results for general local groups (that are not necessarily transitive or finite). Here we will assume our building to be irreducible for sake of convenience; we refer to Chapter 3 for more precise statements.

Theorem. *Let M be an irreducible diagram over an index set I. Let $\boldsymbol{F}$ be a collection of permutation groups $F_i \leq \mathrm{Sym}(\Omega_i)$ with $3 \leq |\Omega_i|$, indexed by $i \in I$. Let Δ be a semiregular right-angled building of type M with parameters equal to $|\Omega_i|$ for every $i \in I$. Equip the building Δ with a legal colouring λ taking values in the sets Ω_i.*

Let $\mathcal{U}_\Delta^\lambda(\boldsymbol{F}) \leq \mathrm{Aut}(\Delta)$ be the corresponding universal group, equipped with the permutation topology. We then have the following.

- *$\mathcal{U}(\boldsymbol{F})$ is transitive on the residues of type J if and only if F_i is transitive for every $i \in I \setminus J$.*
  This is Proposition 3.2.2.
- *$\mathcal{U}(\boldsymbol{F})$ is primitive on the residues of type J if and only if $I = J \sqcup \{k\}$ for some $k \in I$, and F_k is primitive and nonregular, and F_i is transitive for all $i \in I \setminus k^\perp$.*
  This is Theorem 3.2.15.
- *$\mathcal{U}(\boldsymbol{F})$ is generated by chamber stabilisers if and only if F_i is generated by point stabilisers for every $i \in I$ and transitive for every i in some vertex cover of M.*
  This is Theorem 3.2.20.
- *$\mathcal{U}(\boldsymbol{F})$ is totally disconnected.*
  This is Proposition 1.2.27.
- *$\mathcal{U}(\boldsymbol{F})$ is discrete if and only if F_i acts freely on Ω_i for every $i \in I$.*
  This is Proposition 3.3.1.
- *$\mathcal{U}(\boldsymbol{F})$ is closed in $\mathrm{Aut}(\Delta)$ if and only if F_i is closed in $\mathrm{Sym}(\Omega_i)$ for every $i \in I$.*
  This is Proposition 3.3.3.

- *Assuming that $\mathcal{U}(\boldsymbol{F})$ is closed, then $\mathcal{U}(\boldsymbol{F})$ is locally compact if and only if, for every $i \in I$, all suborbits of F_i are finite.*
  This is Proposition 3.3.4.
- *Assuming that Conjecture 3.3.9 is true, and assuming that $\mathcal{U}(\boldsymbol{F})$ is closed and locally compact, then $\mathcal{U}(\boldsymbol{F})$ is compactly generated if and only if F_i is compactly generated and has only finitely many orbits for every $i \in I$.*
  This is a rephrasing of Theorems 3.3.6 and 3.3.7 and Conjecture 3.3.9.
- *Assuming that $\mathcal{U}(\boldsymbol{F})$ is closed and nondiscrete, then $\mathcal{U}(\boldsymbol{F})^+$ is simple.*
  This is Corollary 3.4.10.

In particular, if we pick closed local groups F_i such that every point stabiliser has finite orbits, then $\mathcal{U}(\boldsymbol{F})$ is a totally disconnected locally compact group. If at least one point stabiliser is nontrivial, then $\mathcal{U}(\boldsymbol{F})$ is nondiscrete and the subgroup generated by chamber stabilisers is a simple group.

In Chapter 4, we extended the construction by Le Boudec to the setting of right-angled buildings and argued why we want to restrict to locally finite buildings for our initial study. We find that the diagram may enforce additional restrictions on the local groups, due to the existence of *ladders*. A brief summary of our results follows below.

Theorem. *Let M be an irreducible diagram over an index set I. Let $\boldsymbol{F}$ and $\acute{\boldsymbol{F}}$ be two collections of permutation groups $F_i \leq \acute{F}_i \leq \widehat{F}_i \leq \mathrm{Sym}(\Omega_i)$ with $3 \leq |\Omega_i| < \infty$, indexed by $i \in I$. Let Δ be a semiregular right-angled building of type M with parameters equal to $|\Omega_i|$ for every $i \in I$. Equip Δ with a legal colouring λ taking values in the sets Ω_i. We stress that Δ is assumed to be locally finite.*

Let $\mathcal{G}_\Delta^\lambda(\boldsymbol{F}, \acute{\boldsymbol{F}}) \leq \mathrm{Aut}(\Delta)$ be the corresponding restricted universal group, equipped with the topology from Definition 4.3.1. We then have the following.

- *$\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ is transitive on the residues of type J if and only if F_i is transitive for every $i \in I \setminus J$.*
  This is Proposition 4.2.2.
- *$\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ is totally disconnected and locally compact.*
  This is Proposition 4.3.2.
- *$\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ is discrete if and only if F_i acts freely on Ω_i for every $i \in I$.*
  This is Corollary 4.3.3.
- *$\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ is compactly generated.*
  This is Corollary 4.3.7.
- *Assuming that $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ is nondiscrete and that $F_i = \acute{F}_i$ for every rung $i \in I$ of some ladder, then $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ is virtually simple if and only if $\mathcal{U}(\acute{\boldsymbol{F}})$ is simple — or explicitly, $\acute{F}_i$ is generated by point stabilisers for every $i \in I$ and transitive for every i in some vertex cover of M.*
  This is Theorem 4.4.4.

We note there is again a lot of freedom in the choice of local groups, and that it should be possible to further generalise results of [LB16] to right-angled buildings.

