J. Eur. Math. Soc. 19, 3091–3141

DOI 10.4171/JEMS/736

Andreas Baudisch · Amador Martin-Pizarro · Martin Ziegler

A model-theoretic study of right-angled buildings

Received November 10, 2014

Abstract. We study the model theory of right-angled buildings with infinite residues. For every Coxeter graph we obtain a complete theory with a natural axiomatisation, which is ω-stable and equational. Furthermore, we provide sharp lower and upper bounds for its degree of ampleness, computed exclusively in terms of the associated Coxeter graph. This generalises and provides an alternative treatment of the free pseudospace.

Keywords. Model theory, ampleness, Coxeter, buildings

Contents

1. Introduction

A *Coxeter group* (W, Γ) consists of a group W with a fixed set Γ of generators and defining relations $(\gamma \cdot \delta)^{m_{\gamma,\delta}} = 1$, where $m_{\gamma,\gamma} = 1$ and $m_{\gamma,\delta} = m_{\delta,\gamma}$, for $\gamma \neq \delta$, is either ∞ or an integer larger than 1. We will exclusively consider finitely generated Coxeter groups, with Γ finite. A *word* w is a finite sequence on the generators from Γ,

A. Baudisch: Institut für Mathematik, Humboldt-Universität zu Berlin,
D-10099 Berlin, Germany; e-mail: baudisch@mathematik.hu-berlin.de

A. Martin-Pizarro: Université de Lyon, CNRS UMR 5208, Université Lyon 1,
Institut Camille Jordan, 43 boulevard du 11 novembre 1918, F-69622 Villeurbanne Cedex, France;
e-mail: pizarro@math.univ-lyon1.fr

M. Ziegler: Mathematisches Institut, Albert-Ludwigs-Universität Freiburg,
D-79104 Freiburg, Germany; e-mail: ziegler@uni-freiburg.de

Mathematics Subject Classification (2010): Primary 03C45; Secondary 51E24

and w is *reduced* if its length is minimal with respect to all words representing the same element of W. A *chamber system* (X, W, Γ) for the Coxeter group (W, Γ) is a set X, equipped with a family of equivalence relations $(\sim_\gamma, \gamma \in \Gamma)$. If $w = \gamma_1 \cdots \gamma_n$ is a reduced word, a *reduced path of type* w from x to y in X is a sequence $x = x_0, \ldots, x_n = y$ such that x_{i-1} and x_i are $\sim_{\gamma_i}$-related and different for every $1 \leq i \leq n$. A chamber system (X, W, Γ) is a *building* if each $\sim_\gamma$-class contains at least two elements and, for every pair x and y in X, there exists an element $g \in W$ such that there is a reduced path of type w from x to y if and only if the word w represents g. It follows that g is uniquely determined by x and y, and that the reduced path connecting x and y is uniquely determined by its type w. We refer the reader to [6] for a pleasant introduction to buildings.

The Coxeter group (W, Γ) is *right-angled* if for every $\gamma \neq \delta$, the value $m_{\gamma,\delta}$ is either 2 or ∞. So W is determined by its *Coxeter diagram*: a graph with vertex set Γ such that γ and δ have an edge connecting them, which we denote by $R(\gamma, \delta)$, if $m_{\gamma,\delta} = \infty$. In an abuse of notation, we will denote this graph by Γ as well. The elements of Γ will be referred to as *colours* or *levels*. Note that, for involutions γ and δ, the relation $(\gamma \cdot \delta)^2 = 1$ means that γ and δ commute. We call a word w' a *permutation* of w if it can be obtained from w by a sequence of swaps of commuting consecutive generators. In right-angled Coxeter groups, a word w is reduced if and only if no permutation of w has the form $w_1 \cdot \gamma \cdot \gamma \cdot w_2$ for some generator γ. Every element of W is represented by a unique reduced word, up to permutation. In a building (X, W, Γ), the the class $x/\sim_\gamma$ is the γ*-residue* of x. A right-angled Coxeter group admits a unique (up to isomorphism) countable building $\mathrm{B}_0(\Gamma)$ with infinite residues [7, Proposition 5.1], which we call *rich*.

A right-angled building can also be described in terms of an incidence geometry or, as we will say, a *coloured graph*. The vertices of colour γ are equivalence classes of $\sim^\gamma$, the transitive closure of all $\sim_{\gamma'}$, for $\gamma' \neq \gamma$. Two vertices are connected by an edge if the corresponding classes intersect. The coloured graph associated to the rich building given by the diagram

$[0, n]$

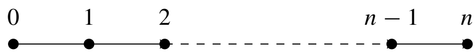

is, as noticed by Tent [15], the prime model of the theory of the *free n-dimensional pseudospace*. A model of a theory is prime if it elementarily embeds into every model. The n-dimensional pseudospace, also considered in [2], witnessed the strictness of the ample hierarchy for ω-stable theories. We are indebted to Tent for pointing out the connection between the free pseudospace and Tits buildings, which was the starting point of the present work.

Recall that a countable complete theory is ω-stable if there is a *rank function* R defined on the collection of definable sets of a (sufficiently saturated) model M with the following principle: If $X \subset M^n$ is definable, then $R(X) > \alpha$ if and only if X contains an infinite family of pairwise disjoint definable sets Y_i with $R(Y_i) \geq \alpha$. The smallest rank function is called *Morley rank*. The Morley rank of a type is the smallest Morley rank of its formulae. This notion agrees with the Cantor–Bendixson rank on the space of types over an ω-saturated model, equipped with the Stone topology. For an algebraically closed field with no additional structure, definable sets are exactly the Zariski constructible ones, and Morley rank coincides with Zariski dimension.

If M is a *group of finite Morley rank*, that is, it carries a definable group structure and the Morley rank is always finite, this notion of dimension is well-behaved. For example, given a definable fibration $S \subset X \times Y \to Y$, the subset consisting of those y in Y such that the fibre over y has dimension k is definable for every k in N. If all fibres have constant Morley rank k, then $\mathrm{RM}(X) = \mathrm{RM}(Y) + k$.

Motivated by a famous conjecture about the structure of *strongly minimal* sets (that is, irreducible definable sets of rank 1), the Algebraicity Conjecture states that every simple group of finite Morley rank can be seen as an algebraic group over an algebraically closed field, which is itself interpretable in the mere group structure. Though the general conjecture on strongly minimal sets was proven to be false [9], work on the Algebraicity Conjecture, which remains open, has become a fruitful research area, combining ideas from model theory with the classification of simple finite groups.

If an ω-stable theory does not interpret a certain incidence configuration present in euclidean space, then it interprets neither infinite fields nor specific possible counterexamples to the Algebraicity Conjecture, called *bad groups* [12]. The notion of n-ampleness [13, 4] for a theory generalises the incidence configuration given in euclidean $(n+1)$-space by flags of affine subspaces of increasing dimension, from a single point to a hyperplane. Ampleness introduces thus a geometrical hierarchy, according to which algebraically closed fields are n-ample for every n. The first two levels of this hierarchy suffice to describe the structure of definable groups [10, 12]: they are virtually abelian if the ω-stable group is not 1-ample, and virtually nilpotent if the group has finite Morley rank and is not 2-ample. However, little is known from 2 onwards. In particular, whether the ample hierarchy was strict remained long unknown. Evans conjectured that his example could be accordingly modified to illustrate the strictness. Extending the construction in [3], where a 2-ample theory was produced which interprets no infinite group (and thus no infinite field is interpretable), the aforementioned free n-dimensional pseudospace was constructed [2, 15] for every n, whose theory is ω-stable and n-ample yet not $(n+1)$-ample. In particular, the n-dimensional pseudospace is a graph with $n+1$ colours, labelled from 0 to n, such that the induced subgraph on consecutive colours is an infinite pseudoplane, whose theory was known to be 1-ample but not 2-ample.

In this article, we will provide an alternative approach to the above construction, which incorporates as well the rich buildings of every right-angled Coxeter group. As explained in Section 2, a right-angled building B can be recovered from its associated coloured graph M: The elements of B correspond to the *flags* of M, coloured subgraphs of M isomorphic to Γ. This article deals with the model theory of $\mathrm{M}_0(\Gamma)$, the coloured graph associated to the rich building $\mathrm{B}_0(\Gamma)$ given by Γ. Though models of the theory of $\mathrm{M}_0(\Gamma)$ need not arise from buildings, given flags F and G in a model, a notion of a *reduced path* between F and G can be defined, whose corresponding word consists of letters which are non-empty connected subsets of Γ. We show that the coloured graphs associated to rich buildings are *simply connected*: any two reduced paths between two given flags have the same word, up to permutation. A combinatorial study of the reduction of a path between two flags allows us to show the following crucial result (Theorem 3.26):

Theorem A. *Simple connectedness is an elementary property.*

Let PS_Γ denotes the collection of sentences stating the following two elementary properties: Simple connectedness and that, given any flag G and a colour γ, there are infinitely many flags which differ from G only at the vertex of colour γ. The following theorem (Theorem 4.12 and Corollaries 4.14 and 6.13) implies that PS_Γ axiomatises the complete theory of $\mathrm{M}_0(\Gamma)$:

Theorem B. *The theory* PS_Γ *is complete and* ω*-stable of Morley rank* $\omega^{(K-1)}$*, where* K *is the cardinality of a connected component of* Γ *of largest size. The coloured graph* $\mathrm{M}_0(\Gamma)$ *associated to the countable rich building* $\mathrm{B}_0(\Gamma)$ *is the unique prime model of* PS_Γ.

Morley rank defines a notion of independence, which agrees in the ω-stable case with non-forking, as introduced by Shelah. A remarkable feature of non-forking independence, which rules out the existence of infinite definable groups, is *total triviality*: whenever we consider a base set of parameters D, given tuples a, b and c such that a is independent both from b and from c over D, then it is independent from b, c over D. Recall that a canonical basis of a type p over a model is some set, fixed pointwise by exactly those automorphisms α of a sufficiently saturated model N fixing the global non-forking extension $\mathbf{p}$ of p over N. If $\mathrm{Cb}(p)$ exists, then it is unique, up to interdefinability, though generally canonical bases only exist as *imaginary* elements in the expansion T^{eq} of an ω-stable theory T. Within the wider class of stable theories, there is a distinguished subclass consisting of the *equational* ones, where each definable set in every cartesian product N^n of a model N is a boolean combination of instances of n-equations $\varphi(x, y)$, that is, the tuple x has length n and the family of finite intersections of instances $\varphi(x, a)$ has the descending chain condition. Whilst all known examples of stable theories arising naturally in nature are equational, the only stable non-equational theory constructed so far [8] is an expansion of the free 2-dimensional pseudospace. We obtain the following (Corollaries 7.25 and 7.28 and Proposition 7.26):

Theorem C. *The theory* PS_Γ *is equational, totally trivial with weak elimination of imaginaries: every type over a model has a canonical basis consisting of real elements.*

To conclude, we provide lower and upper bounds on the ample degree of the theory PS_Γ, which can be described in terms of the underlying Coxeter graph Γ. Set r to be the minimal valency of the non-isolated points of Γ and n the largest integer such that the graph $[0, n]$, as before, embeds as a full subgraph of Γ. We deduce (Theorem 8.6):

Theorem D. *The theory* PS_Γ *is* n*-ample but not* $(|\Gamma| - r + 1)$*-ample.*

These bounds are sharp and attained by $[0, n]$, by the circular graph on $n + 2$ points or by the extremal case of the complete graph K_N on $N \geq 2$ elements, whose theory is 1-ample but not 2-ample.

2. Buildings and geometries

Definition 2.1. A graph consists of set of vertices together with a symmetric and irreflexive binary relation. Two vertices a and b are *adjacent* if the pair (a, b) lies in the relation.

Given a finite graph Γ, its associated *right-angled Coxeter group* (W, Γ) consists of the group W generated by the elements of Γ with defining relations:

$$\begin{aligned} \gamma^2 &= 1 && \text{for all } \gamma \text{ in } \Gamma, \\ \gamma\delta &= \delta\gamma && \text{if } \gamma \text{ and } \delta \text{ are not adjacent.} \end{aligned}$$

By convention, no element γ commutes with itself.

From now on, all Coxeter groups are right-angled.

Fix a Coxeter group (W, Γ). A word $v = \gamma_1 \cdots \gamma_n$ in the generators is *reduced* if there is no pair $i \neq j$ such that γ_i equals γ_j and commutes (i.e. is not adjacent) with every letter occurring between γ_i and γ_j. Two words are *equivalent* if they represent the same element of W. Clearly, every word is equivalent to a reduced one. A reduced word w *commutes* with $\gamma \in \Gamma$ if every element of w does.

A word v' is a *permutation* of v if it can be obtained from v by a sequence of commutations on pairs of commuting generators. A permutation of a reduced word is again reduced.

The following is easy to see.

Lemma 2.2. *Two reduced words u and v are equivalent if and only if u is a permutation of v.*

For s a subset of Γ, let $\langle s\rangle$ denote the subgroup of W generated by s.

Corollary 2.3. *Given two subsets s and t of Γ,*

$$\langle s \cap t\rangle = \langle s\rangle \cap \langle t\rangle.$$

Definition 2.4. A *chamber system* (X, W, Γ) for the Coxeter group (W, Γ) consists of a set X equipped with a family of equivalence relations $\sim_\gamma$ for each $\gamma \in \Gamma$. Given a word $w = \gamma_1 \cdots \gamma_n$, a *path of type w* from x to y in X is a sequence $x = x_0, \ldots, x_n = y$ such that x_{i-1} and x_i are different and $\sim_{\gamma_i}$-related for every $1 \leq i \leq n$. A path of type w is *reduced* if w is.

A chamber system (X, W, Γ) is a *building* if each $\sim_\gamma$-class contains at least two elements, and for any x and y in X, there exists $g \in W$ with the property that there is a reduced path of type w from x to y if and only if the word w represents g.

We will refer to a chamber system (X, W, Γ) uniquely by the underlying set X if the corresponding Coxeter group (W, Γ) is clear. The following two lemmas can be easily shown.

Lemma 2.5. *In a building X, a reduced path of type w connecting x and y is uniquely determined by w, x and y.*

We will indicate the existence of a path of type w connecting x to y by $x \xrightarrow{w} y$. In particular, $x \xrightarrow{\gamma} y$ if and only if $x \neq y$ are $\sim_\gamma$-related.

Lemma 2.6. *A chamber system X is a building if and only if the following four conditions hold:*

(a) *Every $\sim_\gamma$-class has at least two elements.*
(b) *Any two elements of X are connected by a path.*
(c) *Given commuting generators γ and δ, if the elements x and y are connected by a path of type $\gamma\delta$, then they are also connected by a path of type $\delta\gamma$.*
(d) *There is no non-trivial closed reduced path.*

A chamber system satisfying conditions (b) and (c) is called *strongly connected*. A strongly connected chamber system is a *quasi-building* if it satisfies (d). Lemma 2.5 holds for quasi-buildings as well.

Remark 2.7. Given $b \xrightarrow{w} a$ in a quasi-building A and $\lambda \in \Gamma$ commuting with w, if $a \sim_\lambda a^* \in A$, then there is a unique b^* in A with $b \sim_\lambda b^* \xrightarrow{w} a^*$.

Proof. By iterating Lemma 2.6(c), the reduced path $b \xrightarrow{w} a \sim_\lambda a^*$ yields some b^* such that $b \sim_\lambda b^* \xrightarrow{w} a^*$. Furthermore, b^* is uniquely determined by b, a^* and the reduced word $\lambda \cdot w$, by Lemma 2.5. □

We will now produce certain extensions of a given quasi-building A. Fix λ in Γ and an equivalence class $a/{\sim_\lambda}$ in A. We will extend A to a quasi-building containing a new element in $a/{\sim_\lambda}$. Let B be the set of all x in A which are connected to a by a reduced path of type w, where w and λ commute. In particular, the generator λ does not occur in w. Furthermore, if b and c in B are $\sim_\gamma$-related, then λ and γ commute.

Observe that a lies in B. For every $b \in B$, introduce a new element b^*. Denote

$$A(a^*) = A \cup \{b^*\}_{b \in B},$$

and extend the chamber structure of A to $A(a^*)$ by setting

$$b^* \sim_\gamma c^* \;\Leftrightarrow\; b \sim_\gamma c$$

for all b, c in B, and

$$b^* \sim_\gamma a' \;\Leftrightarrow\; \lambda = \gamma \text{ and } b \sim_\lambda a'$$

for all $b \in B$ and $a' \in A$. In particular, if b is in B and $b \xrightarrow{w} a$, then we obtain a reduced path $b^* \xrightarrow{w} a^*$.

The extension $A(a^*)$ is called *simple*.

Lemma 2.8. *The chamber system $A(a^*)$ is a quasi-building.*

Proof. Properties (b) and (c) of Lemma 2.6 can be easily shown. For example, suppose that δ and γ_1 commute, and suppose $b^* \sim_{\gamma_1} d^* \sim_\delta c^*$. Then $b \sim_{\gamma_1} d \sim_\delta c$, so there is some d' in A such that $b \sim_\delta d' \sim_{\gamma_1} c'$, whence $b^* \sim_\delta (d')^* \sim_{\gamma_1} c^*$. The other cases are treated in a similar fashion. For (d), note first that any two elements of B are connected by a word which commutes with λ. Thus, a reduced path cannot change sides twice between A and $A(a^*) \setminus A$, for otherwise it would contain a reduced subpath of the form $\gamma \cdot w \cdot \gamma$, where w commutes with γ. A closed reduced path is hence either fully contained in A—and thus trivial since A is a quasi-building—or fully contained in $A(a^*) \setminus A$, in which case it is in bijection with a closed reduced path in A, which must then be trivial. □

Corollary 2.9. *For every Coxeter group* (W, Γ)*, there is a countable building in which all* $\sim_\gamma$*-equivalence classes are infinite.*

We will now show that every subset of a quasi-building has a strongly connected hull, which is attained by a sequence of simple extensions.

Proposition 2.10. *Given a strongly connected subset A of a quasi-building X and a^* in $X \setminus A$ with $a^* \sim_\lambda a \in A$, the smallest strongly connected subset of X containing $A \cup \{a^*\}$ is isomorphic to $A(a^*)$.*

Proof. Observe that A is a quasi-building, since X is. Let $B \subset A$ be as in the construction of $A(a^*)$, that is, the set of elements $b \in A$ with $b \xrightarrow{w} a$ where w commutes with λ.

Every b in B yields a unique b^* in X such that $b \xrightarrow{\lambda} b^* \xrightarrow{w} a^*$, by Remark 2.7. Since w is uniquely determined, up to permutation, by b (and a), the element b^* depends only on b. By symmetry, b is determined by b^*. Thus, the elements b^* are pairwise distinct. Note that none of the b^*'s belong to A, since $b^* \xrightarrow{w} a^* \xrightarrow{\gamma} a$ and A is strongly connected.

Therefore $A(a^*)$ may be identified with a subset of X which is contained in every strongly connected extension of $A \cup \{a^*\}$. We need only show that the chamber structure of $A \cup \{a^*\}$ agrees with the structure induced by X. Given distinct elements b and c in B with $b \sim_\delta c$, we need only show that $b^* \sim_\delta c^*$, since the converse follows by replacing a, b and c with a^*, b^* and c^*.

The set A is strongly connected, so there are reduced words u and v such that $a \xrightarrow{u} b$ and $a \xrightarrow{v} c$ in A.

Claim. *Let a, b and c be distinct elements of a quasi-building X with $b \sim_\delta c$. Suppose that $a \xrightarrow{u} b$ and $a \xrightarrow{v} c$ for reduced words u and v. Then there are three possibilities:*

(1) *$u \cdot \delta$ is reduced.*
(2) *$v \cdot \delta$ is reduced.*
(3) *There exists a reduced word $w \cdot \delta$ and an element a' such that $a \xrightarrow{w} a'$, $a' \xrightarrow{\delta} b$ and $a' \xrightarrow{\delta} c$.*

The third case implies that both u and v are equivalent to $w \cdot \delta$.

Proof of the Claim. If $u \cdot \delta$ is not reduced, up to permutation, we may assume that $u = w \cdot \delta$. Choose a' in X with $a \xrightarrow{w} a' \xrightarrow{\delta} b$. Thus $a' \sim_\delta c$. Either $a' = c$, so w is equivalent to v, which gives case (2), or $a' \xrightarrow{\delta} c$, which gives case (3). This proves the claim.