Nederlandstalige samenvatting

> *I have learned all kinds of things from my many mistakes. The one thing I never learn is to stop making them.*
>
> — Joe Abercrombie, *Last Argument of Kings*

Historische context

Het centrale thema van deze thesis vindt zijn oorsprong in een gelauwerd artikel van Marc Burger en Shahar Mozes, [BM00a], over de lokale en globale structuur van groepen die werken op bomen. De groepen in kwestie zijn uitgerust met een natuurlijke topologie onder dewelke de groep volledig onsamenhangend en lokaal compact is.

De theorie van topologische groepen is bijzonder ruim en kent dan ook verscheidene raakgebieden en toepassingen binnen de wiskunde en daarbuiten. Vaak beperkt men zich tot *lokaal compacte* groepen, waar de eenheid een compacte omgeving heeft. De samenhangscomponent van de eenheid van zo'n lokaal compacte groep is steeds een normaaldeler en de corresponderende quotiëntgroep een volledig onsamenhangende groep. De studie van deze groepen valt dan ook min of meer uiteen in twee deelgebieden — *samenhangende* lokaal compacte groepen enerzijds, en *volledig onsamenhangende* lokaal compacte groepen anderzijds.

Samenhangende lokaal compacte groepen zijn reeds sinds de jaren 50 goed begrepen te noemen. Resultaten van Andrew Gleason, Deane Montgomery, Leo Zippin en Hidehiko Yamabe in verband met het vijfde probleem van Hilbert geven een bevredigende structurele karakterisatie en drukken uit dat deze groepen, grof gezegd, benaderd kunnen worden door middel van Liegroepen (Stelling 1.2.17 in deze thesis).

Volledig onsamenhangende lokaal compacte groepen daarentegen zijn veel minder goed begrepen, en een gelijkaardige algemene karakterisatie lijkt niet meteen binnen handbereik. Gedurende lange tijd was de enige gekende algemene structurele eigenschap een stelling van van Dantzig uit 1936, die het bestaan van open compacte deelgroepen stipuleert (Stelling 1.2.19). Pas in de jaren 90 werd er verdere vooruitgang geboekt met nieuwe technieken van George Willis, zoals een schaalfunctie, een theorie van nette deelgroepen, en een ruimte van richtingen geassocieerd aan de groep.

De vroege resultaten van van Dantzig legden de focus op technieken om de globale structuur van een volledig onsamenhangende lokaal compacte groep te bestuderen aan de hand van de structuur van diens compacte open deelgroepen. Eén voorbeeld is een resultaat van Willis, dat stelt dat als de globale groep topologisch enkelvoudig en compact voortgebracht is, een open compacte deelgroep dan nooit oplosbaar kan zijn ([Wil07]). Ook de resultaten van Burger en Mozes vallen onder deze categorie, en bestuderen de globale structuur van groepen van automorfismen van bomen via hun lokale acties.

Hun meest iconische constructie is de *universele groep* $\mathcal{U}(F)$, in functie van een nog vrij te kiezen eindige permutatiegroep F die de *lokale groep* wordt genoemd. Deze universele groep bestaat uit alle automorfismen van een reguliere boom die zich lokaal gedragen als permutaties in deze lokale groep F (we geven een precieze definitie in Hoofdstuk 1.5.1). Deze groepen worden voorzien van de permutatietopologie. De constructie van Burger en Mozes leidt aldus tot een uitgebreide familie volledig onsamenhangende lokaal compacte groepen, die daarenboven compact voortgebracht zijn en onder milde voorwaarden op F ook niet-discreet. Bovendien voldoen deze universele groepen nog eens aan een even iconische “onafhankelijkheidsvoorwaarde” van Jacques Tits, waarmee men kan aantonen dat ze een enkelvoudige deelgroep van index twee bezitten.

Automorfismegroepen van bomen zijn om meerdere redenen aanlokkelijk. Zo kunnen lokaal compacte groepen steeds worden uitgerust met de zogenaamde *Haarmaat*, dat aan deelverzamelingen een invariant volume toekent. In de globale groep kan men dan zoeken naar *roosters*: een discrete deelgroep $\Gamma \leq G$ waarvoor de quotiëntruimte Γ/G een eindig invariant volume heeft. Zoals Burger en Mozes opmerkten, zijn er heel wat interessante roosters te vinden in de automorfismegroepen van producten van bomen. Hyman Bass en Alexander Lubotzky vonden ook interessante resultaten in automorfismegroepen van enkele bomen. We verwijzen naar [BM00a, Car02].

Een ander voordeel is dat bomen en hun automorfismen zeer flexibel zijn. In [Smi17] stelde Simon Smith een variatie voor op de constructie van Burger en Mozes, die gebruikmaakt van twee lokale permutatiegroepen en een semireguliere boom. Dankzij die variant kon Smith als eerste een overaftelbare familie van niet-isomorfe, niet-discrete, enkelvoudige, compact voortgebrachte, volledig onsamenhangende lokaal compacte groepen construeren. We geven meer details in Sectie 1.5.2.

Ook Adrien Le Boudec wist met een eenvoudige twist een nieuwe familie aan groepen op te stellen ([LB16]). Hij liet in de universele groepen van Burger en Mozes een eindig aantal singulariteiten toe, waar de lokale actie de voorgeschreven permutatiegroep niet hoeft te volgen. Zijn constructie leidde tot nieuwe topologische groepen zonder roosters. We geven meer details in Sectie 1.5.3.

Waltraud Lederle volgde een nog drastischere aanpak en bestudeerde varianten op de groepen van Burger–Mozes die iets weg hebben van de Neretingroep. Deze groepen bestaan uit *sferomorfismen* van een boom, waarbij (op een precieze manier) stukken uit de boom weggeknipt mogen worden. Opnieuw waren interessante topologische groepen zonder roosters het resultaat. We geven meer details in Sectie 1.5.4.

Al deze constructies hebben als terugkerend thema dat de uiteindelijke structuur van de verkregen groep in sterke mate afhangt van de gekozen lokale data.

Op suggestie van Pierre-Emmanuel Caprace bestudeerde Ana Silva nog een andere veralgemening van de Burger–Mozesgroepen, waar de bomen worden veralgemeend naar *rechthoekige gebouwen*. Gebouwen zijn, heel grof gezegd, meetkundige structuren die gecoördinatiseerd kunnen worden met een Coxetergroep. Gebouwentheorie overkoepelt onder andere projectieve en affiene ruimtes, veralgemeende veelhoeken, en oneindige bomen. Intuïtief kan men stellen de voornaamste reden dat de constructie van Burger en Mozes (en alle varianten) zo krachtig is, de eigenschap is dat elke lokale permutatie van de bogen rond een top van een reguliere boom kan worden uitgebreid naar een automorfisme van de volledige boom. Rechthoekige gebouwen zijn een brede veralgemening van meetkundige structuren die eenzelfde flexibiliteit kennen, zodat universele groepen ook over deze gebouwen zinvol zijn. Naast de permutationele eigenschappen van de lokale groepen spelen nu ook de combinatorische eigenschappen van (het diagram van) het gebouw een grote rol.

Wij zetten in deze thesis de studie van deze groepen verder. De focus van Silva lag op lokaal eindige rechthoekige gebouwen, waarbij de lokale permutatiegroepen ook transitief werken. Wij werken zonder dergelijke aannames en veralgemenen de gekende resultaten verder. Daarnaast passen ook enkele resultaten van Smith en Le Boudec probleemloos in deze context.

Overzicht van de resultaten

We gaan van start met een inleidend hoofdstuk. Hierin bouwen we de nodige achtergrondkennis op in abstracte en topologische groepentheorie, geven we een snelle opfrissing van wat grafentheorie, vermelden we enkele algemene eigenschappen van automorfismen van bomen, geven we een korte schets van de resultaten van Burger–Mozes, Smith, Le Boudec en Lederle, en wijden we ten slotte nog een sectie aan algemene gebouwentheorie.

Het tweede hoofdstuk richten we onze pijlen specifiek op rechthoekige gebouwen. We vermelden een aantal algemene hulpresultaten over automorfismen, gallerijen, projecties, evenwijdigheid, … Ook definiëren we kleuringen van gebouwen, essentieel om de lokale acties van een automorfisme te kunnen volgen. We voeren ook enkele nieuwe concepten in. Zo voorzien *implosies* een manier om rechthoekige gebouwen samen te trekken op een fundamenteel andere manier dan de vertrouwde projectie- en retractieafbeeldingen dat doen, en definiëren we *stadsproducten*, die ons toelaten om nieuwe gebouwen te construeren door gebouwen aaneen te plakken langs een ander gebouw.

In Hoofdstuk 3 voeren we uiteindelijk de universele groepen over rechthoekige gebouwen in volle algemeenheid in. We hebben een technisch lemma nodig dat toelaat om partiële automorfismen uit te breiden naar automorfismen “zo goed lijkend op elementen uit de universele groep als mogelijk” (Stelling 3.1.9). We berekenen de banen van de universele groepen op het gebouw en karakteriseren wanneer de actie transitief is. Daarna veralgemenen we een resultaat van Smith en karakteriseren we wanneer de actie op de residuen van het gebouw primitief is (Stelling 3.2.15). We bekijken ook de deelgroep van de universele groep voortgebracht door alle kamerstabilisatoren. Deze deelgroep is het grootste obstakel voor enkelvoudigheid; we karakteriseren dan ook wanneer deze triviaal is en bewijzen verderop, aan de hand van een algemeen enkelvoudigheidscriterium, dat deze in ieder geval zelf enkelvoudig is.

Op zowel de lokale groepen als de universele groepen zelf leggen we de permutatietopologie. Deze topologie maakt de universele groep in ieder geval volledig onsamenhangend. Daar we er niet langer van uitgaan dat de lokale groepen eindig zijn, wordt deze topologie een pak subtieler. We bekijken hoe de lokale en globale topologie elkaar beïnvloeden en karakteriseren bijvoorbeeld wanneer de universele groep lokaal compact is. We stellen voldoende voorwaarden op onder dewelke de groep compact voortgebracht is. We vermoeden dat deze voorwaarden ook nodig zijn, en motiveren dit aan de hand van enkele bijzondere gevallen en partiële resultaten.

Tot slot beschrijven we universele groepen over stadsproducten als universele groepen van lagere rang, waar de lokale data opnieuw bestaat uit universele groepen maar over de factorgebouwen.

We geven een overzicht van Hoofdstuk 3 in meer detail. We veronderstellen voor de eenvoud dat het gebouw irreducibel is, en verwijzen naar de afzonderlijke stellingen voor de precieze resultaten.

Stelling. *Zij M een irreducibel diagram over indexverzameling I. Zij $\boldsymbol{F}$ een collectie permutatiegroepen $F_i \leq \mathrm{Sym}(\Omega_i)$ met $3 \leq |\Omega_i|$ en geïndexeerd door $i \in I$. Zij Δ een semiregulier rechthoekig gebouw van type M met parameters gelijk aan $|\Omega_i|$ voor elke $i \in I$. Voorzie Δ van een legale kleuring λ die kleuren aanneemt in de verzamelingen Ω_i.*

Zij $\mathcal{U}_\Delta^\lambda(\boldsymbol{F}) \leq \mathrm{Aut}(\Delta)$ de bijhorende universele groep uitgerust met de permutatietopologie. Dan geldt het volgende.