If we apply the Claim to our situation, we obtain three possibilities:

(1) The word $u \cdot \delta$ is reduced. Up to permutation, we have $v = u \cdot \delta$. In particular, the elements γ and δ commute. Since $b \sim_\delta c \sim_\gamma c^*$, there is some c' in X with $b \sim_\delta c' \sim_\gamma c$. Note that since $a^* \xrightarrow{v} c^*$ and $a^* \xrightarrow{u} b^* \xrightarrow{\delta} c'$, we have $a^* \xrightarrow{v} c'$ as well. Thus $c^* = c'$, by Remark 2.7, so $c^* \sim_\delta b^*$.

(2) The word $v \cdot \delta$ is reduced, which is treated similarly to the first case.

(3) For some reduced word $w \cdot \delta$, there is an a' such that $a \xrightarrow{w} a'$, $a' \xrightarrow{\delta} b$ and $a' \xrightarrow{\delta} c$. Since A is strongly connected, the element a' lies in A. By the first case, we have $(a')^* \sim_\delta b^*$ and $(a')^* \sim_\delta c^*$, whence $b^* \sim_\delta c^*$.

If $b^* \sim_\delta a'$ for some b in B and a' in A, then the path $b \sim_\lambda b^* \sim_\delta a'$ cannot be reduced, since b^* does not lie in A. Thus $\delta = \lambda$ and $b^* \sim_\lambda a'$. □

Corollary 2.9 and Proposition 2.10 immediately yield the following result:

Corollary 2.11 (cf. [7, Proposition 5.1]). *Every Coxeter group (W, Γ) has, up to isomorphism, a unique countable building $\mathrm{B}_0(\Gamma)$, in which all $\sim_\gamma$-equivalence classes are infinite.*

We will study the model theory of buildings using the following expansion of the natural language:

Definition 2.12. Let X be a chamber system for (W, Γ) and s a subset of Γ. We denote by $\sim_\emptyset$ the diagonal in $X \times X$. Otherwise, for $\emptyset \neq s \subset \Gamma$, the relation $\sim_s$ is the transitive closure of all $\sim_\gamma$ with $\gamma \in s \subset \Gamma$. The $\sim_s$-class of an element is called its *s-residue*. In particular, its *γ-residue* is its $\sim_\gamma$-class, which is often called the *γ-panel* in the literature.

For γ in Γ, set $\sim^\gamma = \sim_{\Gamma\setminus\{\gamma\}}$. The chamber system $(X, \sim^\gamma)_{\gamma\in\Gamma}$ is called the associated *dual chamber system* of X.

It is easy to see that $x \sim_s y$ if and only if $x \xrightarrow{w} y$ for some $w \in \langle s\rangle$. If X is a quasi-building, the word w is uniquely determined as an element of W, so Corollary 2.3 implies

$$\sim_{s_1\cap s_2} = \sim_{s_1} \cap \sim_{s_2},$$

and in particular

$$\sim_\gamma = \bigcap_{\beta\neq\gamma} \sim^\beta. \tag{†}$$

Thus, for a quasi-building X, the chamber system $(X, \sim_\gamma)_{\gamma\in\Gamma}$ is definable in its associated dual chamber system $(X, \sim^\gamma)_{\gamma\in\Gamma}$. Clearly, the latter is only definable if countable disjunctions are allowed.

The aim of this article is to study the complete theory of $\mathrm{B}^0(\Gamma)$, the associated dual chamber system of $\mathrm{B}_0(\Gamma)$ defined in Corollary 2.11.

Lemma 2.13. *The dual chamber system $(X, \sim^\gamma)_{\gamma\in\Gamma}$ of a quasi-building X has the following elementary properties:*

(1) *If $x, y \in X$ with $x \sim^\gamma y$ for all $\gamma \in \Gamma$, then $x = y$.*

(2) *If $(x_\gamma)_{\gamma\in\Gamma}$ is a* coherent *sequence in X, i.e. whenever γ and δ are adjacent, there exists some $y_{\gamma,\delta}$ in X with $x_\gamma \sim^\gamma y_{\gamma,\delta} \sim^\delta x_\delta$, then there exists $z \in X$ with $z \sim^\gamma x_\gamma$ for all $\gamma \in \Gamma$.*

Proof. Property (1) clearly follows from condition (†).

In order to prove (2), note that any singleton is a quasi-building and has property (2). By Proposition 2.10, we need only show that if A is a quasi-building with property (2), then so is $A(a^*)$.

Let $(x_\gamma)_{\gamma\in\Gamma}$ be a coherent sequence in $A(a^*)$, where $a^* \sim_\lambda a \in A$. Each x_γ lies either in A or is $\sim_\lambda$-related to some element in A. Since only the $\sim^\gamma$-class of x_γ matters, we may assume that all x_γ belong to A, for $\gamma \neq \lambda$. If x_λ is in A, then the result follows, since (2) holds in A. Otherwise, if $x_\lambda \notin A$, we have $x_\lambda \xrightarrow{w} a^*$ for a reduced word w which commutes with λ. Thus $x_\lambda \sim^\lambda a^*$ and we may assume that $x_\lambda = a^*$.

By assumption, for every γ adjacent to λ, there is some $y_{\lambda,\gamma}$ in $A(a^*)$ such that $a^* \sim^\lambda y_{\lambda,\gamma} \sim^\gamma x_\gamma$. The element $y_{\lambda,\gamma}$ cannot lie in A, for otherwise the reduced path $a \sim_\lambda a^* \sim^\lambda y_{\lambda,\gamma}$ implies that so does a^*. Thus $y_{\lambda,\gamma}$ is of the form $b^*_{\lambda,\gamma}$ for some $b_{\lambda,\gamma} \in A$.

It follows that $a \sim^\lambda b_{\lambda,\gamma} \sim^\gamma x_\gamma$. Replacing x_λ in the sequence (x_γ) by a yields a new sequence contained in A and coherent. Thus, we find c in A such that $c \sim^\lambda a$ and $c \sim^\gamma x_\gamma$ for $\gamma \neq \lambda$. Observe that $c^* \sim^\lambda a^*$ and $c^* \sim^\gamma x_\gamma$, so c^* is as desired. □

Definition 2.14. A chamber system $(X, \sim^\gamma)_{\gamma\in\Gamma}$ is a *dual quasi-building* if it has properties (1) and (2) from Lemma 2.13.

The following remark can easily be verified.

Remark 2.15. A chamber system $(X, \sim^\gamma)_{\gamma\in\Gamma}$ is a dual quasi-building if and only if for every coherent sequence $(x_\gamma)_{\gamma\in\Gamma}$ there is a unique z with $z \sim^\gamma x_\gamma$ for all $\gamma \in \Gamma$.

Definition 2.16. A Γ-*graph* M is a coloured graph with colours $\mathcal{A}_\gamma(M)$ for γ in Γ, and no edges between elements of $\mathcal{A}_\gamma(M)$ and $\mathcal{A}_\delta(M)$ if γ and δ are not adjacent.

A *flag* F of the Γ-graph M is a subgraph $F = \{f_\gamma\}_{\gamma\in\Gamma}$, where each f_γ lies in $\mathcal{A}_\gamma(M)$, such that the map $\gamma \mapsto f_\gamma$ induces a graph isomorphism between Γ and F.

The Γ-graph M is a Γ-*space* if the following two additional properties are satisfied:

(1) Every vertex belongs to a flag of M.
(2) Any two adjacent vertices in M can be expanded to a flag of M.

In particular, if Γ is the complete graph $\mathbb{K}_3$, then the following $\mathbb{K}_3$-graph is not a $\mathbb{K}_3$-space:

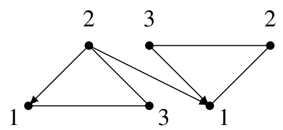

Theorem 2.17. *The class of dual quasi-buildings for* (W, Γ) *and the class of* Γ*-spaces are bi-interpretable.*

Proof. In a chamber system $(X, \sim^\gamma)_{\gamma\in\Gamma}$, we interpret a Γ-graph $\mathcal{M}(X)$ as follows: for every $\gamma \in \Gamma$, the colour $\mathcal{A}_\gamma$ is $X/\sim^\gamma$, the set of $\sim^\gamma$-classes of elements in X. We consider the $\mathcal{A}_\gamma$ as being pairwise disjoint. For the graph structure on $\mathcal{M}(X)$, we declare that two elements u and v are adjacent if $u \in \mathcal{A}_\gamma$ and $v \in \mathcal{A}_\delta$, the colours γ and δ are adjacent and there is some $z \in X$ with $z \sim^\gamma u$ and $z \sim^\delta v$.

Since every $x \in X$ gives rise to the flag $\phi(x) = \{x/\sim^\gamma \mid \gamma \in \Gamma\}$ of $\mathcal{M}(X)$, it is straightforward to see that $\mathcal{M}(X)$ is a Γ-space.

Given a Γ-graph M, we define a chamber system $(\mathcal{X}(M), \sim^\gamma)_{\gamma\in\Gamma}$, whose underlying set is the collection of flags of M. Two flags F and G are $\sim^\gamma$-related if their γ-vertices agree.

We will first show that $\mathcal{X}(M)$ is a dual quasi-building. We need only check property (2). Assume that (F_γ) is a coherent system of flags. Let f_γ denote the γ-vertex of F_γ. If γ and δ are adjacent, there is a flag G such that $F_\gamma \sim^\gamma G \sim^\delta F_\delta$. In particular, the flag G contains an edge between f_γ and f_δ. So $H = (f_\gamma)$ is a flag of M and $H \sim^\gamma F_\gamma$ for all $\gamma \in \Gamma$.

For a chamber system $(X, \sim^\gamma)_{\gamma\in\Gamma}$, the correspondence $x \mapsto \phi(x)$ defines a map $\phi : X \to \mathcal{X}(\mathcal{M}(X))$. It is easy to see that $\phi(x) \sim^\gamma \phi(y)$ if and only if $x \sim^\gamma y$. If X is a dual quasi-building, property (1) implies that ϕ is injective, and property (2) that ϕ is surjective. Thus X and $\mathcal{X}(\mathcal{M}(X))$ are definably isomorphic.

Given a Γ-graph M, the correspondence $a : \mathcal{M}(\mathcal{X}(M)) \to M$ which associates to each class $F/\sim^\gamma$ the γ-vertex of F is a bijection between $\mathcal{M}(\mathcal{X}(M))$ and the collection of vertices of M which belong to a flag of M. For adjacent γ and δ, there is an edge between $F/\sim^\gamma$ and $G/\sim^\delta$ if and only if $a(F/\sim^\gamma)$ and $a(G/\sim^\delta)$ belong to a common flag of M. This shows that a is a definable isomorphism if M is a Γ-space.

Thus, the classes of dual quasi-buildings and of Γ-spaces are bi-interpretable, as desired. □

In order to describe the model-theoretical properties of the dual quasi-building $\mathrm{B}^0(\Gamma)$ (introduced right before Lemma 2.13), we may therefore consider the first-order theory of

$$\mathrm{M}_0(\Gamma) = \mathcal{M}(\mathrm{B}^0(\Gamma)),$$

its associated Γ-space. Our reason to do this is that Γ-spaces, for certain Coxeter groups Γ, will be familiar to the readers of [3, 15, 2]. In particular, many of the tools developed in [2] can be easily generalised and adapted to this context. However, the whole model-theoretical study of $\mathrm{B}^0(\Gamma)$ could be done without passing to the corresponding Γ-space.

3. Simply connected Γ-spaces

Recall that by a Coxeter group we mean a right-angled finitely generated Coxeter group. From now on, fix a Coxeter group (W, Γ), which we will denote simply by W, with underlying Coxeter graph Γ. In order to describe the first-order theory of the structure $\mathrm{M}_0(\Gamma)$ obtained before, we will need to study non-standard paths between flags.

Notation. A *letter* is a non-empty connected subset of the graph Γ. Characters such as s and t will exclusively refer to letters. A *word* u is a finite sequence of letters.

Every generator γ in Γ defines the letter $\{\gamma\}$. In this way, every word in the generators can be considered as a word in the above sense.

Definition 3.1. Two letters s and t *commute* if $s \cup t$ is not a letter, i.e. if the elements of s commute with all elements of t. In particular, no letter commutes with itself. Two words commute if their letters respectively do. A word is *commuting* if it consists of pairwise commuting letters. A *permutation* of a word is obtained by repeatedly permuting adjacent commuting letters. Two words u and v are *equivalent*, denoted by $u \approx v$, if one can be permuted into the other.

The following is easy to see.

Remark 3.2. A commuting word $w = s_1 \cdots s_n$ is determined up to equivalence by its *support*

$$|w| = s_1 \cup \cdots \cup s_n,$$

where the s_i's are the connected components of $|w|$.

We will often write w instead of $|w|$ if w is a commuting word.

Throughout this section, we will work inside some ambient Γ-space.

Definition 3.3. A *weak flag path* P from the flag F to the flag G is a finite sequence $F = F_0, F_1, \ldots, F_n = G$ of flags such that the colours where F_{i+1} and F_i differ form a letter s_{i+1}. To such a path, we associate the word $u = s_1 \cdots s_n$ and denote this by $F \underset{u}{\rightarrow} G$.

In the light of Theorem 2.17, we transfer Definition 2.12 to Γ-spaces and say that two flags F and G are A-*equivalent*,

$$F \sim_A G,$$

if the set of colours where F and G differ is contained in $A \subset \Gamma$. Similarly to [2, Lemma 6.3], by decomposing any subset of Γ into the disjoint union of its connected components, we obtain the following.

Lemma 3.4. *Two flags F and G are A-equivalent if and only if they can be connected by a weak flag path whose word consists of letters contained in A. In particular, setting $A = \Gamma$, any two flags can be connected by a weak flag path.*

Proof. For $F = \{f_\gamma \mid \gamma \in \Gamma\}$ and $G = \{g_\gamma \mid \gamma \in \Gamma\}$, and let $s_1, \ldots, s_n$ be the connected components of $\{\gamma \in \Gamma \mid f_\gamma \neq g_\gamma\}$. Set

$$F_i = \{f_\gamma \mid \gamma \notin s_1 \cup \cdots \cup s_i\} \cup \{g_\gamma \mid \gamma \in s_1 \cup \cdots \cup s_i\}.$$

Then $F_0, \ldots, F_n$ is a weak path which connects F and G, and it has word $s_1 \cdots s_n$. □

The proof yields the following two corollaries.

Corollary 3.5. *Given flags F and G, there exists a commuting word u, unique up to equivalence, such that $F \underset{u}{\rightarrow} G$. For every permutation u' of u, there is a unique weak flag path from F to G with word u'.*

Uniqueness of u follows from the fact that the letters of u are the connected components of the set of colours where F and G differ.

Corollary 3.6. *In a path $F \underset{u}{\to} H \underset{v}{\to} G$, where u and v commute, the flag H is uniquely determined by F, G, u and v.*

To each relation $F \underset{s}{\to} G$ we associate the natural bijection $\mathcal{A}_s(F) \to \mathcal{A}_s(G)$. It is easy to see that a weak flag path $P : F_0 \underset{s_0}{\to} F_1 \underset{s_1}{\to} \cdots \underset{s_{n-1}}{\longrightarrow} F_n$ is completely determined by F_0 (likewise by F_n) and the sequence of associated maps. If $s_1 \cdots s_n$ is commuting, this sequence of maps is equivalent to the collection $\mathcal{A}_{s_i}(F_0) \to \mathcal{A}_{s_i}(F_n)$, which gives an alternative proof of Corollary 3.5. More generally, the following lemma holds.

Lemma 3.7 (Permutation of a path). *Given a weak flag path $P : F_0 \underset{u}{\to} F_n$ and a permutation u' of u, there is a unique weak flag path $P' : F_0 \underset{u'}{\to} F_n$ such that the associated map of each letter in P is the same as the associated map of the corresponding occurrence of that letter in P'.*

Such a path P' is a *permutation* of P.

Definition 3.8.

- A *splitting* of a letter s is a (possibly trivial) word whose letters are properly contained in s. Given words u and v, we write $u \prec v$ if u is equivalent to a word obtained from v by replacing at least one occurrence of a letter in v by a splitting. We write $u \preceq v$ if either $u \prec v$ or $u \approx v$.
- Whenever $F \underset{s}{\to} G$ and there is no weak flag path from F to G whose word is a splitting of s, we write $F \overset{s}{\to} G$. A *flag path* from F to G with word $u = s_1 \cdots s_n$, denoted by $F \overset{u}{\to} G$, is a weak flag path $F = F_0, \ldots, F_n = G$ such that $F_i \xrightarrow{s_{i+1}} F_{i+1}$ for $i = 0, \ldots, n-1$.

It is easy to see that the relation $\prec$ is transitive, irreflexive and well-founded [2, Lemma 5.26]. A permutation of a flag path is again a flag path. If $F \underset{s}{\to} G$, whether $F \overset{s}{\to} G$ depends on the ambient Γ-space.

Lemma 3.9. *If $F \underset{u}{\to} G$, then $F \overset{v}{\to} G$ for some $v \preceq u$.*

Proof. Suppose $F \underset{u}{\to} G$. If this is not a flag path, it contains a step $F' \underset{s}{\to} G'$ which can be replaced by $F' \underset{w}{\to} G'$, where w is a splitting of G. This yields a weak flag path $F \underset{u'}{\to} G$ with $u' \prec u$. Since $\prec$ is well-founded, this procedure stops with a flag path $F \overset{v}{\to} G$ for some word $v \preceq u$. □

Notation. The notation $s \subset t$ means that s is a subset of t, possibly with $s = t$. We will write $s \subsetneq t$ to emphasise that s is a proper subset of t.

Definition 3.10.

- A word $v = s_1 \cdots s_n$ is *reduced* if there is no pair $i \neq j$ such that $s_i \subset s_j$ and s_i commutes with all letters in v between s_i and s_j.

- A flag path is *reduced* if its associated word is.
- The reduced word v is a *reduct* of u if it can obtained from u by the following rules:

  Commutation: Permute consecutive commuting letters.
  Absorption: If s is contained in t, replace a subword $s \cdot t$ (or $t \cdot s$) by t.
  Splitting: Replace a subword $s \cdot s$ by a splitting of s.

  We will denote this by $u \xrightarrow{*} v$ [2, Definition 5.24]. Clearly $u \xrightarrow{*} v$ implies $v \preceq u$.

It is easy to see that a word u is reduced if and only if any permutation of u is. Similarly, a path P is reduced if and only if any permutation of P is.

Consider indices $i \neq j$ in a word $v = s_1 \cdots s_n$ such that $s_i \subset s_j$ and s_i commutes with all letters in v between s_i and s_j. Using Commutation and Absorption, we can delete the occurrence of the letter s_i. If $s_i = s_j$, we may also replace s_j by a splitting of s_j. We call such an operation a *generalised* Absorption or Splitting. It is easy to see that every reduct of a word can be obtained by a sequence of generalised Absorptions and Splittings, followed by a permutation.

Lemma 3.11. *If $F \xrightarrow{u} G$, then $F \xrightarrow{v} G$ for some reduced v with $u \xrightarrow{*} v$.*

Proof. If the path $F \xrightarrow{u} G$ is not reduced, possibly after permutation, we may assume that it contains a subpath $F' \xrightarrow{s} H' \xrightarrow{t} G'$, where $s \subset t$ (or $t \subset s$). One of the following reduction steps now applies:

Proper Absorption: If $s \subsetneq t$, remove H' since $F' \xrightarrow{t} G'$, for otherwise there would be a a splitting x of t such that $F' \underset{x}{\rightarrow} G'$, which implies $H' \underset{s \cdot x}{\longrightarrow} G'$, contradicting $H' \xrightarrow{t} G'$.
Absorption/Splitting: If $s = t$, note that $F' \sim_s G'$. Lemmata 3.4 and 3.9 yield:

  Absorption: $F' \xrightarrow{s} G'$, or
  Splitting: $F' \xrightarrow{x} G'$ for some splitting x of s.

  Therefore, the flag H' can be removed.

Note that both Absorption and Splitting yield words which are $\prec$-smaller than u. Thus, the process must eventually stop. □

Remark 3.12. We will see in Remark 4.15 that, for every reduction $u \xrightarrow{*} v$, there is a flag path of word u in a suitable Γ-space which can be reduced to a path with word v by the above procedure.

Corollary 3.13. *Any two flags can be connected by a reduced path.*

The following property of the ambient space will ensure that all reduced paths between two given flags have equivalent words (Proposition 3.19).

Definition 3.14. A Γ-space M is *simply connected* if there are no non-trivial closed reduced flag paths.

Lemma 3.15. *The Γ-space M is simply connected if and only if the word of any closed flag path can be reduced to the trivial word* 1.

Proof. Suppose the given condition holds. Given a closed reduced flag path with word u, since $u \xrightarrow{*} 1$, we have $u = 1$, as u is already reduced. For the other direction, given a closed path P with word u, apply Lemma 3.11 to obtain a closed reduced path whose word v is a reduct of u. If M is simply connected, the word v must be 1, thus $u \xrightarrow{*} 1$. □

Theorem 3.16. *The* Γ*-space* $\mathrm{M}_0(\Gamma)$ *is simply connected.*

Proof. By the definition of $\mathrm{M}_0(\Gamma)$, as explained in the proof of Theorem 2.17, two flags F and G in $\mathrm{M}_0(\Gamma)$ have the same γ-vertex if they can be connected by a flag path whose letters are singletons different from γ. Thus, if $F \xrightarrow{s} G$, then s must be a singleton. All paths are singleton paths. Since $\mathrm{B}_0(\Gamma)$ is a building, there are no non-trivial closed reduced singleton paths. □

An interesting feature of simply connected Γ-spaces is that the word of a reduced flag path connecting two given flags is unique, up to equivalence. For the proof of the next proposition, we need a definition and a lemma.

Definition 3.17. The letter t is (properly) *left-absorbed*, resp. *right-absorbed*, by the word $s_1 \cdots s_n$ if t is (properly) contained in some s_i and commutes with $s_1 \cdots s_{i-1}$, resp. with $s_{i+1} \cdots s_n$. A word u is *left-absorbed*, resp. *right-absorbed*, by v if each letter in u is.

A word is reduced if and only if it cannot be written as $u \cdot t \cdot v$, where t is either right-absorbed by u or left-absorbed by v. The word $u = s_1 \cdots s_n$ left-absorbs t if and only if u^{-1} right-absorbs t, where $u^{-1} = s_n \cdots s_1$.

Lemma 3.18. *Let* $u \cdot s$ *be reduced and* x *be a splitting of* s*. Every reduct of* $u \cdot x$ *is equivalent to* $u \cdot x_1$ *for some* $x_1 \preceq x$.

Proof. Since $u \cdot s$ is reduced, a generalised Absorption or Splitting for $u \cdot x$ cannot happen for a pair $s_i \subset s_j$, where s_i is contained in u. So letters contained in u will never be deleted. □

Proposition 3.19. *If the flags* F *and* G *in a simply connected* Γ*-space* M *are connected by reduced flag paths with respective words* u *and* v*, then* $u \approx v$.

Proof. We prove it by $\prec$-induction on u and v. If $F = G$, the claim is equivalent to simple connectedness. We may therefore assume that $F \neq G$. Since $u \cdot v^{-1}$ belongs to a closed non-trivial flag path, it cannot be a reduced word. So assume that $u = u_1 \cdot s$ and s is right-absorbed by v. Splitting the first path accordingly,

$$F \xrightarrow{u_1} F_1 \xrightarrow{s} G,$$

we distinguish two cases:

(1) The letter s is properly right-absorbed by v. Then v is the only reduct of $v \cdot s$, and therefore $F \xrightarrow{v} G \xrightarrow{s} F_1$ gives $F \xrightarrow{v} F_1$.

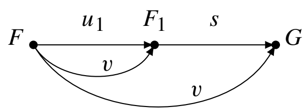

Since $u_1 \prec u$, induction yields $u_1 \approx v$. In particular, the letter s is right-absorbed by u_1, contradicting u being reduced.

(2) After a permutation v has the form $v_1 \cdot s$. We split the second path as

$$F \xrightarrow{v_1} G_1 \xrightarrow{s} G.$$

We then have either $G_1 \xrightarrow{s} F_1$ or $G_1 \xrightarrow{x} F_1$ for a reduced splitting x of s. If $G_1 \xrightarrow{s} F_1$, then $F \xrightarrow{v} F_1$, which contradicts $F \xrightarrow{u_1} F$ as before. So $G_1 \xrightarrow{x} F_1$.

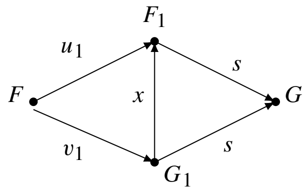

By Lemma 3.18 the path $F \xrightarrow{v_1} G_1 \xrightarrow{x} F_1$ reduces to a path $F \xrightarrow{v_1 \cdot x_1} F_1$ with $x_1 \preceq x$. Since $v_1 \cdot x_1 \prec v$, induction yields $v_1 \cdot x_1 \approx u_1$. So $v_1 \cdot x_1 \cdot s \approx u$ is reduced, which is only possible if $x_1 = 1$. Hence $v_1 \approx u_1$ and therefore $v \approx u$.

□

Definition 3.20. Given flags F and G in a simply connected Γ-space A, we say that the word u *connects* F *and* G if u is the word of a reduced path from F to G. Since u is uniquely determined up to equivalence, we denote it by $\mathrm{w}_A(F, G)$

In order to show that simple connectedness is an elementary property, we will first give a general description of a reduction of a flag path.

For $i \neq j$, a pair of letters s_i, s_j occurring in v is called *reduced* if either s_i and s_j are not comparable, or neither s_i nor s_j commute with all letters in between s_i and s_j. The word v is reduced if and only if all pairs of letters occurring in v are. A pair of two disjoint subwords w_1 and w_2 of a word v, possibly not reduced, is *reduced in* v if all pairs of letters s and t, where s occurs in w_1 and t occurs in w_2, are reduced in v. By a sequence of generalised Absorptions and Splittings applied to letters in w_1 and w_2, we may replace w_1, w_2 by a pair w_1^*, w_2^* which is reduced in the resulting word v^*. We call such a process a *reduction* of w_1, w_2 in v. If v is the word of a flag path, we call the corresponding transformation of the path also a reduction of w_1, w_2.

Lemma 3.21 (Reduction Lemma). *Let $w_1 \cdot w \cdot w_2$ be the (possibly non-reduced) word of a flag path, where both w_1 and w_2 are reduced. Then there are words u_1, v_1, u_2, v_2 and a reduction of w_1, w_2 in the path with resulting word $w_1^* \cdot w \cdot w_2^*$ such that:*

- $w_1 \approx u_1 \cdot v_1$ *and* $v_2 \cdot u_2 \approx w_2$,
- $w_1^* = u_1 \cdot x_1$ *for some* $x_1 \preceq v_1$, *and* $x_2 \cdot u_2 = w_2^*$ *for some* $x_2 \preceq v_2$,
- v_1 *and* v_2 *commute with* w,
- $|v_1|$ *and* $|v_2|$ *are contained in* $|w_1| \cap |w_2|$.

Proof. Since both w_1 and w_2 are reduced, we may assume that the words w_1^* and w_2^* are obtained by a sequence of generalised Absorptions and Splitting, each involving a letter in the first word and a letter in the last word, where after every step we apply a reduction to both the first and the last words. It is enough to show by induction that, at every intermediate step of the reduction, the word $w_1' \cdot w \cdot w_2'$ satisfies the conclusion of the lemma. Start by setting $u_j = w_j$ and $v_j = x_j$ the empty word, for $j = 1, 2$. Assume that $u_1, v_1, x_1, v_2, u_2, x_2$ witness this at the i^{th} step. In particular $w_1' = u_1 \cdot x_1$ and $x_2 \cdot u_2 = w_2'$, where $x_i \preceq v_i$ for $i = 1, 2$.

We will treat the case of a generalised Splitting and leave to the reader the easier case of a generalised Absorption. So suppose that $w_1'' \cdot w \cdot w_2''$ is obtained from $w_1' \cdot w \cdot w_2'$ by a generalised Splitting followed by a reduction of the first and last words. Then there is a letter s occurring in both w_1' and w_2' which commutes with w as well as with the letters of w_1' to its right (resp. the letters of w_2' to its left). Suppose furthermore that the word w_1'' is obtained from w_1' by deleting s. Note that w_1'' is reduced.

The word w_2'' is obtained from w_2' by replacing s by a splitting y of s followed by a further reduction. By Lemma 3.18, we may assume that w_2'' is reduced and obtained, up to a permutation, by replacing s with some word $y_2 \preceq y \prec s$. If s occurred in u_2, then let u_2' be the word obtained by removing s from u_2 and set $v_2' = v_2 \cdot s$ and $x_2' = x_2 \cdot y_2$. If s occurred in x_2, then replace s by y_2 and leave v_2 and u_2 unchanged.

Likewise, modify the words u_1, v_1 and x_1 accordingly. □

Definition 3.22. Given a letter s and a natural number n, the reduced word w *satisfies* $E(s, n)$ if $|w| \subset s$ and no permutation of w is a product of n words u_i with $|u_i| \subsetneq s$.

The properties $E(s, n)$ get stronger as n increases. In particular, the word w satisfies $E(s, 0)$ if and only if $|w| \subset s$ and $w \neq 1$. Similarly, the word w satisfies $E(s, 1)$ if and only if $|w| = s$.

Corollary 3.23. *Let $u = s_1 \cdots s_n$ be a reduced word and $w_1, \ldots, w_n$ reduced words with $|w_k| = s_k$. Consider two indices $i < j$ and a reduction w_i^*, w_j^* of the pair w_i, w_j in $w_i \cdots w_j$ as in the Reduction Lemma 3.21.*

1. *If w_i satisfies $E(s_i, m)$ for some $m > 0$, then w_i^* satisfies $E(s_i, m-1)$, and similarly for w_j and w_j^*.*
2. *Assume that w_i and w_j satisfy $E(s_i, 2)$ and $E(s_j, 2)$, respectively. If a pair $w_{i'}, w_{j'}$ is already reduced in $w = w_1 \cdots w_n$, then the corresponding pair remains reduced in $w^* = w_1 \cdots w_i^* \cdots w_j^* \cdots w_n$.*

Proof. In order to prove (1), choose $u_i, v_i, x_i, u_j, v_j, x_j$ as in the Reduction Lemma. Since $|v_i|$ is contained in $s_i \cap s_j$ and v_i commutes with $s_{i+1} \cdots s_{j-1}$, $|v_i|$ is a proper subset of s_i, for otherwise the pair s_i, s_j would not be reduced in u. As $w_i \approx u_i \cdot v_i$, if w_i has property $E(s_i, m)$, then u_i has property $E(s_i, m-1)$, and so does w_i^*.

For (2), assume that w_i and w_j satisfy $E(s_i, 2)$ and $E(s_j, 2)$, respectively. Therefore w_i^* and w_j^* have property $E(s_i, 1)$ and $E(s_j, 1)$, respectively, so $|w_i^*| = s_i$ and $|w_j^*| = s_j$. Suppose now that the pair $w_{i'}, w_{j'}$ is reduced in w. Since the words w_i^* and w_j^* commute with the same letters as w_i and w_j, respectively, the pair $w_{i'}, w_{j'}$ remains reduced in w^* if $\{i', j'\}$ and $\{i, j\}$ are disjoint. By symmetry, it suffices to consider the following three other cases:

Case $i' < j' = i$: We have to show that the pair $w_{i'}, u_i \cdot x_i$ is reduced in w^* if the pair $w_{i'}, u_i \cdot v_i$ is reduced in w. This follows easily from $|x_i| \subset s_i = |u_i|$, as w_i satisfies $E(s_i, 2)$.

Case $i = i' < j' < j$: Here we have to show that the pair $u_i \cdot x_i, w_{j'}$ is reduced in w^* if $u_i \cdot v_i, w_{j'}$ is reduced in w. This follows easily from the fact that v_i and x_i commute with $w_{j'}$.

Case $i = i' < j < j'$: Again we have to show that the pair $u_i \cdot x_i, w_{j'}$ is reduced in w^* if $u_i \cdot v_i, w_{j'}$ is reduced in w. This follows easily from $|x_i|, |v_i|, |v_j| \subset |u_j|$. □

Proposition 3.24. *Let $u = s_1 \cdots s_n$ be a reduced word. Given reduced flag paths $F_{i-1} \xrightarrow{w_i} F_i$ such that each w_i has property $E(s_i, n)$, the path*

$$F_0 \xrightarrow{w_1} F_1 \xrightarrow{w_2} \cdots \xrightarrow{w_n} F_n$$

has a reduction of length $\geq n$.

Proof. Choose any enumeration of all pairs $i < j$ of indices between 1 and n and apply the Reduction Lemma to each pair in order. Observe that k occurs in at most $n - 1$ reductions. At every step of the reduction, the resulting words satisfy $E(s_k, 2)$, so the resulting path

$$F_0 \xrightarrow{w_1^*} F_1^* \xrightarrow{w_2} \cdots \xrightarrow{w_n^*} F_n$$

is reduced, by Corollary 3.23(2). None of the words w_i^* is trivial, by Corollary 3.23(1). □

Together with Remark 4.15, the previous proposition will imply the following (a priori) stronger form.

Remark 3.25. Let $u = s_1 \cdots s_n$ be a reduced word and $w_1, \ldots, w_n$ reduced words with property $E(s_i, n)$. Then every reduct of $w_1 \cdots w_n$ has length at least n.

Theorem 3.26. *Simple connectedness is an elementary property for Γ-spaces.*

Proof. For each natural number n and letter s, consider the following elementary property of flags F and G: We have $F \sim_s G$, but there exist no proper subsets $A_1, \ldots, A_n$ of s and flags $F_1, \ldots, F_{n-1}$ such that

$$F \sim_{A_1} F_1 \sim_{A_2} \cdots \sim_{A_{n-1}} F_{n-1} \sim_{A_n} G.$$

We then write $F \xrightarrow{s,n} G$. Observe that if $F \xrightarrow{s,n} G$, then there is a path $F \xrightarrow{w} G$ for some reduced w which satisfies property $E(s, n)$. Indeed, since $F \sim_s G$, there exists a reduced word w connecting F to G with support contained in s. Any such word satisfies $E(s, n)$.

It suffices to show that a Γ-space M is simply connected if and only if for all natural numbers n and all non-trivial reduced words $u = s_1 \cdots s_n$, there is no sequence $F_0 \xrightarrow{s_1,n} \cdots \xrightarrow{s_n,n} F_n = F_0$.

Clearly, right-to-left is obvious, since $F \xrightarrow{s} G$ implies $F \xrightarrow{s,n} G$, by Lemma 3.4. Suppose now that M is simply connected and let $F_0 \xrightarrow{s_1,n} \cdots \xrightarrow{s_n,n} F_n$ be a weak flag path for some non-trivial reduced word $u = s_1 \cdots s_n$. By the above discussion, there are words w_i, each satisfying property $E(s_i, n)$, respectively, such that

$$F_0 \xrightarrow{w_1} \cdots \xrightarrow{w_n} F_n.$$

Proposition 3.24 shows that this path has a reduction of length at least n, so $F_0 \neq F_n$, since M is simply connected. □

Simple connectedness allows us to generalise [2, Remark 4.9], which will be needed for the proof of Proposition 4.7.

Lemma 3.27. *Given two adjacent colours γ and δ, if M is a simply connected Γ-space, the subgraph $\mathcal{A}_{\gamma,\delta}(M) = \mathcal{A}_\gamma(M) \cup \mathcal{A}_\delta(M)$ has no non-trivial circles.*

Proof. Since any edge in $\mathcal{A}_{\gamma,\delta}$ lies within a flag in M, by property (2) of Definition 2.16, a path with no repetitions in $\mathcal{A}_{\gamma,\delta}$ induces a (possibly non-reduced) flag path of the form

$$F_0 \xrightarrow{w_1} \cdots \xrightarrow{w_n} F_n,$$

where the reduced words $w_1, \ldots, w_n$ have the properties

$$\delta \in |w_{2k+1}| \subset \Gamma \setminus \{\gamma\}, \qquad \gamma \in |w_{2k}| \subset \Gamma \setminus \{\delta\}.$$

By repeatedly applying Lemma 3.21 to each pair w_i, w_j for $i \neq j$, it is easy to see that the above conditions still hold in the reduct $w_1^* \cdots w_n^*$ of the word $w_1 \cdots w_n$. In particular, the word $w_1^* \cdots w_n^*$ is not trivial, and thus $F_0 \neq F_n$. Hence, the original path in $\mathcal{A}_{\gamma,\delta}$ was not closed. □

4. The theory PS_Γ

Definition 4.1. In the language of graphs enriched with unary predicates for the colours $\{\mathcal{A}_\gamma\}_{\gamma\in\Gamma}$, let the theory PS_Γ be a collection of sentences stating that the structure is a Γ-space with the following properties:

(1) simple connectedness,
(2) for any colour γ in Γ, the $\sim_\gamma$-class of any flag G is infinite (observe that the relation $\sim_\gamma$ is definable in this language).

Axiom (1) is a first-order property, by Theorem 3.26. Clearly, so is axiom (2). The Γ-space $\mathrm{M}_0(\Gamma)$, as defined on page 3100, is a model of PS_Γ by Theorem 3.16, so PS_Γ is consistent.

The rest of this section is devoted to proving the completeness of PS_Γ.

We first generalise [2, Definition 4.3].

Definition 4.2. Fix some letter s, and let F be a flag in a Γ-graph A. Create a new flag $F^* = \{f^*_\gamma\}_{\gamma\in\Gamma}$ which agrees with F on the colours of $\Gamma \setminus s$ but $f^*_\gamma \notin A$ for $\gamma \in s$. We define a Γ-graph $A(F^*)$ with vertices $A \cup F^*$ and edges those of A and of F^*. A Γ-graph $B \supset A$ is a *simple extension* of A of *type* (s, F) if it is A-isomorphic to $A(F^*)$.

Note that $F^* \underset{s}{\rightarrow} F$ by construction.

Remark 4.3. If A is a Γ-space, then so is $A(F^*)$.

These simple extensions generalise those defined after Remark 2.7, as the following easy remark shows.

Remark 4.4. Let $\partial s = s \cup \{\delta \in \Gamma \mid \delta \text{ is adjacent to some } \gamma \in s\}$ denote the set of all γ in Γ which do not commute with s. Let H be a flag in A which agrees with F on the colours in ∂s. For each $\gamma \in s$, replace h_γ in H by f^*_γ in order to obtain a flag H^* of $A(F^*)$. It is easy to see, since s is connected, that that this construction defines a 1-to-1-correspondence between the flags H of A with $H \sim_{\Gamma\setminus\partial s} F$ and the new flags H^* of $A(F^*)$. Note that H^* is uniquely determined by

$$H^* \sim_{\Gamma\setminus\partial s} F^* \quad \text{and} \quad H^* \sim_s H.$$

In order to prove that the theory PS_Γ is complete, we will need the appropriate interpretation of a strongly connected subset in this context.

Definition 4.5. A non-empty subgraph D of a Γ-space M is *nice* if it satisfies the following conditions:

- Any point a in D lies in a flag in D.
- Given flags F and G in D and a letter s, if $F \xrightarrow{s} G$ in D, then $F \xrightarrow{s} G$ in M.

Any nice set is the union of all the flags contained in it. Niceness is a transitive property. Proposition 3.19 implies that a non-empty subset D of a simply connected Γ-space M is nice if and only if the following hold:

- Any point a in D lies in a flag in D.
- Given flags F and G in D and a reduced word u, if $F \xrightarrow{u} G$ in M, then there exists such a path in D with the same word.

In particular, if D is nice in M, then D is simply connected whenever M is.

Remark 4.6. (1) The Γ-space A is nice in $A(F^*)$.
(2) A nice subset of a Γ-space is itself a Γ-space.

Proof. For (1), given flags G and H in A with $G \xrightarrow{t} H$ in A, suppose there is a splitting $x = t_1 \cdots t_n$ of t such that

$$G = F_0 \underset{t_1}{\rightarrow} \cdots \underset{t_n}{\rightarrow} F_n = H$$

in $A(F^*)$. We replace each F_i by a flag F'_i in A as follows: If F_i belongs to A, set $F'_i = F_i$. Otherwise, by Remark 4.4, the flag F_i has the form H^*_i and we set $F'_i = H_i$. Note that $F'_{i-1} \sim_{t_i} F'_i$, so G and H can be connected in A be a weak flag path whose word is $\preceq$-smaller than $t_1 \cdots t_n$, contradicting $G \xrightarrow{t} H$.