- *$\mathcal{U}(\boldsymbol{F})$ is transitief op de residuen van type J als en slechts als F_i transitief is voor elke $i \in I \setminus J$.* Dit is Propositie 3.2.2.

- *$\mathcal{U}(\boldsymbol{F})$ is primitief op de residuen van type J als en slechts als $I = J \sqcup \{k\}$ voor zekere $k \in I$, terwijl F_k primitief en niet-regulier is en F_i transitief voor elke $i \in I \setminus k^{\perp}$.*
  Dit is Stelling 3.2.15.
- *$\mathcal{U}(\boldsymbol{F})$ is voortgebracht door kamerstabilisatoren als en slechts als F_i voortgebracht is door puntstabilisatoren voor elke $i \in I$ en transitief is voor elke i in een toppenbedekking van M.*
  Dit is Stelling 3.2.20.
- *$\mathcal{U}(\boldsymbol{F})$ is volledig onsamenhangend.*
  Dit is Propositie 1.2.27.
- *$\mathcal{U}(\boldsymbol{F})$ is discreet als en slechts als F_i vrij werkt op Ω_i voor elke $i \in I$.*
  Dit is Propositie 3.3.1.
- *$\mathcal{U}(\boldsymbol{F})$ is gesloten in $\mathrm{Aut}(\Delta)$ als en slechts als F_i gesloten is in $\mathrm{Sym}(\Omega_i)$ voor elke $i \in I$.*
  Dit is Propositie 3.3.3.
- *Als $\mathcal{U}(\boldsymbol{F})$ gesloten is, dan is $\mathcal{U}(\boldsymbol{F})$ lokaal compact als en slechts als F_i uitsluitend eindige subbanen heeft, voor elke $i \in I$.*
  Dit is Propositie 3.3.4.
- *Onder voorbehoud van Vermoeden 3.3.9 en als $\mathcal{U}(\boldsymbol{F})$ gesloten en lokaal compact is, dan is $\mathcal{U}(\boldsymbol{F})$ compact voortgebracht als en slechts als F_i compact voortgebracht is en eindig veel banen heeft, voor elke $i \in I$.*
  Dit is een herformulering van Stellingen 3.3.6 en 3.3.7 en Vermoeden 3.3.9.
- *Als $\mathcal{U}(\boldsymbol{F})$ gesloten en niet-discreet is, dan is $\mathcal{U}(\boldsymbol{F})^+$ enkelvoudig.*
  Dit is Gevolg 3.4.10.

In het bijzonder, kiezen we lokale groepen F_i waarvoor elke puntstabilisator slechts eindige banen heeft, dan is $\mathcal{U}(\boldsymbol{F})$ een volledig onsamenhangende lokaal compacte groep. Is bovendien minstens één puntstabilisator niet-triviaal, dan is $\mathcal{U}(\boldsymbol{F})$ niet-discreet en is de deelgroep voortgebracht door alle kamerstabilisatoren een enkelvoudige groep.

In Hoofdstuk 4 definiëren we een analogon voor de groepen van Le Boudec in de gebouwensetting, waar eindig veel singulariteiten worden toegelaten. De combinatoriek van het gebouw leidt snel tot extra restricties (in de vorm van *ladders*). We motiveren waarom we ons hier opnieuw beperken tot lokaal eindige gebouwen en bestuderen de eigenschappen van de verkregen groepen. Opnieuw voorzien we deze van een topologie: dit keer een op maat gemaakte topologie, die toelaat om heel wat eigenschappen van de universele groepen rechtstreeks over te zetten. Ten slotte veralgemenen we een resultaat uit [CRW19] en karakteriseren we wanneer de groepen virtueel enkelvoudig zijn.

We geven een overzicht van Hoofdstuk 4 in meer detail. We veronderstellen voor de eenvoud dat het gebouw irreducibel is, en verwijzen naar de afzonderlijke stellingen voor de precieze resultaten.

Stelling. *Zij M een irreducibel diagram over indexverzameling I. Zij $\boldsymbol{F}, \acute{\boldsymbol{F}}$ twee collecties permutatiegroepen $F_i \leq \acute{F}_i \leq \widehat{F}_i \leq \mathrm{Sym}(\Omega_i)$ met $3 \leq |\Omega_i| < \infty$ en geïndexeerd door $i \in I$. Zij Δ een semiregulier rechthoekig gebouw van type M met parameters gelijk aan $|\Omega_i|$ voor elke $i \in I$. Voorzie Δ van een legale kleuring λ die kleuren aanneemt in de verzamelingen Ω_i. We benadrukken dat Δ lokaal eindig is.*

Zij $\mathcal{G}^{\lambda}_{\Delta}(\boldsymbol{F}, \acute{\boldsymbol{F}}) \leq \mathrm{Aut}(\Delta)$ de bijhorende gerestringeerde universele groep, uitgerust met de topologie van Definitie 4.3.1. Dan geldt het volgende.

- *$\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ is transitief op de residuen van type J als en slechts als F_i transitief is voor elke $i \in I \setminus J$.*
  Dit is Propositie 4.2.2.