In order to show (2), consider two elements a and b in a nice subset A of M, connected by a flag G in M. Choose flags F and H in A containing a and b respectively. Let γ be the colour of a, and δ the colour of b. Since $F \sim_{\Gamma\setminus\{\gamma\}} G \sim_{\Gamma\setminus\{\delta\}} H$, there is a reduced path $F \xrightarrow{u} G' \xrightarrow{v} H$ in M such that γ does not occur in u, and δ does not occur in v. By niceness, we may therefore assume that G' belongs to A. Clearly G' contains a and b. □

We can now prove the analogue of [2, Lemma 4.21].

Proposition 4.7. *Given a flag F in a nice subset A of a simply connected Γ-space M, and a flag F^* in M which is s-equivalent to F for some letter s, the following are equivalent:*

(a) *The Γ-graph $A \cup F^*$ is a simple extension of A of type (s, F) and is nice in M.*
(b) *Whenever G is a flag in A and x a splitting of s, we have $F^* \underset{x}{\not\to} G$ in M.*

Proof. (a)⇒(b): Set $B = A \cup F^*$. If

$$F^* \underset{x}{\to} G$$

in M, then we cannot have $F^* \xrightarrow{s} G$ in B, for B is nice. Thus, there is a splitting x' of s such that $F^* \xrightarrow{x'} G$ in B. All flags in B which are s-equivalent to F are either F^* itself or contained in A, by Remark 4.4, so $F^* \sim_t G'$ for some G' in A, where $t \subsetneq s$ is the first letter of x'. This is impossible, for no vertex of F^* with colour in s lies in A.

(b)⇒(a): Since F lies in A, the hypothesis implies that $F^* \xrightarrow{s} F$. We will first show that, for any flag G in A, if $F^* \xrightarrow{u} G$ in M, then $s \preceq u$. We may assume that u is reduced. Since A is nice, there is a reduced path $F \xrightarrow{v} G$ in A which remains reduced in M. The word u is thus a reduct of $s \cdot v$. If in the reduction splitting ever occurs, it produces a flag in A which connects to F^* by a splitting of s, contradicting the assumption. Thus, the reduction involves only commutation and possibly absorption of s by u. Hence $s \preceq u$.

In particular, for any flag G in A with $F^* \xrightarrow{u} G$, there exists a letter t in u containing s. Actually, by Lemma 3.9 it suffices to assume $F^* \underset{u}{\to} G$.

In order to show that $B = A \cup F^*$ is a simple extension of A, we need to show that there is no new edge consisting of an element b in $\mathcal{A}_\gamma(B) \setminus A$ and some a in $\mathcal{A}_\delta(A) \setminus F$. Suppose otherwise that there exists a flag F' in M passing through a and b. Take a flag G in A containing a.

Note that γ is in s. Suppose first that δ lies in s as well. Since $F^* \sim_{\Gamma\setminus\{\gamma\}} F' \sim_{\Gamma\setminus\{\delta\}} G$, we obtain reduced words u and v such that γ does not occur in u, the colour δ does not occur in v and

$$F^* \xrightarrow{u} F' \xrightarrow{v} G.$$

The reduction $F^* \xrightarrow{w} G$ in M has the property that every letter in w does not contain either γ or δ, contradicting the previous discussion.

If δ does not lie in s, then, with the choice of flags as before, we obtain the path

$$F \xrightarrow{s} F^* \xrightarrow{u} F' \xrightarrow{v} G.$$

As before, since A is nice, this implies that $F \xrightarrow{w} G$ in A, where each letter in w avoids either γ or δ. This induces a path in $\mathcal{A}_{\gamma,\delta}(A)$ between a' and a, where a' is the δ-vertex of F. This, together with the connection $a - b - a'$, yields a non-trivial circle, contradicting Lemma 3.27.

Let us now show that B is nice in M. Given flags $G_1 \xrightarrow{t} G_2$ in B, we distinguish the following cases:

- Both flags lie in A. Then $G_1 \xrightarrow{t} G_2$ also in A, and thus in M, since A is nice.
- None of the flags lies in A. By Remark 4.4, we have $G_1 \sim_{\Gamma\setminus\partial s} G_2$, and hence $t \subset \Gamma \setminus \partial s$. Thus we find H_1 and H_2 in A such that $H_1 \sim_{\Gamma\setminus\partial s} H_2$ and $G_i \sim_s H_i$ for $i = 1, 2$. This implies $H_1 \xrightarrow{t} H_2$ in B and also in A. Therefore $H_1 \xrightarrow{t} H_2$ in M, for A is nice, which implies that $G_1 \xrightarrow{t} G_2$ in M as well.
- Exactly one flag, say G_1, is not fully contained in A. Again by Remark 4.4, we see that $F^* \xrightarrow{w} G_1$ for a word w which commutes with s and a flag H_1 in A with $G_1 \underset{s}{\rightarrow} H_1$. Since $F^* \underset{w\cdot t}{\longrightarrow} G_2$, some letter of $w \cdot t$ must contain s, so $s \subset t$. If $s = t$ but $G_1 \xrightarrow{x} G_2$ in M for some $x \prec t = s$, then $F^* \xrightarrow{w} G_1 \xrightarrow{x} G_2$ in M, whose reduction yields a word where no letter contains s. Thus $G_1 \xrightarrow{t} G_2$ in M. Otherwise, if $s \subsetneq t$, then $H_1 \xrightarrow{t} G_2$ in B. Since A is nice, we have $H_1 \xrightarrow{t} G_2$ in M, which implies $G_1 \xrightarrow{t} G_2$ in M. □

In particular, setting $s = \{\gamma\}$ for γ in Γ, we deduce the following result.

Corollary 4.8. *Given a flag G in a nice subset A of a simply connected Γ-space M and γ in Γ, if the flag F is $\{\gamma\}$-equivalent to G and the γ-vertex of F does not lie in A, then the set $B = A \cup F$ is nice and a simple extension of A of type $(\{\gamma\}, G)$.*

The next proposition shows that a simply connected space is the increasing union of simple extensions of nice subsets (cf. [2, Theorem 4.22]).

Proposition 4.9. *Given a nice subset A of a simply connected Γ-space M and b in M, there exists a nice subset B containing b which can be obtained from A by a finite number of simple extensions.*

Proof. Given a flag F in M containing b, choose a reduced path

$$F = F_0 \xrightarrow{s_1} \cdots \xrightarrow{s_n} F_n$$

connecting F to a flag F_n in A such that the word $u = s_1 \cdots s_n$ is $\prec$-minimal. We prove the claim by $\prec$-induction on u. If $u = 1$, there is nothing to show. Otherwise, minimality of u implies that there is no path which connects F_{n-1} to a flag in A whose word is a splitting of s_n. It follows from Proposition 4.7 that $A' = A \cup F_{n-1}$ is a simple extension of A of type (s_n, F_n), so A' is nice in M. Now F can be connected to some flag in A' by a reduced word path, whose word u' is $\prec$-minimal such, with $u' \preceq s_1 \cdots s_{n-1} \prec u$. By induction, the element b is contained in a nice set B which can be obtained from A' (and thus from A) by a finite number of simple extensions. □

Lemma 4.10. *Let s be a letter in Γ which is not a singleton. There are distinct γ and γ' in s such that both $s \setminus \{\gamma\}$ and $s \setminus \{\gamma'\}$ are connected.*

Proof. By repeatedly removing edges, it is enough to prove this for a *spanning tree* s, that is, whenever we remove an edge between two points in s, the resulting graph is no longer connected. In particular, such an s contains no cycles. The assertion now follows, since any non-trivial tree has at least two extremal points. □

Corollary 4.11. *Given a letter s and a flag G contained in some finite nice subset A of an ω-saturated model M of PS_Γ, the model M contains a simple extension of A of type (s, G) which is nice in M.*

Proof. For $s = \{\gamma\}$, pick any flag G^* which is γ-equivalent to G and its γ-vertex does not lie in (the finite set) A, by property (2) of Definition 4.1. The set $A \cup G^*$ is nice in M and a simple extension of A of type (s, G), by Corollary 4.8.

Suppose now that $|s| \geq 2$. By Proposition 4.7 and saturation, it is enough to produce, for every n, a flag G_n with $G_n \underset{s}{\rightarrow} G$ and $G_n \underset{x}{\not\rightarrow} G'$, whenever G' is a flag in A and x a splitting of s of length at most n.

By Lemma 4.10, find two subletters s_0 and s_1 of s of cardinality $|s| - 1$ and not commuting with each other. Since A contains only finitely many flags, simple connectedness yields an upper bound N for the length of the word of any reduced flag path between any two flags in A. Set $A_0 = A$ and $G_0 = G$. By induction on $|s|$, there is a sequence $\{(G_i, A_i)\}_{i \leq N+n}$ of pairs such that G_i is a flag in the (finite) nice set A_i and $A_{i+1} = A_i \cup G_{i+1}$ is a simple extension of type $(s_{f(i)}, G_i)$, where $f(i)$ in $\{0, 1\}$ is the residue of i modulo 2. The flag path

$$G_{N+n+1} \xrightarrow{s_{f(N+n)}} \cdots \xrightarrow{s_0} G_0$$

is reduced, since s_0 and s_1 do not commute.

Clearly $G_{N+n+1} \underset{s}{\rightarrow} G_0$. Suppose there is some flag G' in A such that

$$G_{N+n+1} \underset{x}{\rightarrow} G'$$

in A for some splitting x of s of length at most n. Reducing this path, we obtain a reduced path $G' \xrightarrow{u} G$, where u is also a splitting of s. By niceness of A, we may assume that this path lies in A, so u has length at most N. Proposition 3.19 implies that $x \cdot u \overset{*}{\rightarrow} s_{f(N+n)} \cdots s_0$. However, at every step of the reduction, the number of letters of size exactly $|s| - 1$ is bounded by $N + n$, contradicting our choice of $s_{f(N+n)} \cdots s_0$. □

We can now conclude that the theory PS_Γ is complete and that the type of a nice set is determined by its quantifier-free type.

Theorem 4.12. *Any two ω-saturated models of PS_Γ have the back-and-forth property with respect to the collection of partial isomorphisms between finite nice substructures. In particular, any partial isomorphism $f : A \to A'$ between finite nice subsets of two models of PS_Γ is elementary. The theory PS_Γ is complete.*

Proof. Let M and M' be two ω-saturated models and let $f : A \to A'$ be a partial isomorphism between finite nice substructures. Given b in M, Proposition 4.9 yields a nice subset B containing $A \cup \{b\}$ such that $A \leq B$ in finitely many steps. By ω-saturation of M' and Corollary 4.11 (applied finitely many times), we obtain a nice subset B' of M' containing A' such that f extends to an isomorphism between B and B'.

Since any model M is nice in any elementary extension, replacing the models by appropriate saturated extensions, we produce a back-and forth system. Completeness of PS_Γ then follows, since any two flags have the same quantifier-free type. □

Corollary 4.13. *The type of a nice set A is determined by its quantifier-free type.*

Proof. For finite sets, this follows from Theorem 4.12. For infinite nice sets, note that they are direct unions of finite nice subsets. □

Corollary 4.14. *The theory PS_Γ is ω-stable and the model $\mathrm{M}_0(\Gamma)$ is the unique (up to isomorphism) countable prime model.*

Proof. In order to show that PS_Γ is ω-stable, we need to count 1-types over a countable subset of A, which we may assume is nice inside a given saturated model [16, Theorem 5.2.6]. Every simple extension of A is uniquely determined, up to A-isomorphism, by its type (s, G), by Corollary 4.13. Therefore, if A is countable, there are, up to A-isomorphism, only countably many simple extensions. Now, every 1-type is realised in a finite tower of simple extensions over A, by Proposition 4.9, so there are only countably many types, as desired.

In order to show that the countable model $\mathrm{M}_0(\Gamma)$ is the the prime model of PS_Γ, it suffices to show that it is constructible. It follows from the proof of Theorem 3.16 that the only words of reduced paths in $\mathrm{M}_0(\Gamma)$ are finite products of singletons. Given γ in Γ and a flag G in a nice subset A, the type over A of the simple extension $B = A \cup F$ of type $(\{\gamma\}, G)$ is determined by its quantifier-free type, which amounts to saying, by Corollary 4.8, that F and G are γ-equivalent but the γ-vertex of F does not lie in A. Therefore, if A is finite, the type of B over A is isolated. □

Uniqueness of prime models and Theorem 2.17 yield the uniqueness result in Corollary 2.11.

Remark 4.15. Let M be an ω-saturated model of PS_Γ. For every word v, there is a path $F \xrightarrow{v} G$ in M. Whenever a (possibly non-reduced) word u can be reduced to v, there is a flag path from F to G with word u.

5. Non-splitting reductions

In order to describe the geometrical complexity of PS_Γ, we will need several auxiliary results on the combinatorics of reduction of words when no splitting occurs, generalising some of the results of [2]. For the sake of self-containment, we will provide different proofs whenever possible.

Definition 5.1. A letter s is a *beginning* (resp. *end*) of the word u if $u \approx s \cdot v$ (resp. $u \approx v \cdot s$). The *initial segment* (resp. *final segment*) of u is the commuting subword whose letters are beginnings (resp. ends) of u.

By abuse of language, we say that the word u is an *initial subword* of v if u is an initial subword (in the proper sense) of some permutation of v. Likewise for *final subword*. The initial segment of v is the largest commuting initial subword of v. Inductively on the sum of their lengths, it is easy to see that any two words u and v have a *largest common initial subword*, resp. a *largest common final subword*, none of which need be commutative.

Common initial subwords can be removed, as seen easily.

Lemma 5.2. *If $u \cdot v' \approx u \cdot v''$, then $v' \approx v''$. Likewise, if $v' \cdot u \approx v'' \cdot u$, then $v' \approx v''$.*

Lemma 5.3. *Let a be a final subword of $c \cdot d$ such that every end of a commutes with d. Then a and d commute and a is a final subword of c.*

Proof. Proceed by induction on the length of a. If a is the trivial word, there is nothing to show. Otherwise, write $a = s \cdot a_1$. By induction, the subword a_1 commutes with d and is a final subword of $c \approx c' \cdot a_1$. Since $c \cdot d \approx c' \cdot d \cdot a_1$, Lemma 5.2 implies that s is an end of $c' \cdot d$.

If s occurs in d, then s and a_1 commute, so s is an end of a and hence it must commute with d, by hypothesis, which is a contradiction, since no letter commutes with itself. Therefore, the letter s must occur in c', so in particular it must commute with d. Thus, the word a commutes with d and is a final subword of c, as desired. □

Definition 5.4. A *non-splitting reduction* of the word u, denoted by $u \to v$, is a reduced word v obtained from u by a reduction process, where only Commutation and Absorption occur (see Definition 3.10).

By induction on the length of the reduction $u \to v$, the following can be easily shown.

Lemma 5.5. *If v is a non-splitting reduction of u, then every $u' \preceq u$ has a non-splitting reduction $v' \preceq v$.*

Corollary 5.6 (cf. [2, Proposition 5.3]). *Every word admits exactly one non-splitting reduction, up to permutation.*

Proof. Assume that v_1 and v are two non-splitting reducts of u. In particular $v_1 \preceq u$, so $v_1 \to v'$ for some $v' \preceq v$, by Lemma 5.5. Thus v' must be equivalent to v_1, and hence $v_1 \preceq v$. Similarly, we obtain $v \preceq v_1$, so $v_1 \approx v$. □

Notation. Given reduced words u and v, we denote by $[u \cdot v]$ the *non-splitting reduct* of $u \cdot v$, which is defined up to permutation.

Corollary 5.7 (cf. [2, Lemma 5.29]). *Given reduced words u, v and x with $x \preceq v$, we have $[u \cdot x] \preceq [u \cdot v]$.*

Corollary 5.8. *If u and v are reduced words, then $u \preceq [u \cdot v]$.* □

Recall (Definition 3.17) that a letter t is properly left-absorbed, resp. right-absorbed, by the word $s_1 \cdots s_n$ if and only if t is properly contained in some s_i and commutes with $s_1 \cdots s_{i-1}$, resp. with $s_{i+1} \cdots s_n$.

Proposition 5.9 (Symmetric Decomposition Lemma, cf. [2, Corollary 5.23]). *Given reduced words u and v, there are unique decompositions (up to permutation)*

$$u = u_1 \cdot u' \cdot w, \qquad w \cdot v' \cdot v_1 = v$$

such that:

(a) *w is a commuting word,*
(b) *u' is properly left-absorbed by v_1,*
(c) *v' is properly right-absorbed by u_1,*
(d) *u', w and v' pairwise commute,*
(e) *$u_1 \cdot w \cdot v_1$ is reduced.*

Furthermore,

$$[u \cdot v] = u_1 \cdot w \cdot v_1.$$

Proof. We show the existence of such a decomposition by induction on the sum of the lengths of u and v. If $u \cdot v$ is already reduced, then set $u_1 = u$ and $v_1 = v$, and define v', u' and w to be the trivial word.

Otherwise, up to permutation and changing the roles of u and v , we may assume that $u = \tilde{u} \cdot s$, where s is left-absorbed by v. In particular $[s \cdot v] = v$. By induction, we find words $\tilde{u}_1$, $\tilde{u}'$, $\tilde{w}$, $\tilde{v}'$ and $\tilde{v}_1$ with the desired properties such that

$$\tilde{u} = u_1 \cdot \tilde{u}' \cdot \tilde{w}, \qquad \tilde{w} \cdot v' \cdot \tilde{v}_1 = v.$$

Observe that s cannot be left-absorbed by $\tilde{w}$ nor by $\tilde{v}'$, for u is reduced. Thus s commutes with $\tilde{w} \cdot \tilde{v}'$ and is absorbed by $\tilde{v}_1$. There are two possibilities:

(1) The letter s is properly absorbed by $\tilde{v}_1$. Set $u' = \tilde{u}' \cdot s$, $w = \tilde{w}$ and $v_1 = \tilde{v}_1$.
(2) Up to permutation, the letter s is a beginning of $\tilde{v}_1 = s \cdot v_1^*$. Set $u' = \tilde{u}'$, $w = \tilde{w} \cdot s$ and $v_1 = v_1^*$.