- $\mathcal{G}(\boldsymbol{F}, \boldsymbol{\acute{F}})$ *is volledig onsamenhangend en lokaal compact.*
  Dit is Propositie 4.3.2.
- $\mathcal{G}(\boldsymbol{F}, \boldsymbol{\acute{F}})$ *is discreet als en slechts als* F_i *vrij werkt op* Ω_i *voor elke* $i \in I$.
  Dit is Propositie 4.3.3.
- $\mathcal{G}(\boldsymbol{F}, \boldsymbol{\acute{F}})$ *is compact voortgebracht.*
  Dit is Gevolg 4.3.7.
- *Als* $\mathcal{G}(\boldsymbol{F}, \boldsymbol{\acute{F}})$ *niet-discreet is en als* $F_i = \acute{F}_i$ *voor elke sport* $i \in I$ *van een ladder in het diagram, dan is* $\mathcal{G}(\boldsymbol{F}, \boldsymbol{\acute{F}})$ *virtueel enkelvoudig als en slechts als* $\mathcal{U}(\boldsymbol{\acute{F}})$ *enkelvoudig is — of expliciet,* $\acute{F}_i$ *voortgebracht is door puntstabilisatoren voor elke* $i \in I$ *en transitief voor elke* i *in een toppenbedekking van* M.
  Dit is Stelling 4.4.4.

In zowel Hoofdstuk 3 als 4 valt regelmatig op dat de combinatoriek van het gebouw tot boeiende bijkomende condities en eigenschappen leidt, zodat de globale structuur van de verkregen groepen in sterke mate afhangt van zowel de lokale permutatiegroepen als de Coxeterdiagrammen.

We sluiten deze thesis af met enkele open vragen, die doorheen het doctoraatsonderzoek opdoken en waar we niet de nodige tijd of inzichten voor bleken te hebben.

List of symbols

> *I strive for nothing if not consistency.*
>
> — Brandon Sanderson, *The Final Empire*

Abstract groups and group actions

g, h	typical group elements
id	the identity element
$g \cdot h$	group multiplication
$\mathrm{Sym}(\Omega)$	the symmetric group on Ω
$\mathrm{Sym}(n)$	shorthand for $\mathrm{Sym}(\{1, \ldots, n\})$
$H \leq G$	a subgroup of G
$N \trianglelefteq G$	a normal subgroup of G
${}^h g$	the conjugate $h \cdot g \cdot h^{-1}$
$[g, h]$	the commutator $g \cdot h \cdot g^{-1} \cdot h^{-1}$
$g \,.\, x$	the group action of g on x
G_x	the stabiliser of an element x
$G_{\{Y\}}$	the setwise stabiliser of a set Y
$G_{(Y)}$	the pointwise stabiliser of a set Y
$G \,.\, x$	the orbit of an element x
X/G	the orbit space
G^+	the subgroup $\langle G_x \mid x \in X \rangle$ generated by point stabilisers
$\widehat{G}$	the Young overgroup of G

Actions on trees

T	a typical tree
∂T	the boundary of T
γ	a typical path in T (finite or infinite)
$\ell(g)$	the displacement of g
$\mathsf{A}(g)$	the axis of g

Chamber systems

I	an index set
I^*	the free monoid on I
i, j, k	typical elements of I
J	a typical subset of I
Δ	a chamber system or building
c, d	typical chambers of Δ
$\sim_i$	an i-adjacency relation on Δ
γ	a gallery in Δ
$\mathrm{Aut}(\Delta)$	the group of (type-preserving) automorphisms of Δ
$\mathcal{P}$	a panel in Δ
$\mathcal{P}_i(c)$	the panel of type i containing c
$\mathcal{R}$	a residue in Δ
$\mathcal{R}_J(c)$	the residue of type J containing c
$\mathrm{Res}_J(\Delta)$	the set of residues of Δ of type J
$\mathrm{Res}_j(\Delta)$	the set of panels of Δ of type j
dist	the distance function in Δ
$\mathsf{B}_n(c)$	the ball of radius n and centre c
$\mathsf{S}_n(c)$	the sphere of radius n and centre c

Coxeter systems

M	a Coxeter matrix over I
m_{ij}	a Coxeter matrix's entries
(W, S)	a Coxeter system
ς	the evaluation map $I^* \to W$

$p(i,j)$	the word of length m_{ij} of alternating letters i and j
$\simeq$	the homotopy relation on I^*
M_J	the Coxeter matrix induced by J
(W_J, S_J)	the Coxeter system induced by J
$\mathbb{L}(W,S)$	the nerve of the Coxeter system
$\mathbb{K}(W,S)$	the cone over the barycentric subdivision of the nerve $\mathbb{L}(W,S)$

General buildings

δ	a building's W-distance function
$\mathcal{A}$	an apartment in Δ
q_i	a semiregular building's parameter
$\mathrm{proj}_{\mathcal{R}}$	the projection map onto $\mathcal{R}$
$\rho_{c,\mathcal{A}}$	the retraction map onto $\mathcal{A}$ with centre c
$\mathbb{K}(\Delta)$	the Davis realisation of Δ

Right-angled buildings

$J^{\perp}$	the set of $i \in I$ with $m_{ij} = 2$ for all $j \in J$
$j^{\perp}$	an abbreviation for $\{j\}^{\perp}$
$X_J(c)$	the J-wing of c
$X_j(c)$	the j-wing of c
$\mathcal{T}$	a tree-wall
Γ_i	the i-tree-wall tree
Ω_i	a set of i-colours of cardinality q_i
λ	a (legal) colouring $(\lambda_i)_{i\in I}$
λ_i	an i-colouring, component of λ
(Δ', τ)	an implosion of Δ
✠$_M$	a city product of diagrams or buildings, over diagram M
Φ	the skeletal building of a city product

Universal groups

$\sigma_\lambda(g,\mathcal{P})$	the local action of g at $\mathcal{P}$
$\boldsymbol{F}$	a family of local groups $F_i \leq \mathrm{Sym}(\Omega_i)$
$\mathcal{U}_\Delta^\lambda(\boldsymbol{F})$	the universal group of $\boldsymbol{F}$ over Δ, usually abbreviated as $\mathcal{U}(\boldsymbol{F})$ or $\mathcal{U}$
$\mathcal{U}\vert_{\mathcal{P}}$	the panel group of $\mathcal{U}$ w.r.t. $\mathcal{P}$
$\mathcal{U}^+$	the subgroup $\langle \mathcal{U}_c \mid c \in \Delta\rangle$ generated by chamber stabilisers
$V_i(c)$	the subgroup with support in $X_i(c)$
$W_i(c)$	the pointwise stabiliser of $X_i(c)$