Let us finish by showing the uniqueness of the above decomposition. Note first that $[u \cdot v] = u_1 \cdot w \cdot v_1$, since $u \cdot v \to u_1 \cdot w \cdot v_1$. The word w is exactly the intersection of the final segments of u and v^{-1}, since $u_1 \cdot w \cdot v_1$ is reduced. Furthermore, the word $u_1 \cdot w$ is the largest common initial subword of u and $[u, v]$, for otherwise there exists a letter s with $v_1 = s \cdot \tilde{v}_1$ and $u' = s \cdot \tilde{u}'$, contradicting u' being properly left-absorbed by v_1. Similarly, the word $w \cdot v_1$ is the largest common final subword of v and $[u, v]$, providing the desired result. □

Corollary 5.10 (cf. [2, Corollary 5.14]). *Let u and v be reduced words. Then u is left-absorbed by v if and only if $[u \cdot v] = v$.*

Proof. Note that $[u \cdot v] = v$ if and only if $u_1 \cdot w \cdot v_1 \approx v' \cdot w \cdot v_1$. Lemma 5.2 implies $u_1 \approx v'$. So u_1 is properly right-absorbed by itself, which can only happen if $u_1 = 1$. Thus $u \approx u' \cdot w$ is left-absorbed by v. □

Corollary 5.11 (cf. [2, Corollary 5.22]). *A reduced word is commuting if and only if it left-absorbs itself.*

Proof. Suppose first that we have reduced words u and v such that u right-absorbs v and v left-absorbs u. The proof of Corollary 5.10 yields $u_1 = v' = 1 = v_1 = u'$, so $u = v = w$ is commuting. The statement now follows easily, since Corollary 5.10 implies that if u left-absorbs itself, then it also right-absorbs itself. □

Corollary 5.12. *If the reduced word u is left-absorbed by the reduced word v, we can write* (*up to permutation*)

$$u = u' \cdot w, \qquad w \cdot v_1 = v,$$

where

(1) *w is a commuting word,*
(2) *u' is properly left-absorbed by v_1,*
(3) *u' and w commute.* □

Corollary 5.13. *If u, v and x are reduced words such that $[u \cdot v] = u \cdot x$, then x is a final subword of v. In particular, x commutes with any word commuting with v.*

Proof. Decompose

$$u = u_1 \cdot u' \cdot w, \qquad w \cdot v' \cdot v_1 = v,$$

as in Proposition 5.9. Then $u_1 \cdot w \cdot v_1 = [u \cdot v] = u_1 \cdot w \cdot u' \cdot x$, so $v_1 \approx u' \cdot x$, by Lemma 5.2. Thus $u' = 1$ and x is a final subword of v, as desired. □

Lemma 5.14. *If the reduced word u is left-absorbed by $v_1 \cdot v_2$, then $u \approx u_1 \cdot u_2$, where each u_i is left-absorbed by v_i, and u_2 commutes with v_1 (and therefore with u_1).*

Proof. If u is empty, there is nothing to show. Otherwise, write $u = u' \cdot s$. By induction on the length of u, obtain a decomposition $u' \approx u'_1 \cdot u'_2$ such that u'_i is absorbed by v_i and u'_2 commutes with v_1. We distinguish two cases: If s is absorbed by v_1, then u'_2 commutes with s, so decompose $u \approx (u'_1 \cdot s) \cdot u'_2$. Otherwise, s commutes with v_1 and is left-absorbed by v_2. Set $u \approx u_1 \cdot (u_2 \cdot s)$. □

Corollary 5.15. *Given a reduced word $u \approx u_1 \cdot \tilde{u}$, where $\tilde{u}$ denotes the final segment of u, we have $[u \cdot u^{-1}] = u_1 \cdot \tilde{u} \cdot u_1{}^{-1}$.*

Proof. It suffices to show that the word $u_1 \cdot \tilde{u} \cdot u_1{}^{-1}$ is reduced. Otherwise, there is an end s of u_1 which commutes with $\tilde{u}$. In that case, the letter s is an end of u and hence it occurs in $\tilde{u}$, which is a contradiction. □

For the proof of the next lemma, we will require the following notation: Recall (Remark 4.4) that $\Gamma \setminus \partial s$ is the set of those colours commuting with s. For a letter t, denote by $\mathrm{C}_t(s)$ the commuting word with support $t \setminus \partial s$ (Remark 3.2). For a word $v = t_1 \cdots t_n$, set

$$\mathrm{C}_v(s) = \mathrm{C}_{t_1}(s) \cdots \mathrm{C}_{t_n}(s).$$

Lemma 5.16 (Division Lemma). *Given reduced words $u \preceq v$, there exists a reduced word w, unique up to permutation, such that for every reduced word x,*

$$[x \cdot u] \preceq v \;\Leftrightarrow\; x \preceq w.$$

We denote w by v/u. Since $v/u \preceq v/u$, Corollary 5.8 implies that $v/u \preceq v$. Furthermore, since $1 \preceq v/u$, the condition $u \preceq v$ is necessary for the existence of v/u.

Proof of Lemma 5.16. Note that we may assume that u consists of a single letter: given $u = u_1 \cdot u_2$, suppose that the statement holds for both u_1 and u_2. Since $u \preceq v$, we have $u_2 \preceq v$, so v/u_2 exists. Now $u_1 \cdot u_2 \preceq v$ implies $u_1 \preceq v/u_2$. Thus, set $v/u = (v/u_2)/u_1$, which exists. Note that

$$[x \cdot u] = [[x \cdot u_1] \cdot u_2] \preceq v \;\Leftrightarrow\; [x \cdot u_1] \preceq v/u_2 \;\Leftrightarrow\; x \preceq v/u,$$

as desired. Therefore, assume $u = s \preceq v$.

We can write $v = v_1 \cdot t \cdot v_2$ where $s \subset t$ and no letter in v_2 contains s. Set $w = [v_1 \cdot t \cdot \mathrm{C}_{v_2}(s)]$. No letter in $v_1 \cdot t$ is left-absorbed during the non-splitting reduction $v_1 \cdot t \cdot \mathrm{C}_{v_2}(s) \to w$, for v is reduced.

Clearly $[w \cdot s] = w \preceq v$. Given $x \preceq w$ reduced, Corollary 5.7 implies that $[x \cdot s] \preceq [w \cdot s] \preceq v$, which gives one implication. For the other, assume that $[x \cdot s] \preceq v$. We distinguish two cases : if s is absorbed by x, then $x = [x \cdot s] \preceq v = v_1 \cdot t \cdot v_2$, so we may decompose $x \approx x_1 \cdot x_2$, where $x_1 \preceq v_1 \cdot t$ and $x_2 \preceq v_2$. Since x_2 does not right-absorb s, Lemma 5.14 implies that x_2 commutes with s, which is right-absorbed by x_1. This implies $x_2 \preceq \mathrm{C}_{v_2}(s)$, and thus $x \preceq w$.

If s is not absorbed by x, then Proposition 5.9 yields a decomposition $x \approx x_1 \cdot x'$, where x' is properly absorbed by s and $[x \cdot s] = x_1 \cdot s$. The word $x_1 \cdot s$ right-absorbs s, so $x_1 \cdot s \preceq w$ by the previous discussion. Since $x' \prec s$, we conclude that $x \approx x_1 \cdot x' \preceq w$, as desired. □

Definition 5.17. According to Lemma 5.16, we denote by $\mathcal{S}_{\mathrm{R}}(u)$ the largest reduced word, unique up to permutation, such that for every reduced word x,

$$[u \cdot x] \preceq u \;\Leftrightarrow\; x \preceq \mathcal{S}_{\mathrm{R}}(u).$$

Corollary 5.18. *Given a reduced word u, the word $\mathcal{S}_{\mathrm{R}}(u) \preceq u$ is commutative. A reduced word x is right-absorbed by u if and only if $x \preceq \mathcal{S}_{\mathrm{R}}(u)$. In particular, if $\tilde{u}$ denotes the final segment of u, then $\tilde{u} \preceq \mathcal{S}_{\mathrm{R}}(u)$.*

Proof. Note that x is right-absorbed by u if and only if $[u \cdot x] = u$. Since $u \preceq [u \cdot x]$ by Corollary 5.8, this is equivalent to $[u \cdot x] \preceq u$, that is, $x \preceq \mathcal{S}_{\mathrm{R}}(u)$.

In particular, the word $\mathcal{S}_{\mathrm{R}}(u)$ is right-absorbed by u. Thus $[u \cdot [\mathcal{S}_{\mathrm{R}}(u) \cdot \mathcal{S}_{\mathrm{R}}(u)]] = [[u \cdot \mathcal{S}_{\mathrm{R}}(u)] \cdot \mathcal{S}_{\mathrm{R}}(u)] = [u \cdot \mathcal{S}_{\mathrm{R}}(u)] = u$, so $\mathcal{S}_{\mathrm{R}}(u) \preceq [\mathcal{S}_{\mathrm{R}}(u) \cdot \mathcal{S}_{\mathrm{R}}(u)] \preceq \mathcal{S}_{\mathrm{R}}(u)$, which implies that $\mathcal{S}_{\mathrm{R}}(u)$ is commutative, by Corollary 5.11. □

6. Forking

We work inside a big sufficiently saturated model M of the theory PS_Γ, as a universal domain. Recall, by Corollary 4.14, that PS_Γ is ω-stable. Given a finite tuple a and subsets $C \subset B$ of M, the extension $\mathrm{tp}(a/C) \subset \mathrm{tp}(a/B)$ is *non-forking* if $\mathrm{RM}(\mathrm{tp}(a/B)) = \mathrm{RM}(\mathrm{tp}(a/C))$. More generally, given subsets A, B and C of M, the set A is *independent from B over C*, denoted by

$$A \mathop{\smile\!\!\!\!\!\!|}_{C} B,$$

if, for every finite tuple a in A, the extension $\mathrm{tp}(a/C) \subset \mathrm{tp}(a/B \cup C)$ is non-forking. This gives rise to a well-behaved notion of independence, which has, among many other, the following remarkable properties (cf. [16, Corollary 8.5.4 and Theorem 8.5.5]):

Symmetry: If $A \mathop{\smile\!\!\!\!\!\mid}_C B$, then $B \mathop{\smile\!\!\!\!\!\mid}_C A$.
Extension: Given a tuple a and subsets $C \subset B$ of M, there is a non-forking extension of $\mathrm{tp}(a/C)$ to $B \cup C$, that is, there is some realisation of $\mathrm{tp}(a/C)$ which is independent from B over C.
Stationarity: Every type p over an elementary substructure N of M is *stationary*, that is, given a subset B of M, there is a unique non-forking extension of p to $N \cup B$.
Invariant Extension: If p is a type over a sufficiently saturated elementary substructure N of M which is *invariant* over a small subset $C \subset N$, that is, every automorphism of N fixing C fixes p (as a collection of formulae), then p is the unique non-forking extension of $p \restriction C$ to N.

Indeed, the last property follows from the fact that a global type, which is invariant over C, does not fork over C [16, Exercise 7.1.4], plus the fact that all non-forking extensions of $p \restriction C$ to N are conjugate under automorphisms of N [16, Theorem 8.5.6].

In a similar fashion to [2, Section 7], we will describe non-forking over nice sets and canonical bases in PS_Γ. We will also show that this theory, which has weak elimination of imaginaries, has trivial forking and furthermore is totally trivial, as defined in [5].

Recall the terminology introduced in Definition 3.20: A word u connects the flag F to the flag G if it is the word of a reduced path from F to G. This word is unique, up to permutation, and denoted by $\mathrm{w}(F, G)$.

The following result describes the type of a flag over a nice set and will help us to determine the nature of non-forking.

Proposition 6.1. *Given a flag F and a reduced path with word u which connects F to a flag G lying in a nice set D, the following are equivalent:*

(a) *For any flag G' in D,*
$$\mathrm{w}(F, G') = [u \cdot \mathrm{w}(G, G')],$$
that is, the word connecting F to G' is equivalent to $[u \cdot v]$, the non-splitting reduct of $u \cdot v$, where v is the reduced word connecting G to G'.
(b) *The word u is the $\preceq$-smallest word connecting F to some flag in D.*
(c) *The word u is $\preceq$-minimal among words connecting F to some flag in D.*

This generalisation of [2, Proposition 7.2] has essentially the same proof. Note that a word u satisfying (b) is unique, up to permutation, for $\prec$ is irreflexive.

Proof. (a)⇒(b) follows from Corollary 5.8. The implication (b)⇒(c) is trivial.

For (c)⇒(a), let v be the word of a reduced path from G to some flag G' in D. By niceness, we may assume that the path is fully contained in D. Choose a decomposition $u = u_1 \cdot u' \cdot w$ and $w \cdot v' \cdot v_1 = v$, as in Proposition 5.9, with corresponding paths

$$F \xrightarrow{u_1 \cdot u'} F_1 \xrightarrow{w} G \xrightarrow{w} G_1 \xrightarrow{v' \cdot v_1} G',$$

where G_1 is some flag in D. The word b connecting F_1 to G_1 is a reduct of $w \cdot w$, so $b \preceq w$. The word c connecting F to G_1 is hence a reduct of $u_1 \cdot u' \cdot b$, so

$$c \preceq u_1 \cdot u' \cdot b \preceq u_1 \cdot u' \cdot w \preceq u.$$

Minimality of u implies $c \approx u$, so $b \approx w$. Hence, in the reduction of the path $F \to G'$, no splitting occurred and the resulting word is $u_1 \cdot w \cdot v_1 = [u \cdot v]$, by Proposition 5.9. □

Definition 6.2. A *base-point* of the flag F over the nice set D is a flag G in D such that any of the conditions of Proposition 6.1 hold.

Recall Definition 5.17 of $\mathcal{S}_{\mathrm{R}}(u)$. By Corollaries 5.10 and 5.18, we easily deduce the following:

Corollary 6.3. *Let G be a base-point of F over the nice set D and u the word which connects F to G. Let v be a reduced word connecting G to some flag G_1 in D. The flag G_1 is a base-point of F over D if and only if v is right-absorbed by u if and only if $u = [u \cdot v]$ if and only if $v \preceq \mathcal{S}_{\mathrm{R}}(u)$.*

Lemma 6.4 (cf. [2, Lemma 7.4]). *Let G be a base-point of F over the nice set D and denote by P a reduced path $F = F_0, \ldots, F_n = G$ with word u. Then $D \cup P$ is nice. It is uniquely determined by G and u in the following strong sense: If $P' = F'_0, \ldots, F'_n$ is a second path with word u from F to G, then there is a (unique) isomorphism $D \cup P \to D \cup P'$ which is the identity on D and maps each F_i onto F'_i.*

We will express the last property by saying that the “extension $F_0 \cdots F_n/D$” is uniquely determined, up to isomorphism.

Proof of Lemma 6.4. If u is trivial, there is nothing to show. Otherwise, write $u = u' \cdot s$. Minimality of u shows that no splitting of s can connect F_{n-1} to a flag in D. Thus Proposition 4.7 implies that the set $D' = D \cup F_{n-1}$ is nice and a simple extension of D of type (s, G).

We will now show that F_{n-1} is a base-point of F over D', by Proposition 6.1(c). Otherwise, there is a reduced word $v \prec u'$ which connects F to a flag H' in D'. By minimality of u, the flag H' cannot lie in D. By Remark 4.4, there exists some flag H in D such that $H' \xrightarrow{s} H$. This gives a reduced path from F to H whose word is some reduction of $v \cdot s$, contradicting the minimality of u.

By induction, if P_0 denotes the subpath $F = F_0, \ldots, F_{n-1}$, then the set $D' \cup P_0$ is nice and the extension $F_0 \cdots F_{n-1}/D'$ is uniquely determined, up to isomorphism, by G' and u'. Therefore, the extension $F_0 \cdots F_n/D$ is uniquely determined, up to isomorphism, by G and u. □

Corollary 6.5. *Given a reduced word u and a flag G in a nice set D, there is a flag F such that u connects F to G, which is a base-point of F over D. The type of F over D is uniquely determined by G and u.*

Recall that the word $\mathrm{w}(F, G)$ is determined only up to a permutation. So $\mathrm{tp}(F/D)$ depends only on the equivalence class of u.

Proof of Corollary 6.5. Observe that if such a flag F exists and P denotes the reduced path from F to G with word u, Lemma 6.4 and Corollary 4.13 imply that the type of P over D is uniquely determined by G and u.

Thus, we need only show the existence of such a flag F, by induction on the length of u. If $u = 1$, there is nothing to do. Otherwise, write $u = s \cdot u'$ and choose a flag F' connecting to G by u' such that G is a base-point of F' over D. Let P' denote the reduced path $F' \xrightarrow{u'} G$. By Lemma 6.4, the set $D' = D \cup P'$ is nice. Corollary 4.11 yields a simple extension $D' \cup F$ of type (s, F'). Proposition 4.7 implies that F' is a base-point of F over D'. We need only show that G is a base-point of F. Hence, let G' be an arbitrary flag in $D \subset D'$. We have

$$\mathrm{w}(F, G') = [s \cdot \mathrm{w}(F', G')] = [s \cdot [u' \cdot \mathrm{w}(G, G')]] = [u \cdot \mathrm{w}(G, G')],$$

as desired. □

If we denote the type of F over G, resp. over D, by

$$\mathrm{p}_u(G), \quad \text{resp.} \quad \mathrm{p}_u(G)|D,$$

we deduce the following.

Proposition 6.6. *Let G be a flag in the nice set D and u a reduced word. Then $\mathrm{p}_u(G)|D$ is the unique non-forking extension of $\mathrm{p}_u(G)$ to D.*

Proof. By the Extension principle, we may replace D by a sufficiently saturated elementary substructure containing it, which is again nice in M. Since $\mathrm{p}_u(G)|D$ is invariant over G, the Invariant Extension principle yields the desired result. □

Since the type $\mathrm{p}_u(G)$ admits a non-forking extension to D, which must coincide with $\mathrm{p}_u(G)|D$ by the previous result, we obtain the following immediate observation.

Corollary 6.7 (cf. [2, Lemmata 7.4 and 7.6]). *Given a flag F and a nice set D, the flag G in D is a base-point of F over D if and only if $F \mathop{\smile\!\!\!\!\!\!\mid}_G D$.*

Recall that the *canonical base* $\mathrm{Cb}(p)$ of a stationary type p is some set fixed pointwise by exactly those automorphisms of M fixing the global non-forking extension $\mathbf{p}$ of p to M. If $\mathrm{Cb}(p)$ exists, then it is unique, up to interdefinability, and $\mathbf{p}$ is the unique non-forking extension to M of its restriction to $\mathrm{Cb}(p)$. Furthermore, if p is a stationary type over B and $\mathrm{Cb}(p)$ exists, then p does not fork over $A \subset B$ if and only if $\mathrm{Cb}(p)$ is algebraic over A.

Canonical bases exist as *imaginary* elements in the expansion T^{eq} of an ω-stable theory T [16, Chapter 8.4].

As in Remark 3.2, given a reduced word u, we do not distinguish between the word $\mathcal{S}_{\mathrm{R}}(u)$ and its support. For a flag G, the class of G modulo $\mathcal{S}_{\mathrm{R}}(u)$ can be identified with a subset of the real sort, namely with the set of vertices of G whose colours do not belong to $\mathcal{S}_{\mathrm{R}}(u)$. To simplify the notation, we will denote this class by $G/\mathcal{S}_{\mathrm{R}}(u)$.

Corollary 6.8. *The class $G/\mathcal{S}_{\mathrm{R}}(u)$ is a canonical base of* $\mathrm{p}_u(G)$.

Proof. Two types $\mathrm{p}_u(G)$ and $\mathrm{p}_u(G_1)$ have a common global non-forking extension if and only if $\mathrm{p}_u(G)|D = \mathrm{p}_u(G_1)|D$ for some nice set D which contains G and G_1. This is equivalent, by Corollary 6.3, to $\mathrm{w}(G, G_1) \preceq \mathcal{S}_{\mathrm{R}}(u)$, which is equivalent to $G_1 \sim_{\mathcal{S}_{\mathrm{R}}(u)} G$. □

Definition 6.9. Given reduced words u and v, we say that u is a *proper left-divisor* of v if $u \not\approx v$ and there is a reduced word w such that $[u \cdot w] = v$.

If u is a proper left-divisor of v, it follows that $u \prec v$. In particular, being a proper left-divisor is a well-founded relation. Let $\mathrm{R}_{\mathrm{div}}$ be its foundation rank, and likewise let $\mathrm{R}_{\prec}$ denote the foundation rank of reduced words with respect to $\prec$.

The foundation rank of types associated to the forking relation is called *Lascar rank*, denoted by U. This means that a type has Lascar rank at least $\alpha + 1$ if and only if it has a forking extension of Lascar rank at least α.

The following result can be proved exactly as [2, Lemmata 7.10 and 7.11].

Lemma 6.10. *For every flag G and every reduced word u,*

$$\mathrm{U}(\mathrm{p}_u(G)) = \mathrm{R}_{\mathrm{div}}(u) \le \mathrm{RM}(\mathrm{p}_u(G)) \le \mathrm{R}_{\prec}(u).$$

In general $\mathrm{R}_{\mathrm{div}}(u)$, $\mathrm{RM}(\mathrm{p}_u(G))$ and $\mathrm{R}_{\prec}(u)$ need not agree [2, Remark 7.14]. They coincide however in the following special case, the proof of which is a straightforward modification of the proof of [2, Lemma 7.12], together with Lemma 4.10.

Lemma 6.11 (cf. [2, Corollary 7.13]). *For every reduced word $u = s_1 \cdots s_n$ such that $|s_i| \ge |s_{i+1}|$ for $i = 1, \ldots, n-1$,*

$$\mathrm{R}_{\mathrm{div}}(u) = \mathrm{R}_{\prec}(u) = \omega^{|s_1|-1} + \cdots + \omega^{|s_n|-1}.$$

Remark 6.12. For an arbitrary reduced word u, the Morley rank of the type $\mathrm{p}_u(G)$ can be easily computed thanks to the following observation: The rank of $\mathrm{p}_u(G)$ is strictly larger than α if and only if either

(1) the word u has a proper left-divisor v such that $\mathrm{p}_v(G)$ has rank at least α, or
(2) the type $\mathrm{p}_u(G)$ is an accumulation point of a family of types $\mathrm{p}_v(G)$, each of rank at least α.

Corollary 6.13. *The theory PS_Γ is ω-stable of Morley rank ω^{K-1}, where K is the cardinality of a connected component of Γ of largest size.*

Proof. Decompose $\Gamma = \bigcup_{i=1}^n \Gamma_i$ into its connected components. As in Corollary 8.4, it is easy to see that each restriction $M_i = \mathcal{A}_{\Gamma_i}(M)$ is a model of PS_{Γ_i}. The structure M can be considered as the disjoint union of the structures M_i's, so the Morley rank of M is the maximum of the Morley ranks of the structures M_i. We may therefore assume that Γ is connected.

Given any vertex a, choose a flag F containing a as well as a flag G independent from F over $\emptyset$. If $\mathrm{p}_u(G)$ is the type of F over G, then the word u must be equal to Γ, since the canonical base $G/\mathcal{S}_{\mathrm{R}}(u)$ is algebraic over the empty set. By Lemma 6.11,

$$\mathrm{U}(F) = \mathrm{RM}(F) = \mathrm{R}_{\prec}(\Gamma) = \omega^{K-1}.$$

By Lascar inequalities [16, Exercise 8.6.5], we have $\mathrm{U}(F/a) + \mathrm{U}(a) \leq \mathrm{U}(F) = \omega^{K-1}$, so $\mathrm{U}(a) = \omega^{K-1}$, since $\mathrm{U}(a) > 0$. Since

$$\omega^{K-1} = \mathrm{U}(a) \leq \mathrm{RM}(a) \leq \mathrm{RM}(F) = \omega^{K-1},$$

we have equality, as desired. □

Two types p and q, possibly over different sets of parameters, are *non-orthogonal* if there is a common extension C of both sets of parameters, and two realisations a and b of the corresponding non-forking extensions of p and q to C such that a forks with b over C. As in [2, Theorem 7.15], we deduce the following.

Remark 6.14. Every type over a nice set D is non-orthogonal to some $\mathrm{p}_s(G)|D$, where G lies in D.

Given a reduced flag path $P : F \xrightarrow{u} G$, we will conclude this section by describing the flags one can obtain from the collection of vertices of the flags occurring in P, as well as describing how the flags in P can vary (or *wobble*), whilst the endpoints are fixed.

Lemma 6.15 (cf. [2, Lemma 6.18]). *Let A be the set of vertices of the flags occurring in a reduced flag path P. Then any flag in A occurs in some permutation of P (cf. Lemma 3.7).*

Proof. Write $P : F \xrightarrow{u} G$ and note first that G is the base-point of F over G. If $u = 1$, then $F = G$ is the only flag in A, so there is nothing to prove. Otherwise, write $u = s \cdot v$ and decompose the path P as $F \xrightarrow{s} H \xrightarrow{v} G$. If we denote by B the set of vertices of the flags occurring in the reduced path $H \xrightarrow{v} G$, Lemma 6.4 shows that B is nice. By induction and Proposition 4.7, the nice set $B \cup F$ is a simple extension of B of type (s, H).

Given any flag K in A, we distinguish two cases: If K lies in B, by induction K occurs in some permutation of $H \xrightarrow{v} G$, which induces a permutation of P. Otherwise, by Remark 4.4, there exist a reduced word w commuting with s and a flag K_1 in B such that $F \xrightarrow{s} H \xrightarrow{w} K_1$ and $F \xrightarrow{w} K \xrightarrow{s} K_1$. By Corollary 3.6, we may assume that the second path is a permutation of the first. By induction, K_1 belongs to a permutation $H \xrightarrow{w} K_1 \xrightarrow{v_2} G$ of $H \xrightarrow{v} G$. So $F \xrightarrow{w} K \xrightarrow{s} K_1 \xrightarrow{v_2} G$ is a permutation of P. □

The following generalises Remark 4.4 and Lemma 6.15.

Lemma 6.16. *Let G be a base-point of F over the nice set D, and P a reduced path connecting F to G. For every flag K' in the nice set $D \cup P$, there is a flag K occurring in some permutation of P and a flag G' in D such that $w = \mathrm{w}(K, G)$ commutes with $v = \mathrm{w}(G, G')$ and $K \xrightarrow{v} K' \xrightarrow{w} G'$.*

Proof. If P is trivial, set $K = G$ and $G' = K'$. Otherwise, decompose P into $F \xrightarrow{s} F' \xrightarrow{u'} G$ and set $P' : F' \xrightarrow{u'} G$. Note that G is also a base-point of F' over D. If K' is contained in $D \cup P'$, find, by induction on the length of P, a flag K occurring in some permutation of P' and a flag G' in D, as desired. Otherwise, Remark 4.4 implies the existence of a flag K'_1 in $D \cup P'$ such that $w' = \mathrm{w}(F', K'_1)$ commutes with s and $F \xrightarrow{w'} K' \xrightarrow{s} K'_1$. By induction, we find a flag K_1 occurring in a permutation of P' and a flag G' in D such that $w_1 = \mathrm{w}(G, G')$ commutes with $v_1 = \mathrm{w}(K_1, G)$ and $K_1 \xrightarrow{w_1} K'_1 \xrightarrow{v_1} G'$.

Let $P_1 : K_1 \xrightarrow{v_1} G$ be the part of a permutation of P' which connects K_1 to G. Lemma 6.4 implies that the set $D_1 = D \cup P_1$ is nice. Furthermore, the flag K_1 is a base-point of F' over D_1. Set $u_1 = \mathrm{w}(F', K_1)$. Since $F' \xrightarrow{u_1} K_1 \xrightarrow{w_1} K'_1$, we conclude that $w' = [u_1 \cdot w_1]$. In particular, the letter s must commute with both u_1 and w_1. Since $F \xrightarrow{s} F' \xrightarrow{u_1} K_1$, there exists a unique flag K such that $F \xrightarrow{u_1} K \xrightarrow{s} K_1$. Clearly, K occurs in a permutation of P with word $u_1 \cdot s \cdot v_1$ and $\mathrm{w}(K, G) = s \cdot v_1 = \mathrm{w}(K', G')$. We need only show that $\mathrm{w}(K, K')$ commutes with $s \cdot v_1$. Since $K \xrightarrow{u_1^{-1}} F \xrightarrow{w'} K'$, the word $\mathrm{w}(K, K')$ is some reduction of $u_1^{-1} \cdot w'$ (possibly with splitting), so $\mathrm{w}(K, K')$ commutes with s, since both u_1 and w do. The flag path $K \xrightarrow{s} K_1 \xrightarrow{w_1} K'_1 \xrightarrow{s} K'$ must reduce to one with word $\mathrm{w}(K, K')$, which commutes with s, so the word

$$s \cdot w_1 \cdot s \approx w_1 \cdot s \cdot s$$

must reduce to $w_1 = \mathrm{w}(K, K')$, which commutes with $s \cdot v_1$, as desired. □

Corollary 6.17. *Let A be the nice set consisting of the vertices of the flags occurring in a reduced flag path $P : F \xrightarrow{u} G$. Every flag K in A is uniquely determined by $\mathrm{w}(K, G)$. Thus, the only automorphism of A fixing one of the endpoints of P is the identity.*

Proof. Given flags K and K' in A with $\mathrm{w}(K', G) = \mathrm{w}(K, G) = w$, Lemma 6.15 shows that there are permutations $Q : F \xrightarrow{v} K \xrightarrow{w} G$ and $Q' : F \xrightarrow{v'} K' \xrightarrow{w'} G$ of P. Since $w \approx w'$, Lemma 5.2 implies that $v' \approx v$. We may therefore permute Q' in order to decompose it as $Q_1 : F \xrightarrow{v} K' \xrightarrow{w} G$. The correspondence between permutations of a flag path and permutations of the associated word in Lemma 3.7 implies that $Q = Q_1$, so $K' = K$. □

Definition 6.18. The *wobbling* of a reduced product $u \cdot v$ is $\mathrm{Wob}(u, v) = \mathcal{S}_{\mathrm{R}}(u) \cap \mathcal{S}_{\mathrm{R}}(v^{-1})$, that is, the collection of those γ in Γ which are both right-absorbed by u and left-absorbed by v.

Note that $\mathrm{Wob}(u, v)$ cannot be equal to $|u|$ or to $|v|$, for $u \cdot v$ is reduced.

Lemma 6.19 (Wobbling Lemma, cf. [2, Lemma 6.19]). *Given two reduced paths between the flags F and G with the same word $u = s_1 \cdots s_i \cdots s_n$,*

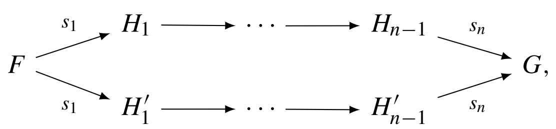

the flags H_i and H_i' are $\mathrm{Wob}(s_1 \cdots s_i, s_{i+1} \cdots s_n)$*-equivalent for all i in* $\{1, \ldots, n-1\}$.

In particular, the tuple $H_i/\mathrm{Wob}(s_1 \cdots s_i, s_{i+1} \cdots s_n)$ *enumerating the vertices of H_i with colours in* $\Gamma \setminus \mathrm{Wob}(s_1 \cdots s_i, s_{i+1} \cdots s_n)$ *lies in* $\mathrm{dcl}(F, G)$.

Proof. Given two different flag paths as in the above picture, we prove the statement by induction on $i < n$. For $i = 1$, let w_1 be the reduced word connecting H_1 to H_1'. Since $H_1 \xrightarrow{s_1} F \xrightarrow{s_1} H_1'$, it follows that $w_1 \preceq s_1$ and w_1 is right-absorbed by s_1. If $w_1 = s_1$, this contradicts Proposition 3.19, since $H_1 \xrightarrow{s_2 \cdots s_n} G$. Therefore, the word w_1 is a proper splitting of s_1. Furthermore, since $w_1 \cdot s_2 \cdots s_n \xrightarrow{*} s_2 \cdots s_n$, no letter from $s_2 \cdots s_n$ can be absorbed during the reduction, for the word u is reduced. Corollary 5.10 implies that w_1 is completely absorbed by $s_2 \cdots s_n$, so $|w_1| \subset \mathrm{Wob}(s_1, s_2 \cdots s_n)$, as desired.

Let now $H_i \xrightarrow{w_i} H_i'$, resp. $H_{i+1} \xrightarrow{w_{i+1}} H_{i+1}'$. By induction, the word w_i has support in $\mathrm{Wob}(s_1 \cdots s_i, s_{i+1} \cdots s_n)$. Lemma 5.14 implies that $w_i \approx w_i^1 \cdot w_i^2$, where w_i^1 is left-absorbed by s_{i+1}, and w_i^2 commutes with s_{i+1} and is left-absorbed by $s_{i+2} \cdots s_n$. Thus $w_i^1 \prec s_{i+1}$. Since

$$s_{i+1} \cdot w_i \cdot s_{i+1} \to w_i^2 \cdot s_{i+1} \cdot s_{i+1}$$

reduces to w_{i+1}, we conclude that $w_{i+1} \approx w_i^2 \cdot x$, where $x \preceq s_{i+1}$. Note that $x \prec s_{i+1}$, since $H_{i+1} \xrightarrow{s_{i+2} \cdots s_n} G$, so x is a splitting of s_{i+1} which must be absorbed by $s_{i+2} \cdots s_n$, and so is w_{i+1}. The word w_i^2 is right-absorbed by $s_1 \cdots s_i$, by induction, and commutes with s_{i+1}. Since x is a proper splitting of s_{i+1}, we have $s_1 \cdots s_{i+1} \cdot w_i^2 \cdot x \to s_1 \cdots s_{i+1}$, so w_{i+1} is right-absorbed by $s_1 \cdots s_{i+1}$, as desired. □

Lemma 6.20 (Base-Point Lemma, cf. [2, Lemma 7.18]). *Given a flag H with base-point G in the nice set A, let $H \xrightarrow{v} G$ be a reduced flag path connecting H to G. If H/W lies in* $\mathrm{acl}(A)$ *for some subset $W \subset \Gamma$, then $|v| \subset W$.*

Proof. Choose a flag $H' \models \mathrm{tp}(H/\mathrm{acl}(A))$ with $H \mathop{\smile\!\!\!\!\!\!\mid}_A H'$. Note that G is a base-point in A for H' as well. Thus Proposition 6.6 and transitivity of non-forking imply that $H \mathop{\smile\!\!\!\!\!\!\mid}_G H'$. Furthermore, the flags H and H' are W-equivalent, so the reduced word w connecting H to H' has support in W. By Proposition 6.1, the word w is the non-splitting reduct of $v \cdot v^{-1}$ and equals $v_1 \cdot \tilde{v} \cdot v_1^{-1}$, where $\tilde{v}$ is the final segment of $v \approx v_1 \cdot \tilde{v}$, by Corollary 5.15. Thus, the set $|v|$ is contained in $|w| \subset W$, as desired. □

Corollary 6.21. *Nice sets are algebraically closed.*

Proof. Let b be algebraic over the nice set A. Suppose that b has colour γ and choose some flag H containing it. Pick some base-point G for H over A, and let $P : H \xrightarrow{u} G$ be a reduced flag path connecting H to G. By assumption, the set $H/(\Gamma \setminus \{\gamma\}) = \{b\}$ lies in $\mathrm{acl}(A)$, so $|u| \subset \Gamma \setminus \{\gamma\}$ by Lemma 6.20, that is, the element b equals the γ-vertex of G, which lies in A. □

7. Equationality

In this section, we will show that the theory PS_Γ is equational.

Definition 7.1. A parameter-free formula $\varphi(x, y)$, where the tuple x has length n, is an *n-equation* if the family of finite intersections of instances $\varphi(x, a)$ (where a belongs to a sufficiently saturated model N) has the descending chain condition (DCC).

A complete theory T is *n-equational* if every definable set in N^n is a Boolean combination of instances of n-equations. A theory is *equational* if it is n-equational for every n in $\mathbb{N}$.

Stability, a wider class class containing ω-stable theories [16, Section 8.2], is preserved under naming parameters and bi-interpretability. The same holds for equationality [11]. However, it is unknown whether equationality follows from 1-equationality, which itself implies stability for formulae $\varphi(x, y)$ where x is a single variable, and thus stability [14]. The rest of this section is devoted to showing that the theory PS_Γ is equational. As in the previous section, let M denote a sufficiently saturated model of PS_Γ, inside of which we work.

Definition 7.2. Given a word $u = s_1 \cdots s_n$ and a flag G in M, let $\mathrm{P}_u(X, G)$ be the formula stating the existence of a sequence $X = F_0, \ldots, F_n = G$ of flags with $F_{i-1} \sim_{s_i} F_i$ for $1 \le i \le n$.

It follows from Lemmata 3.4, 3.7 and 3.9 that $M \models \mathrm{P}_u(F, G)$ if and only if there exists some reduction $u' \preceq u$ with $F \xrightarrow{u'} G$. Thus, if $w \preceq u$, then the sentence

$$\forall X\ \forall Y\ (\mathrm{P}_w(X, Y) \to \mathrm{P}_u(X, Y))$$

holds in all Γ-spaces. When the word w is reduced, the converse holds in models of PS_Γ, as shown below.

Remark 7.3. Let w be a reduced word and $w_1, \ldots, w_n$ be arbitrary words. Then

$$\mathrm{PS}_\Gamma \models \forall X\ \forall Y \left(\mathrm{P}_w(X, Y) \to \bigvee_{i=1}^{n} \mathrm{P}_{w_i}(X, Y)\right)$$

if and only if $w \preceq w_i$ for some i.

Proof. We need only prove left-to-right. By Corollary 6.5, there are flags F and G in our saturated model M such that $F \xrightarrow{w} G$. Thus $\mathrm{P}_{w_i}(F, G)$ for some i, so $F \xrightarrow{w'} G$ for some reduced $w' \preceq w_i$. Proposition 3.19 implies that $w \approx w'$, so $w \preceq w_i$. □

Lemma 7.4. *Given two arbitrary words u and v, there is a finite collection of reduced words $w_1, \ldots, w_n$ such that, for any reduced word w,*

$$w \preceq u, v \iff \bigvee_{i=1}^{n} w \preceq w_i.$$

In particular, $w_i \preceq u, v$ for all i.

Proof. In M, the formula $\mathrm{P}_u(F, G) \wedge \mathrm{P}_v(F, G)$ implies that $\mathrm{P}_w(F, G)$ for some reduced word $w \preceq u, v$. By compactness, there is a finite set $w_1, \ldots, w_n$ of reduced words satisfying $w_i \preceq u, v$ for $i = 1, \ldots, n$ and

$$\mathrm{PS}_\Gamma \models \forall X\, \forall Y \left(\mathrm{P}_u(X, Y) \wedge \mathrm{P}_v(X, Y) \to \bigvee_{i=1}^{n} \mathrm{P}_{w_i}(X, Y)\right).$$

Thus, if $w \preceq u, v$, the formula $\mathrm{P}_w(X, Y)$ implies $\bigvee_{i=1}^{n} \mathrm{P}_{w_i}(X, Y)$. Hence $w \preceq w_i$ for some i, by Remark 7.3. □

Lemma 7.5. *Given flags G_1 and G_2 in M, and reduced words u_1 and u_2 such that neither $\mathrm{P}_{u_1}(X, G_1)$ nor $\mathrm{P}_{u_2}(X, G_2)$ implies the other, if P denotes a reduced path from G_1 to G_2, then the conjunction $\mathrm{P}_{u_1}(X, G_1) \wedge \mathrm{P}_{u_2}(X, G_2)$ is equivalent to*

$$\bigvee_{i=1}^{n} \mathrm{P}_{w_i}(X, H_i)$$

for some flags $H_1, \ldots, H_n$ occurring in some permutation of P and for some reduced words $w_1, \ldots, w_n \prec u_1, u_2$.

If $\mathrm{P}_{u_1}(X, G_1)$ and $\mathrm{P}_{u_2}(X, G_2)$ are disjoint, set $n = 0$. If $G_1 = G_2$, this is the content of Lemma 7.4.

Proof of Lemma 7.5. Choose any realisation $F \models \mathrm{P}_{u_1}(X, G_1) \wedge \mathrm{P}_{u_2}(X, G_2)$ and a base-point H of F over the nice set determined by the reduced path P. The flag H occurs in some permutation of P, by Lemma 6.15.

Set $x = \mathrm{w}(F, H)$ and $v_i = \mathrm{w}(H, G_i)$ for $i = 1, 2$. Observe that $\mathrm{w}(F, G_i) \preceq u_i$ for $i = 1, 2$. Proposition 6.1 implies that $[x \cdot v_i] = \mathrm{w}(F, G_i) \preceq u_i$, so $x \preceq u_i/v_i$ for $i = 1, 2$, by Lemma 5.16. We obtain the following diagram:

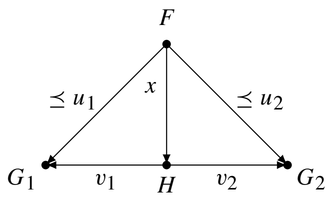

In particular, a flag F realises $\mathrm{P}_{u_1}(X, G_1) \wedge \mathrm{P}_{u_2}(X, G_2)$ if and only if there is some flag H occurring in some permutation of P with

$$\mathrm{P}_{u_1/v_1}(F, H) \wedge \mathrm{P}_{u_2/v_2}(F, H).$$

Lemma 7.4 applied to u_1/v_1 and u_2/v_2 yields reduced words $w_1, \ldots, w_n$ describing the above intersection. We need only show that $w_j \prec u_1, u_2$ for $j = 1, \ldots, n$. Clearly $w_j \preceq u_i/v_i \preceq u_i$. Suppose however that $w_j \approx u_1$ for some $j = 1, \ldots, n$, so $u_1 \preceq u_i/v_i$ and thus $u_1 \preceq [u_1 \cdot v_1] \preceq u_1$. Hence $\mathrm{P}_{u_1}(X, G_1) = \mathrm{P}_{u_1}(X, H)$. Also $u_1 \preceq [u_1 \cdot v_2] \preceq u_2$, so $\mathrm{P}_{u_1}(X, G_1) = \mathrm{P}_{u_1}(X, H) \subset \mathrm{P}_{u_2}(X, H) \subset \mathrm{P}_{u_2}(X, G_2)$, contradicting our hypothesis. □

Since the relation $\prec$ is well-founded, we obtain the following.