Restricted universal groups

$\acute{\boldsymbol{F}}$	a second family of local groups satisfying $F_i \leq \acute{F}_i \leq \widehat{F} \leq \mathrm{Sym}(\Omega_i)$
$\mathcal{G}_\Delta^\lambda(\boldsymbol{F})$	a relaxation of $\mathcal{U}_\Delta^\lambda(\boldsymbol{F})$ allowing for finitely many singularities
$\mathcal{G}_\Delta^\lambda(\boldsymbol{F}, \acute{\boldsymbol{F}})$	the group $\mathcal{G}(\boldsymbol{F}) \cap \mathcal{U}(\acute{\boldsymbol{F}})$, usually abbreviated as $\mathcal{G}(\boldsymbol{F}, \acute{\boldsymbol{F}})$ or $\mathcal{G}$
$S(g)$	the set of singularities of g
$K_{\mathcal{P}}$	a subgroup of $\mathcal{G}_{\{\mathcal{P}\}}$ with only singularities parallel to $\mathcal{P}$
$K_{n,\mathcal{P}}$	a subgroup of $\mathcal{G}_{\{\mathcal{P}\}}$ with only singularities at bounded distance

List of figures

> *Almost it seemed that the words took shape, and visions of far lands and bright things that he had never yet imagined opened out before him.*
>
> — J.R.R. Tolkien, *The Fellowship of the Ring*

Bibliography

> *A mind needs books as a sword needs a whetstone, if it is to keep its edge.*
>
> — George R.R. Martin, *A Game of Thrones*

[AB08] P. Abramenko, K. S. Brown, *Buildings: theory and applications.* Graduate Texts in Mathematics, vol. 248, Springer New York, (2008).

[Ama03] O. É. Amann, *Groups of tree-automorphisms and their unitary representations.* Ph.D. thesis, ETH Zürich, (2003).

[BCGM12] U. Bader, P.-E. Caprace, T. Gelander, S. Mozes, *Simple groups without lattices.* Bulletin of the London Mathematical Society, vol. 44, no. 1, (2012), pp. 55–67.

[BDM21] J. Bossaert, T. De Medts, *Topological and algebraic properties of universal groups for right-angled buildings.* Forum Mathematicum, vol. 33, no. 4, (2021), pp. 867–888.

[BEW15] C. Banks, M. Elder, G. A. Willis, *Simple groups of automorphisms of trees determined by their actions on finite subtrees.* Journal of Group Theory, vol. 18, no. 2, (2015), pp. 235–261.

[BGS85] W. Ballmann, M. Gromov, V. Schroeder, *Manifolds of nonpositive curvature.* Progress in Mathematics, vol. 61, Birkhäuser Basel, (1985).

[BH99] M. Bridson, A. Häfliger, *Metric spaces of non-positive curvature.* Grundlehren der Mathematischen Wissenschaften, vol. 319, Springer Berlin Heidelberg, (1999).

[BM00a] M. Burger, S. Mozes, *Groups acting on trees: from local to global structure.* Publications Mathématiques de l'IHÉS, vol. 92, (2000), pp. 113–150.

[BM00b] ———, *Lattices in product of trees.* Publications Mathématiques de l'IHÉS, vol. 92, (2000), pp. 151–194.

[BMPZ17] A. Baudisch, A. Martin-Pizarro, M. Ziegler, *A model-theoretic study of right-angled buildings.* Journal of the European Mathematical Society, vol. 19, no. 10, (2017), pp. 3091–3141.

[Bor01] A. Borel, *Essays in the history of Lie groups and algebraic groups.* History of Mathematics, American Mathematical Society, (2001).

[Bou97] M. Bourdon, *Immeubles hyperboliques, dimension conforme et rigidité de Mostow.* Geometric & Functional Analysis GAFA, vol. 7, no. 2, (1997), pp. 245–268.

[Bou07] N. Bourbaki, *Topologie générale: chapitres 1 à 4.* Springer Berlin Heidelberg, (2007).

[Bra45] J. Braconnier, *Sur les groupes topologiques localement compacts.* Thèses de l'entre-deux-guerres, vol. 280, (1945).

[BRW13] U. Baumgartner, J. Ramagge, G. A. Willis, *Scale-multiplicative semigroups and geometry: automorphism groups of trees.* Groups, Geometry, and Dynamics, vol. 10, no. 3, (2013), pp. 1051–1075.

[Byw19] T. P. Bywaters, *Directions and scale-multiplicative semigroups in restricted Burger–Mozes groups.* arXiv:1911.08032v2 [math.GR], `https://arxiv.org/abs/1911.08032v2`, (2019).

[Cam53] R. Camm, *Simple free products.* Journal of the London Mathematical Society, vol. S1-28, no. 1, (1953), pp. 66–76.

[Cap14a] P.-E. Caprace, *Automorphism groups of right-angled buildings: simplicity and local splittings.* Fundamenta Mathematicae, vol. 224, no. 1, (2014), pp. 17–51.

[Cap14b] ———, *Lectures on* $\mathrm{CAT}(0)$ *spaces and their isometry groups.* In *Geometric group theory,* pp. 91–125, IAS/Park City Mathematics Series, vol. 21, American Mathematical Society, (2014), based on lectures given at PCMI 2012.