Corollary 7.6. *The formulae* $\mathrm{P}_u(X, Y)$ *are equations.*

By bi-interpretatibility, in order to conclude that the theory PS_Γ is equational, we need only show that every formula whose free variables enumerate flags is a Boolean combination of formulae $\mathrm{P}_u(X, G)$. For that, we will first introduce the notion of nice hulls.

A colour-preserving graph homomorphism $f : A \to B$ between two Γ-spaces induces a homomorphism $\chi(f) : \chi(A) \to \chi(B)$ between the chamber systems of flags of A and B. It is easy to see that this defines an isomorphism between the category of Γ-spaces and the category of dual quasi-buildings (cf. Theorem 2.17).

Definition 7.7. Suppose that both A and B are simply connected. Given $X \subset \chi(A)$ and $Y \subset \chi(B)$, the map $\phi : X \to Y$ is *contracting* if and only if for any flags F and G in X,

$$\mathrm{w}_B(\phi(F), \phi(G)) \preceq \mathrm{w}_A(F, G).$$

We say that ϕ is an *isometry* if

$$\mathrm{w}_B(\phi(F), \phi(G)) = \mathrm{w}_A(F, G).$$

The following is easy to see:

Lemma 7.8. *If A and B are simply connected, a map $\chi(A) \to \chi(B)$ is contracting if and only if it is a homomorphism of chamber systems.*

Definition 7.9. If $A \subset B$ are Γ-spaces, a *specialisation* from B to A is a homomorphism $f : B \to A$ such that $f\restriction_A = \mathrm{Id}_A$.

Remark 7.10. If $B = A \cup F$ is a simple extension of A of type (s, G), then the map f fixing all elements of A which sends the γ-vertex of F to the γ-vertex of G, for γ in s, is a specialisation.

Lemma 7.11. *Given Γ-spaces $A \subset B$, with B simply connected, the subspace A is nice in B if and only if B can be specialised to A.*

Proof. Suppose that B can be specialised to A, and let $P : F \xrightarrow{u} G$ be a reduced path in B connecting two flags F and G in A. The specialisation maps P to a connecting path P' in A between F and G with word $u' \preceq u$, so A is nice in B.

If A is nice in B, observe that A specialises to itself. Choose then a maximal specialisation $f : C \to A$, with $C \subset B$ nice in B. If $C \neq B$, Proposition 4.9 yields a proper simple extension C' of C, which is again nice in B. By Remark 7.10, there is a specialisation $C' \to C$. Composing it with f contradicts the maximality of f. □

Given two nice subsets N_1 and N_2 of our fixed saturated model M, an *isomorphism* means a bijection $f : N_1 \to N_2$ such that both f and f^{-1} are homomorphisms of Γ-spaces. If N_1 and N_2 have a common subset A , we say that they are A*-isomorphic*, denoted by $N_1 \simeq_A N_2$, if there is an isomorphism between N_1 and N_2 fixing all elements in A.

Definition 7.12. Let $N \subset M$ be nice and A be some subset of N.

(1) We say that N is a *nice hull* of A if every nice subset of M containing A has a nice subset N' which is A-isomorphic to N.

(2) The Γ-space N is *incompressible* over A if every A-homomorphism $f : N \to N$ is an automorphism of N.
(3) The nice subset N is *strongly incompressible* over A if every A-homomorphism $f : N \to M$ induces an isomorphism of N with a nice subset of M.

We will see in Proposition 7.18 that if N is incompressible over A, then the only A-endomorphism of N is the identity.

Lemma 7.11 implies the following easy observation.

Remark 7.13. If N is incompressible over A, then it contains no proper nice subset $A \subset N' \subsetneq N$. Likewise if N is strongly incompressible over A.

Lemma 7.14. (1) *If the nice set N is strongly incompressible over A, then it is incompressible and a nice hull of A.*
(2) *If N is a nice hull of A, and N' is incompressible over A, then $N \simeq_A N'$.*

Proof. For (1), let N be strongly incompressible over A. To show that N is incompressible over A, consider an A-homomorphism $f : N \to N$. It must induce an A-isomorphism with a nice subset N_1 of N containing A, so $N_1 = N$ by Remark 7.13.

Let us now show that N is a nice hull. Given any nice subset N' of M containing A, choose a specialisation $f : M \to N'$, by Lemma 7.11. The map $f\restriction_N$ must then induce an A-isomorphism of N with a nice subset of N', as desired.

Suppose now N and N' are as stated in (2). Since N is a nice hull, there is some nice subset N'' of N' which is A-isomorphic to N. We conclude that $N'' = N'$ by Remark 7.13. □

We now have all the necessary ingredients to deduce the following result.

Theorem 7.15. *Every subset A of M has a unique, up to A-isomorphism, strongly incompressible extension. If A is finite, so is this extension.*

Proof. We give the proof for A finite, and leave the general case to reader.

We proceed by induction on $|A|$. If A is empty, then any flag is strongly incompressible over $\emptyset$. Otherwise, write $A = A_0 \cup \{a\}$ and choose, by induction, a finite strongly incompressible extension N_0 of A_0. Among all flags passing through a, choose one, say F, with $\preceq$-minimal word $u = \mathrm{w}_M(F, G)$, where G is some base-point of F over N_0. Assume furthermore that u is $\preceq$-minimal among all possible A_0-copies of N_0. Let $P : F \xrightarrow{u} G$ be a reduced flag path connecting F to G. Set $N = N_0 \cup P$, which is a finite nice subset of M, by Lemma 6.4. In order to show that N is strongly incompressible over A, consider an A-homomorphism $f : N \to M$. By induction, the map $f\restriction_{N_0}$ induces an A_0-isomorphism between N_0 and the nice set $f(N_0)$, so $f(N_0)$ is also strongly incompressible over A_0. The map f is contracting, by Lemma 7.8, so $\mathrm{w}(f(F), f(G)) \preceq u$. Since a is contained in $f(F)$, minimality of u implies that $\mathrm{w}(f(F), f(G)) = u$, thus $f(G)$ is a base-point of $f(F)$ over $f(N_0)$. Therefore, the set $f(P)$ determines a reduced path from $f(F)$ to $f(G)$. Hence, the set $f(N)$ is nice, by Lemma 6.4, and f is an A-isomorphism, as desired. □

Together with Lemma 7.14, we obtain:

Corollary 7.16. *Every set A has a nice hull* $\mathrm{N}(A)$*, which is incompressible and unique, up to A-isomorphism. If A is finite, then so is* $\mathrm{N}(A)$.

Thus, the three notions in Definition 7.12 coincide.

Corollary 7.17. *The algebraic closure* $\mathrm{acl}(A)$ *of a finite set A is finite and contained in* $\mathrm{N}(A)$.

Proof. Let $\mathrm{N}(A)$ be the nice hull of a finite set A. Nice sets are algebraically closed, by Corollary 6.21. Thus, the set $\mathrm{acl}(A)$ is contained in $\mathrm{N}(A)$, which is finite. □

Proposition 7.18. *The nice hull* $\mathrm{N}(A)$ *is rigid over A, that is, its only automorphism fixing A pointwise is the identity.*

Proof. Again, we leave the case of infinite A to the reader and assume A is finite.

If A is empty, recall that $\mathrm{N}(\emptyset)$ consists of a single flag, so the result is obvious. By the proof of Theorem 7.15, if $A = A_0 \cup \{a\}$, then $\mathrm{N}(A) = N_0 \cup P$, where N_0 is a nice hull of A_0 and P is a reduced flag path connecting a flag F containing a to its basepoint G over N_0. Furthermore, the word $\mathrm{w}(F, G)$ is $\preceq$-minimal among all words $\mathrm{w}(F, G')$, where G' has the same type as G over A_0. If γ denotes the colour of a, then u has a unique beginning s, which contains γ.

By Remark 7.10, choose a specialisation $\phi : N_0 \cup P \to N_0$, collapsing the whole path P onto G. Let now f be some automorphism of $\mathrm{N}(A)$ fixing A. The map $(\phi \circ f)\restriction_{N_0}$, which is an A_0-automorphism by strong incompressibility of N_0, must then be the identity, by induction on $|A|$. Thus f is the identity on $N_0 \setminus G$.

Since f is an automorphism, we have $\mathrm{w}(f(F), f(G)) = \mathrm{w}(F, G) = u$. As the flag $f(F)$ lies in $N_0 \cup P$, Lemma 6.16 yields some flag K with $P : F \xrightarrow{u_1} K \xrightarrow{u_2} G$ where $\mathrm{w}(K, f(F))$ and u_2 commute. If P' denotes the subpath $K \xrightarrow{u_2} G$, the flag K is a basepoint of F over $N_0 \cup P'$. Thus, the word $\mathrm{w}(F, f(F))$ equals $[u_1 \cdot \mathrm{w}(K, f(F))]$. Since a is contained in $f(F)$, the flags F and $f(F)$ are $(\Gamma \setminus \{\gamma\})$-equivalent, so γ does not occur in $\mathrm{w}(F, f(F))$. If u_1 is not trivial, then s must be its beginning, for u has only s as a beginning, and hence either s or a larger letter containing it must occur in $\mathrm{w}(F, f(F))$. Thus, we conclude that $u_1 = 1$ and the word $\mathrm{w}(F, f(F))$ commutes with u.

As the flag $f(G)$ lies in $N_0 \cup P$, by Lemma 6.16 there is a flag K_1 with $P : F \xrightarrow{v_1} K \xrightarrow{v_2} G$ such that $\mathrm{w}(K, f(G))$ and v_2 commute. We have the following permutations:

$$\begin{aligned} v_1 \cdot v_2 \cdot \mathrm{w}(F, f(F)) &\approx \mathrm{w}(F, f(F)) \cdot u \approx \mathrm{w}(F, f(F)) \cdot \mathrm{w}(f(F), f(G)) \\ &\approx [v_1 \cdot \mathrm{w}(K, f(G))]. \end{aligned}$$

Corollary 5.13 implies that $v_2 \cdot \mathrm{w}(F, f(F))$ is a final subword of $\mathrm{w}(K, f(G))$, which commutes with v_2. We conclude that $v_2 = 1$, and hence $f(G)$ lies in N_0. The map f maps N_0 to itself, so it is the identity on N_0. Since f induces a permutation of P, Corollary 6.17 implies that f is the identity on N. □

Corollary 7.19. *The algebraic closure* $\mathrm{acl}(A)$ *of A is rigid over A, so it equals* $\mathrm{dcl}(A)$.

Proof. Note that $\mathrm{acl}(A)$ is contained in $\mathrm{N}(A)$, by Corollary 7.17, so $\mathrm{N}(A) = \mathrm{N}(\mathrm{acl}(A))$. Every A-automorphism of $\mathrm{acl}(A)$ extends to an automorphism of $\mathrm{N}(A)$, which must then be the identity, by Proposition 7.18. □

Proposition 7.20. *All types are stationary.*

This need no longer hold if we consider types over subsets of M^{eq}.

Proof of Proposition 7.20. We need only prove the statement for 1-types, and thus it suffices to prove it for types of a single flag. Let p be the type of a flag F over the parameter set A and let q be a global non-forking extension of p to M. Since $q = \mathrm{p}_u(G)|M$ for some flag G in A, its canonical parameter is $B = G/\mathcal{S}_{\mathrm{R}}(u)$, by Corollary 6.8, which is interdefinable with a set of real elements. Since B is algebraic over A, it is hence definable over A, by Corollary 7.19. The type p is thus stationary. □

We will show that the theory PS_Γ is equational by proving that the type of finitely many flags is determined by the collection of words connecting each pair of flags. For two flags this follows from Lemma 6.4: A reduced path $P : F \xrightarrow{u} G$ determines a nice subset, whose type is determined by u. For three flags F, G and H as in

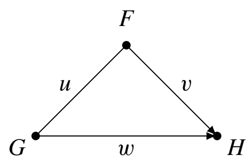

we need the following proposition.

Proposition 7.21. (1) *Given reduced words u, v and w with $u \cdot v \xrightarrow{*} w$, there is a decomposition*

$$u \approx u_1 \cdot \alpha^{-1} \cdot c^{-1}, \qquad v \approx c \cdot \beta \cdot v_1, \qquad w \approx u_1 \cdot x \cdot v_1$$

such that α, β and x pairwise commute, the word x is properly right-absorbed by c, the word α is properly left-absorbed by v_1, and β is right-absorbed by u_1. The words u_1, v_1, c, x, α and β are unique up to permutation.

(2) *Assume further that F, G and H are flags such that $\mathrm{w}(G, F) = u$, $\mathrm{w}(F, H) = v$ and $\mathrm{w}(G, H) = w$. Then there is a reduced path $P : G \xrightarrow{w} H$ and a base-point K of F over P such that*

$$\mathrm{w}(G, K) = u_1, \qquad \mathrm{w}(K, H) = x \cdot v_1, \qquad \mathrm{w}(F, K) = c \cdot \alpha \cdot \beta.$$

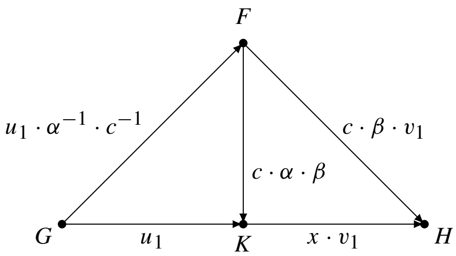

Observe that the existence of a decomposition as in (1) implies $u \cdot v \xrightarrow{*} w$. Indeed, since x is properly right-absorbed by $c = t_m \cdots t_1$, Lemma 5.14 implies that $x \approx x_1 \cdots x_m$, where x_i is a splitting t_i and commutes with t_j whenever $j < i$. Thus,

$$c^{-1} \cdot c = t_1 \cdots t_m \cdot t_m \cdots t_1 \xrightarrow{*} t_1 \cdots t_{m-1} \cdot x_m \cdot t_{m-1} \cdots t_1 \\ \xrightarrow{*} t_1 \cdots t_{m-1} \cdot t_{m-1} \cdot x_m \cdot t_{m-2} \cdots t_1 \xrightarrow{*} \cdots$$

can be reduced to x. Therefore,

$$u \cdot v \approx u_1 \cdot \alpha^{-1} \cdot c^{-1} \cdot c \cdot \beta \cdot v_1 \xrightarrow{*} u_1 \cdot \alpha^{-1} \cdot x \cdot \beta \cdot v_1 \\ \xrightarrow{*} u_1 \cdot \beta \cdot x \cdot \alpha^{-1} \cdot v_1 \xrightarrow{*} u_1 \cdot x \cdot v_1 \approx w.$$

By Lemma 6.15, we may assume in (2) that the flag K occurs in the path P.

Proof of Proposition 7.21. Let us first prove the existence of such a decomposition. By Remark 4.15, find flags F, G and H such that $G \xrightarrow{u} F \xrightarrow{v} H$ and $G \xrightarrow{w} H$. Let P be a reduced flag path connecting G to H. By Lemma 6.15, we may choose a base-point K of F over P which occurs in the path P. Set $\mathrm{w}(G, K) = w_1$, $\mathrm{w}(K, H) = w_2$, and $\mathrm{w}(F, K) = y$. Assume that y is $\preceq$-minimal among all choices of P, which implies

$$\text{no end of } y \text{ is contained in } \mathrm{Wob}(w_1, w_2). \tag{‡}$$

Indeed, if $y = y' \cdot s$, with $s \subset \mathrm{Wob}(w_1, w_2)$, decompose $F \xrightarrow{y'} K' \xrightarrow{s} K$ for some flag K'. Observe that K is also a base-point for K' over the nice set determined by P such by Proposition 6.1 the reduction of $s \cdot w_1^{-1}$, resp. $s \cdot w_2$, is non-splitting and equals w_1^{-1}, resp. w_2. Replacing K by K', we obtain a permutation P' of P such that F connects to P' with word y', contradicting the minimality of y.

Proposition 6.1 implies that $[y \cdot w_1^{-1}] = u^{-1}$ and $[y \cdot w_2] = v$. By Proposition 5.9, up to permutations of w_1 and u, we can write

$$w_1 = u_1 \cdot x', \quad y \approx c_1^{-1} \cdot \beta, \quad u = u_1 \cdot c_1,$$

where x' and β commute, the word x' is properly left-absorbed by c_1, and β is right-absorbed by u_1. Let K' be a flag in P such that $G \xrightarrow{u_1} K' \xrightarrow{x'} K$. Since x' is right-absorbed by y, the flag K' is also a base-point of F over P by Corollary 6.3. Replacing K' by K, we may assume that $x' = 1$.

Likewise, we can write

$$w_2 = x \cdot v_1, \quad y \approx c_2 \cdot \alpha, \quad v = c_2 \cdot v_1,$$

where x and α commute, the word x is properly right-absorbed by c_2, and α is left-absorbed by v_1.

Note that $y = c_2 \cdot \alpha \approx c_1^{-1} \cdot \beta$. However, no end of α can be an end of β, by (‡). Thus, every end of α commutes with β, and Lemma 5.3 implies that α commutes with β and is a final subword of c_1^{-1}. Likewise for β and c_2. After possible permutations of c_1 and c_2, we can write

$$c_1^{-1} = c \cdot \alpha, \quad c_2 = c \cdot \beta.$$

Let us now show that β and x commute. Otherwise, write $x = x_1 \cdot s \cdot x_2$, where x_1 and β commute, but β and s do not. Since x is right-absorbed by $c_2 = c \cdot \beta$, the letter s must be absorbed by β, and therefore right-absorbed by u_1. Since s commutes with x_1, the word $w = u_1 \cdot x \cdot v_1$ is not reduced, which is a contradiction. Hence, x is properly right-absorbed by c.

We now have

$$u = u_1 \cdot \alpha^{-1} \cdot c^{-1}, \quad v = c \cdot \beta \cdot v_1, \quad w = w_1 \cdot w_2 = u_1 \cdot x \cdot v_1, \quad y = c \cdot \alpha \cdot \beta.$$

The only property left to show is that α is properly left-absorbed by v_1. Otherwise, apply Corollary 5.12 to produce, up to permutation, the following decompositions:

$$\alpha = \alpha' \cdot \omega, \quad \omega \cdot v_2 = v_1,$$

where ω is a commuting word, the word α' is properly left-absorbed by v_2, and α' and ω commute. Since $w_2 \approx \omega \cdot x \cdot v_2$, there is a flag K' in some permutation of P such that $K \xrightarrow{\omega} K' \xrightarrow{x \cdot v_2} H$. Now, the word ω is right-absorbed by $y \approx c \cdot \beta \cdot \alpha' \cdot \omega$, so the flag K' is also a base-point of F over P, by Corollary 6.3. Replacing K by K' and substituting

$$u_1 \rightsquigarrow u_1 \cdot \omega, \quad v_1 \rightsquigarrow v_2, \quad \alpha \rightsquigarrow \alpha', \quad \beta \rightsquigarrow \beta \cdot \omega$$

gives the new words u_1, v_1, c, x, α, β and the new base-point K with the desired properties.

For uniqueness, let $u_1 \cdot a$ be the largest common initial subword of $u = u_1 \cdot \alpha^{-1} \cdot c^{-1}$ and $w = u_1 \cdot x \cdot v_1$ (cf. the discussion after Definition 5.1). Since α and x commute and each one is properly left absorbed by v_1, resp. c^{-1}, the word a must commute with α and x, and is an initial subword of both $c^{-1} \approx a \cdot c'$ and $v_1 \approx a \cdot v_1'$. In particular, the words $\alpha \cdot c'$ and $x \cdot v_1'$ are uniquely determined, up to a permutation, for $u_1 \cdot a$ is.

Observe that c' is the largest common final subword of $\alpha \cdot c'$ and $v^{-1} = {v_1'}^{-1} \cdot a^{-1} \cdot \beta^{-1} \cdot a \cdot c'$, for otherwise, the largest common final subword would then contain a letter s from α, which is an end of ${v_1'}^{-1} \cdot a^{-1} \cdot \beta^{-1} \cdot a$, contradicting α being properly left-absorbed by v_1'. Therefore, the words α and $a^{-1} \cdot \beta \cdot a \cdot v_1'$ are uniquely determined.