[Car02] L. Carbone, *The tree lattice existence theorems.* Comptes Rendus Mathématique, vol. 335, no. 3, (2002), pp. 223–228.

[Cas02] B. Casselman, *Computation in Coxeter groups (I: Multiplication).* The Electronic Journal of Combinatorics, vol. 9, no. 1, (2002), p. R25, electronic only.

[Cas20] I. Castellano, *Rational discrete first degree cohomology for totally disconnected locally compact groups.* Mathematical Proceedings of the Cambridge Philosophical Society, vol. 168, no. 2, (2020), pp. 361–377.

[CRW19] P.-E. Caprace, C. Reid, P. Wesolek, *Approximating simple locally compact groups by their dense locally compact subgroups.* International Mathematics Research Notices, vol. 2021, no. 7, (2019), pp. 5037–5110.

[Dav98] M. W. Davis, *Buildings are* $\mathrm{CAT}(0)$. In *Geometry and Cohomology in Group Theory,* pp. 108–123, London Mathematical Society Lecture Note Series, Cambridge University Press, (1998).

[Dav12] M. Davis, *The geometry and topology of Coxeter groups.* London Mathematical Society Monographs, vol. 32, Princeton University Press, (2012).

[dCdLH16] Y. de Cornulier, P. de La Harpe, *Metric geometry of locally compact groups.* EMS Tracts in Mathematics, vol. 25, European Mathematical Society, (2016).

[Dik18] D. Dikranjan, *Introduction to topological groups.* Downloaded from `https://users.dimi.uniud.it/~dikran.dikranjan/`, (2018), version 26.02.2018.

[DMC11] T. De Medts, P.-E. Caprace, *Simple locally compact groups acting on trees and their germs of automorphisms.* Transformation Groups, vol. 16, no. 2, (2011), pp. 375–411.

[DMdS19] T. De Medts, A. F. C. da Silva, *Open subgroups of the automorphism group of a right-angled building.* Geometriae Dedicata, vol. 203, (2019), pp. 1–23.

[DMdSS18] T. De Medts, A. F. C. da Silva, K. Struyve, *Universal groups for right-angled buildings.* Groups, Geometry and Dynamics, vol. 12, no. 1, (2018), pp. 231–287.

[Dra01] A. N. Dranishnikov, *On boundaries of hyperbolic Coxeter groups.* Topology and its Applications, vol. 110, no. 1, (2001), pp. 29–38.

[Gal67] T. Gallai, *Transitiv orientierbare Graphen.* Acta Mathematica Academiae Scientiarum Hungarica, vol. 18, no. 1, (1967), pp. 25–66.

[GGT18] A. Garrido, Y. Glasner, S. Tornier, *Automorphism groups of trees: generalities and prescribed local actions.* In *New Directions in Locally Compact Groups*, pp. 92–116, London Mathematical Society Lecture Note Series, Cambridge University Press, (2018).

[GL18] L. Garncarek, N. Lazarovich, *The Neretin groups.* In P.-E. Caprace, N. Monod (eds.), *New Directions in Locally Compact Groups*, pp. 131–144, London Mathematical Society Lecture Note Series, Cambridge University Press, (2018).

[HHR18] W. Herfort, K. H. Hofmann, F. G. Russo, *Periodic locally compact groups: a study of a class of totally disconnected topological groups.* De Gruyter Studies in Mathematics, vol. 71, De Gruyter, (2018).

[Hig67] D. G. Higman, *Intersection matrices for finite permutation groups.* Journal of Algebra, vol. 6, (1967), pp. 22–42.

[Hig76] P. J. Higgins, *The fundamental groupoid of a graph of groups.* Journal of the London Mathematical Society, vol. S2-13, no. 1, (1976), pp. 145–149.

[Hol19] D. Holt, *Permutation groups generated by finitely many point stabilisers.* MathOverflow, (2019), `https://mathoverflow.net/q/347667`, version 05/12/2019.

[Hos03] T. Hosaka, *Determination up to isomorphism of right-angled Coxeter systems.* Proceedings of the Japan Academy, Series A, Mathematical Sciences, vol. 79, no. 2, (2003), pp. 33–35.

[HP03] F. Haglund, F. Paulin, *Constructions arborescentes d'immeubles.* Mathematische Annalen, vol. 325, no. 1, (2003), pp. 137–164.

[HR13] E. Hewitt, K. A. Ross, *Abstract harmonic analysis: volume I (structure of topological groups, integration theory, group representations).* Die Grundlehren der mathematischen Wissenschaften, Springer Berlin Heidelberg, (2013).

[Ji06] L. Ji, *Buildings and their applications in geometry and topology.* Asian Journal of Mathematics, vol. 10, no. 1, (2006), pp. 11–80.

[KM07] B. Krön, R. G. Möller, *Analogues of Cayley graphs for topological groups.* Mathematische Zeitschrift, vol. 258, no. 3, (2007), pp. 637–675.

[Kra20] L. Kramer, *Locally compact groups.* Downloaded from `https://www.wwu.de/AGKramer/index.php?name=LocallyCompactGroups`, (2020), version 19/05/2020.

[KT12] A. Kubena, A. Thomas, *Density of commensurators for uniform lattices of right-angled buildings.* Journal of Group Theory, vol. 15, no. 5, (2012), pp. 565–611.

[LB16] A. Le Boudec, *Groups acting on trees with almost prescribed local action.* Commentarii Mathematici Helvetici, vol. 91, no. 2, (2016), pp. 253–293.

[Led17] W. Lederle, *Coloured Neretin groups.* arXiv:1701.03027v3 [math.GR], `https://arxiv.org/abs/1701.03027v3`, (2017).

[LMZ94] A. Lubotzky, S. Mozes, R. J. Zimmer, *Superrigidity for the commensurability group of tree lattices.* Commentarii mathematici Helvetici, vol. 69, no. 4, (1994), pp. 523–548.

[Mar19] T. Marquis, *On geodesic ray bundles in buildings.* Geometriae Dedicata, vol. 202, no. 1, (2019), pp. 27–43.

[Mat15] H. Matui, *Topological full groups of one-sided shifts of finite type.* Journal für die reine und angewandte Mathematik (Crelles Journal), vol. 2015, no. 705, (2015), pp. 35–84.

[Möl03] R. G. Möller, *FC^{-}-elements in totally disconnected groups and automorphisms of infinite graphs.* Mathematica Scandinavica, vol. 92, no. 2, (2003), pp. 261–268.

[Möl10] ———, *Graphs, permutations and topological groups.* arXiv:1008.3062v2 [math.GR], `https://arxiv.org/abs/1008.3062v2`, (2010).

[Mou88] G. Moussong, *Hyperbolic coxeter groups.* Ph.D. thesis, Ohio State University, (1988).

[MS99] R. M. McConnell, J. P. Spinrad, *Modular decomposition and transitive orientation.* Discrete Mathematics, vol. 201, no. 1, (1999), pp. 189–241.

[Neb00] C. Nebbia, *Minimally almost periodic totally disconnected groups.* Proceedings of the American Mathematical Society, vol. 128, no. 2, (2000), pp. 347–351.

[OEISa] The On-Line Encyclopedia of Integer Sequences, `http://oeis.org/A000088`.

[OEISb] The On-Line Encyclopedia of Integer Sequences, `http://oeis.org/A001349`.

[OEISc] The On-Line Encyclopedia of Integer Sequences, `http://oeis.org/A079473`.

[OEISd] The On-Line Encyclopedia of Integer Sequences, `http://oeis.org/A000041`.

[Rad02] D. G. Radcliffe, *Rigidity of right-angled Coxeter groups.* arXiv:9901049v3 [math.GR], `https://arxiv.org/abs/math/9901049v3`, (2002).

[Rad20] N. Radu, *New simple lattices in products of trees and their projections.* Canadian Journal of Mathematics, vol. 72, no. 6, (2020), pp. 1624–1690.

[Rei19] C. Reid, *Compactly generated vertex stabilisers in compactly generated t.d.l.c. groups acting on trees.* MathOverflow, (2019), `https://mathoverflow.net/q/341874`, version 18/09/2019.

[Ron09] M. Ronan, *Lectures on buildings.* University of Chicago Press, (2009).

[Smi17] S. M. Smith, *A product for permutation groups and topological groups.* Duke Mathematical Journal, vol. 166, no. 15, (2017), pp. 2965–2999.

[Tao11] T. Tao, *Hilbert's fifth problem and Gleason metrics.* Blog, (2011), `https://terrytao.wordpress.com/2011/06/17/hilberts-fifth-problem-and-gleason-metrics/`, version 17/08/2021.

[Tho06] A. Thomas, *Lattices acting on right-angled buildings.* Algebraic & Geometric Topology, vol. 6, no. 3, (2006), pp. 1215–1238.

[Tit65] J. Tits, *Structures et groupes de Weyl.* Séminaire Bourbaki, vol. 9, (1965), pp. 169–183.

[Tit69] ———, *Le problème des mots dans les groupes de Coxeter.* In *Symposia Mathematica*, pp. 175–185, Academic Press, (1969).

[Tit70] ———, *Sur le groupe des automorphismes d'un arbre.* In *Essays on Topology and Related Topics: Memoires dédiés à Georges de Rham*, pp. 188–211, Springer Berlin Heidelberg, (1970).

[Tit74] ———, *Buildings of spherical type and finite BN-pairs.* Springer Berlin Heidelberg, (1974).

[Tit81] ———, *A local approach to buildings.* In *The Geometric Vein*, pp. 519–547, Springer New York, (1981).

[Tit92] ———, *Twin buildings and groups of Kac–Moody type.* In *Groups, Combinatorics and Geometry*, pp. 249–286, London Mathematical Society Lecture Note Series, Cambridge University Press, (1992).

[vD36] D. van Dantzig, *Zur topologischen Algebra. III. Brouwersche und Cantorsche Gruppen.* Compositio Mathematica, vol. 3, (1936), pp. 408–426.

[VM98] H. Van Maldeghem, *Generalized polygons.* Monographs in Mathematics, vol. 93, Birkhäuser, (1998).

[Wei03] R. M. Weiss, *The structure of spherical buildings.* Princeton University Press, (2003).

[Wei09] ———, *The structure of affine buildings.* Princeton University Press, (2009).

[Wes18] P. R. Wesolek, *An introduction to totally disconnected locally compact groups.* Downloaded from `http://people.math.binghamton.edu/wesolek/`, (2018), version 11.12.2018.

[Wil94] G. A. Willis, *The structure of totally disconnected, locally compact groups.* Mathematische Annalen, vol. 300, no. 2, (1994), pp. 341–364.

[Wil01] ———, *Further properties of the scale function on a totally disconnected group.* Journal of Algebra, vol. 237, no. 1, (2001), pp. 142–164.

[Wil07] ———, *Compact open subgroups in simple totally disconnected groups.* Journal of Algebra, vol. 312, no. 1, (2007), pp. 405–417.

[Woe91] W. Woess, *Topological groups and infinite graphs.* Discrete Mathematics, vol. 95, no. 1, (1991), pp. 373–384.

Index

> *Wicked people never have time for reading. It's one of the reasons for their wickedness.*
>
> — Lemony Snicket, *The Penultimate Peril*