Since x and $a^{-1} \cdot \beta \cdot a$ commute, the word v_1' is the largest common final subword of $x \cdot v_1'$ and $a^{-1} \cdot \beta \cdot a \cdot v_1'$, so x and $a^{-1} \cdot \beta \cdot a$ are uniquely determined. The result now follows by applying the following auxiliary result to the words $u_1 \cdot a$ and $a^{-1} \cdot \beta \cdot a$, in order to determine u_1, a and β, as desired.

Claim. *Given reduced words $e \approx u \cdot a$ and $f \approx a^{-1} \cdot b \cdot a$, where u right-absorbs b, the words a, u and b are uniquely determined by e and f.*

Proof of the claim. We proceed by induction on the length of f. Clearly, if $f = 1$, then $b = a = 1$ and thus $e = u$.

Otherwise, note that e right-absorbs f if and only if $a = 1$, in which case $e = u$ and $f = b$. Therefore, if e does not right-absorb f, write $a = a' \cdot s$. In particular, the word e has an end which is simultaneously a beginning and an end of f.

We show first that the only possible ends t of e which are both a beginning and an end of f are exactly the ends of a. If not, the end t must be an end of u which commutes with a. Likewise $b \approx t \cdot b' \cdot t$, which contradicts u right-absorbing b, since b' and t do not commute.

Removing s yields the reduced words $u \cdot a'$ and $a'^{-1} \cdot b \cdot a'$, and the result now follows by induction. □

Corollary 7.22. *The type* $\mathrm{tp}(F, G, H)$ *of three flags* F*,* G *and* H *is uniquely determined by* $\mathrm{w}(F, G)$*,* $\mathrm{w}(G, H)$ *and* $\mathrm{w}(F, H)$*.*

Proof. Proposition 7.21 yields a reduced path $P : G \xrightarrow{w_1} K \xrightarrow{w_2} H$ from G to H, where K is a base-point of F over the nice set determined by P, such that w_1, w_2 and $\mathrm{w}(F, K)$ are uniquely determined, up to permutation, by the words $\mathrm{w}(F, G)$, $\mathrm{w}(G, H)$ and $\mathrm{w}(F, H)$. By Corollary 6.5, the type of F over P is uniquely determined by $\mathrm{w}(F, K)$ and K. Note that H is a base-point of G over the nice set H, so Lemma 6.4 shows that the type of $P = G, \ldots, K, \ldots, H$ is determined by the word $w_1 \cdot w_2$. Thus, the type of GKH is determined by the equivalence classes of w_1 and w_2, as desired. □

To extend this result to arbitrary sets of flags, we need the following lemma. Given a Γ-space A, recall that $\chi(A)$ denotes the chamber system of the flags in A (see the discussion before Definition 7.7).

Lemma 7.23. *A non-empty collection* X *of flags of* M *equals* $\chi(A)$ *for some nice subset* A *of* M *if and only if whenever* F *and* G *in* X *are connected by a reduced word* u *in* M*, there is a path in* X *with word* u *connecting* F *to* G*.*

If $X = \chi(A)$, the remark after Definition 4.5 implies that A is the union of the flags in X.

Proof of Lemma 7.23. One direction is an equivalent definition of niceness, so we need only show that, if X satisfies the right-hand condition, then $X = \chi(A)$ for some nice subset A. Let A be the collection of all vertices of flags in X. In order to show that $X = \chi(A)$ and A is nice, it suffices to show that any flag in A is a flag from X, that is, we need only show, by induction on $|S|$, that given a collection S of elements in A lying in a flag H in M, there is a flag in X containing S. If S is a singleton, there is nothing to prove. Otherwise, enumerate $S = \{a_1, \ldots, a_r\}$ and let $T \subset \Gamma$ be the collection of colours of $a_1, \ldots, a_{r-1}$, and γ the colour of a_r. Choose a flag F in X containing a_r and, by induction, a flag G in X containing $a_1, \ldots, a_{r-1}$. Observe that

$$F \sim_{\Gamma\setminus\{\gamma\}} H \sim_{\Gamma\setminus T} G,$$

so there are reduced words u and v such that γ does not occur in u, the support of v is disjoint from T and $F \xrightarrow{u} H \xrightarrow{v} G$. Reducing this path in M gives a word $w_1 \cdot w_2$, where γ does not occur in w_1 and each letter in w_2 is disjoint from T. By assumption, there is a reduced path in X of the form $F \xrightarrow{w_1} H' \xrightarrow{w_2} G$. Clearly, the path H' lies in X and contains the set S, as desired. □

We will now extend Corollary 7.22 to arbitrarily many flags.

Theorem 7.24. *The type* $\mathrm{tp}(F_1, \ldots, F_n)$ *of a sequence of flags is uniquely determined by* $\mathrm{w}(F_i, F_j)$ *for* $1 \le i \ne j \le n$.

Proof. Recall the notion of isometry, as in Definition 7.7. Given two collections of flags X and X', we will show that a surjective isometry $\phi : X \to X'$ induces an elementary map between nice subsets of M.

Suppose first that X satisfies the conditions of Lemma 7.23. Then so does X', and there are nice subsets A and A' such that $X = \chi(A)$ and $X' = \chi(A')$. In particular, the map ϕ induces a graph isomorphism $f : A \to A'$, for two flags F and G in A are γ-equivalent if and only if $\mathrm{w}(F, G)$ does not contain γ. By Remark 4.6(2), nice subsets are Γ-spaces, so f is an elementary map, by Corollary 4.13.

Thus, we need only show that ϕ can be extended to some supersets of flags which satisfy the conditions of Lemma 7.23. Given flags F and G in X with respective images F' and G' in X', if $\mathrm{w}(F, G) = u$, choose reduced paths $P : F \xrightarrow{u} G$ and $P' : F' \xrightarrow{u'} G'$ with

$$X \underset{F,G}{\perp} P \quad \text{and} \quad X' \underset{F',G'}{\perp} P'.$$

Let ψ be the isometry $P \to P'$ which maps (F, G) to (F', G'). In order to show that $\phi \cup \psi$ is a well-defined isometry between $X \cup P$ and $X' \cup P'$, consider some flag H in X. It suffices to show that $\phi \cup \psi$ induces an isometry $H \cup P \to \phi(H) \cup P'$. By Corollary 7.22, the map $\phi \restriction_{H,F,G}$ is elementary. So is ψ, by Lemma 6.4. As the type $\mathrm{tp}(H/F, G)$ is stationary by Proposition 7.20, it follows that $\phi \restriction_{H,F,G} \cup \psi$ is an elementary map and thus an isometry.

Iterating the above process countably many times, we obtain the desired superset satisfying the conditions of Lemma 7.23. □

The above result and Corollary 7.6 yield the following, by compactness.

Corollary 7.25. *The theory* PS_Γ *is equational.*

Equationality, or rather Theorem 7.24, allows us to show total triviality and hence weak elimination of imaginaries.

Proposition 7.26 (cf. [2, Lemma 7.22]). *The theory* PS_Γ *is totally trivial: over any set* D *of parameters, if* a, b *and* c *are tuples such that* a *is independent both from* b *and from* c *over* D, *then* a *is independent from* b, c *over* D.

By induction on the length of the tuples (cf. [5, Lemma 4]), it is easy to see that it suffices to check total triviality for singletons a, b and c.

Proof of Proposition 7.26. By taking a non-forking extension to a small submodel containing D, we may assume that D is nice. Since every element is contained in a flag, it suffices to show total triviality when a, b and c enumerate three flags F, H and K.

Let G be a base-point of F over D. In particular, we have $F \perp_G H$ and $F \perp_G K$.

Choose another realisation F' of $\mathrm{tp}(F/G)$ such that $F' \perp_G H, K$. Then $\mathrm{w}(F', G) = \mathrm{w}(F, G)$ and

$$\mathrm{w}(F', H) = [\mathrm{w}(F', G) \cdot \mathrm{w}(G, H)] = [\mathrm{w}(F, G) \cdot \mathrm{w}(G, H)] = \mathrm{w}(F, H),$$

since $F \mathop{\smile\hspace{-0.9em}\mid\hspace{0.4em}}_G H$. Likewise for $\mathrm{w}(F, K)$. Theorem 7.24 implies that F and F' have the same type over G, H, K, so $F \mathop{\smile\hspace{-0.9em}\mid\hspace{0.4em}}_G H, K$, as desired. □

Since PS_Γ is ω-stable, we obtain the following by [5, Proposition 7].

Corollary 7.27. *The theory* PS_Γ *is perfectly trivial, that is, given any set D of parameters and tuples a, b and c such that a and b are both independent over D, they are also independent over $D \cup \{c\}$.*

An ω-stable theory T has *weak elimination of imaginaries* if the canonical base of every stationary type can be chosen to be a subset of the real sort. In this case, types over algebraically closed sets are always stationary.

Corollary 7.28. *The theory* PS_Γ *has weak elimination of imaginaries, and given any stationary type* $\mathrm{tp}(a_1, \ldots, a_n/B)$*, we have*

$$\mathrm{Cb}(\mathrm{tp}(a_1 \ldots a_n/B)) = \mathrm{Cb}(\mathrm{tp}(a_1/B)) \cup \cdots \cup \mathrm{Cb}(\mathrm{tp}(a_n/B)).$$

Proof. Considering a small elementary substructure N, we may assume that the stationary type p equals $\mathrm{tp}(a_1, \ldots, a_n/N)$. Suppose that the canonical base of every unary type over N is interdefinable with a finite subset of the real sort. Thus, we may choose finite subsets $C_1, \ldots, C_n$ such that $\mathrm{Cb}(\mathrm{tp}(a_i/N))$ is interdefinable with C_i. Total triviality (Proposition 7.26) implies that p does not fork over $C = C_1 \cup \cdots \cup C_n$. By Proposition 7.20, the restriction of p to C is stationary, so C is interdefinable with $\mathrm{Cb}(p)$.

Therefore, we need only show that the canonical base of every unary $\mathrm{tp}(a/N)$ over the nice set N is interdefinable with a finite subset of the real sort. Given a flag F containing a independent from N, the type $\mathrm{tp}(F/a)$ is stationary, so $\mathrm{Cb}(\mathrm{tp}(a/N))$ is interdefinable with $\mathrm{Cb}(\mathrm{tp}(F/N))$. Corollary 6.8 now yields the desired result. □

Corollary 7.29. *The canonical base of a type is algebraic over any two independent realisations.*

Proof. Again by Proposition 7.26 and taking small elementary substructures, we need only consider a unary type p over some nice set A. Let a and a' be independent realisations of p. Choose a flag F containing a independent from A over a. Since the type of F over a is stationary, it follows that $\mathrm{Cb}(a/A)$ is interdefinable with $\mathrm{Cb}(F/A)$. By automorphisms, we may find a flag F' containing a' with the same type as F over A and

$$F' \mathop{\smile\hspace{-0.9em}\mid\hspace{0.4em}}_{a'} AF.$$

In particular,

$$FF' \mathop{\smile\hspace{-0.9em}\mid\hspace{0.4em}}_{a,a'} A \quad \text{and} \quad F \mathop{\smile\hspace{-0.9em}\mid\hspace{0.4em}}_{A} F',$$

so if $\mathrm{Cb}(F/A) \subset \mathrm{acl}(F, F') \cap A$, then $\mathrm{Cb}(a/A) = \mathrm{Cb}(F/A) \subset \mathrm{acl}(a, a')$, as desired.

Therefore, we need only prove the statement for the type of a flag F over A. Let $F' \mathop{\smile\hspace{-0.9em}\mid\hspace{0.4em}}_A F$ and choose a base-point G of F over A. Suppose that $P : F \xrightarrow{u} G$ and let

$\tilde{u}$ be the final segment of $u \approx u_1 \cdot \tilde{u}$. Then $w(F, F') = [u \cdot u^{-1}] = u_1 \cdot \tilde{u} \cdot u_1{}^{-1}$, by Corollary 5.15.

We obtain the following diagram:

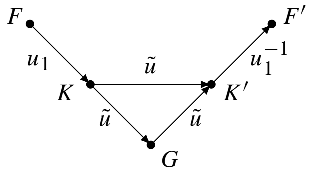

Since $\mathrm{Wob}(u_1, \tilde{u} \cdot u_1^{-1}) \subset \mathcal{S}_{\mathrm{R}}(u)$, Lemma 6.19 implies that $K/{\sim_{\mathcal{S}_{\mathrm{R}}(u)}}$ lies in $\mathrm{acl}(F, F')$. Observe that $K \sim_{\tilde{u}} G$ and $\tilde{u} \preceq \mathcal{S}_{\mathrm{R}}(u)$, by Corollary 5.18. Thus, the canonical base $\mathrm{Cb}(F/A)$ is contained in $\mathrm{acl}(F, F')$ as well. □

8. Ampleness

Recall the definition of n-ampleness [13, 4].

Definition 8.1. A stable theory T is *n-ample* if, working inside a sufficiently saturated model and possibly over parameters, there are real tuples $a_0, \ldots, a_n$ such that:

(1) $\mathrm{acl}^{\mathrm{eq}}(a_0, \ldots, a_i) \cap \mathrm{acl}^{\mathrm{eq}}(a_0, \ldots, a_{i-1}, a_{i+1}) = \mathrm{acl}^{\mathrm{eq}}(a_0, \ldots, a_{i-1})$ for every $0 \leq i < n$,
(2) $a_{i+1} \mathop{\smile\!\!\!\!\!\!\,|}_{a_i} a_0, \ldots, a_{i-1}$ for every $1 \leq i < n$,
(3) $a_n \not\mathop{\smile\!\!\!\!\!\!\,|} a_0$.

Note that T is n-ample if and only if T^{eq} is [2, Corollary 2.4]. Furthermore, if T is n-ample, it is $(n-1)$-ample. A theory is 1-based if and only if it is not 1-ample. It is CM-trivial if and only if is not 2-ample.

In order to find an upper bound for the ample degree of PS_Γ, we will use the following result.

Lemma 8.2 ([2, Remarks 2.3 and 2.5]). *If T is n-ample, there are tuples $a_0, \ldots, a_n$ enumerating small elementary substructures of an ambient saturated model such that for every $0 \leq i < n-1$:*

(a) $a_n \mathop{\smile\!\!\!\!\!\!\,|}_{a_{i-1}} a_i$,
(b) $\mathrm{acl}^{\mathrm{eq}}(a_i, a_{i+1}) \cap \mathrm{acl}^{\mathrm{eq}}(a_i, a_n) = \mathrm{acl}^{\mathrm{eq}}(a_i)$,
(c) $a_n \not\mathop{\smile\!\!\!\!\!\!\,|}_{\mathrm{acl}^{\mathrm{eq}}(a_i) \cap \mathrm{acl}^{\mathrm{eq}}(a_{i+1})} a_i$.

Recall that, given a subset X (of some cartesian power) of a structure M, the *induced structure* on X is the set of all relations on every cartesian power of X which are definable in M without parameters.

Lemma 8.3. *Let X be a subset, definable without parameters, of a model M of a stable theory T. If the theory of X equipped with the induced structure is n-ample, then so is T.*

Proof. We may assume that M is sufficiently saturated. Let $a_0, \ldots, a_n$ in X witness that X is n-ample for some n in $\mathbb{N}$. Since X is equipped with the full induced structure from M, properties (2) and (3) of Definition 8.1 hold when we consider these tuples in M. Furthermore, working in M, we have

$$\mathrm{acl}^{\mathrm{eq}}(a_0, \ldots, a_i) \cap \mathrm{acl}^{\mathrm{eq}}(a_0, \ldots, a_{i-1}, a_{i+1}) \cap X^{\mathrm{eq}} = \mathrm{acl}^{\mathrm{eq}}(a_0, \ldots, a_{i-1}) \cap X^{\mathrm{eq}},$$

where X^{eq} denotes those imaginary elements of M^{eq} having some representative in X. The following result together with (1) yields the desired result.

Claim. *If A and B are subsets of X, then*

$$\mathrm{acl}^{\mathrm{eq}}(A) \cap \mathrm{acl}^{\mathrm{eq}}(B) = \mathrm{acl}^{\mathrm{eq}}\big(\mathrm{acl}^{\mathrm{eq}}(A) \cap \mathrm{acl}^{\mathrm{eq}}(B) \cap X^{\mathrm{eq}}\big).$$

Proof of the Claim. If e lies in $\mathrm{acl}^{\mathrm{eq}}(A) \cap \mathrm{acl}^{\mathrm{eq}}(B)$, it is witnessed by finite definable sets $E_a = \varphi(x, a)$, resp. $F_b = \psi(x, b)$, with a in A, resp. b in B. The canonical parameter d of the finite set $E_a \cap F_b$, which contains e, belongs to X^{eq}, since X^{eq} is definably closed in M^{eq}. It lies in $\mathrm{acl}^{\mathrm{eq}}(A)$, resp. $\mathrm{acl}^{\mathrm{eq}}(B)$, because if E_a, resp. F_b, is fixed, there are only finitely many possibilities for the subset $E_a \cap F_b$. □

As in the previous sections, we will work inside a sufficiently saturated Γ-space M, a model of PS_Γ.

A subgraph $Y \subset \Gamma$ is *full* if whenever two vertices x and y in Y are adjacent in Γ, they are so in Y.

Corollary 8.4. *Let Γ' be a full subgraph of Γ and F a fixed flag in M. Consider the Γ'-residue of F,*

$$X = \{F' \text{ flag in } M \mid F' \sim_{\Gamma'} F\}.$$

The set X is the collection of flags of some nice set A. The restriction $M' = \mathcal{A}_{\Gamma'}(A)$ to vertices with colours in Γ' is a model of $\mathrm{PS}_{\Gamma'}$. If $\mathrm{PS}_{\Gamma'}$ is n-ample, so is PS_Γ.

Proof. If two flags from X are connected by a reduced flag-path P with word u, the $|u|$ is a subset of Γ', by simple connectedness. Thus, all flags in P belong to X, so $X = \chi(A)$ for some nice subset A, by Lemma 7.23.

The restriction $G \mapsto G' = G\restriction_{\Gamma'}$ is a bijection between the flags in A (i.e. the elements of X) and the flags of M'. Since Γ' is full, the Coxeter group generated by Γ' is a subgroup of (W, Γ), so a reduced word on the letters of Γ' remains so as a word in Γ. It follows that M' is a simply connected Γ'-space. For every $\gamma \in \Gamma'$ and G in X, every flag in M which is γ-equivalent to G belongs again to X. Thus M' is a model of $\mathrm{PS}_{\Gamma'}$.

In order to show that if $\mathrm{PS}_{\Gamma'}$ is n-ample, so is PS_Γ, we need only show, by Lemma 8.3, that the induced structure on M' (as a definable subset of M with parameters in F) coincides with the structure of M' as a model of $\mathrm{PS}_{\Gamma'}$. That is, every definable relation on M' which is definable in M over F is then F-definable in the Γ'-space M', or equivalently, a type in M' over F determines a unique type in M over F. Assume therefore that the tuples c_1 and c_2 have the same type in M' over F'. If we choose nice sets D'_i in M', for $i = 1, 2$, containing F', c_i, there is an elementary map $f' : D'_1 \to D'_2$ in M' which maps c_1 to c_2 and is the identity on F'. The sets $D_i = D'_i \cup (F \setminus M')$ are clearly nice in A, and hence in M. The map f' extends to an elementary map between D_1 and D_2, which is the identity on F. Thus the tuples c_1 and c_2 have the same type over F in M, as desired. □

Recall (Definition 5.1) that a letter s is an end of the word u if $u \approx v \cdot s$. The final segment of u, denoted by $\tilde{u}$, is the commuting subword consisting of all ends of u. When a word w is commuting, we will identify it with its support $|w|$.

Lemma 8.5 ([2, Lemma 8.2 and Proposition 8.3]). *Consider nice sets A and B and a flag F such that $F \mathop{\smile\!\!\!\!\!\mid}_B A$ and $\mathrm{acl}(AB) \cap \mathrm{acl}(AF) = \mathrm{acl}(A)$. Let $u = u_B$ (resp. u_A) be the minimal word connecting F to a flag G_B in B (resp. G_A in A). Consider the reduced word v which connects G_B to G_A and the associated symmetric decomposition*

$$u = u_1 \cdot u' \cdot w, \qquad w \cdot v' \cdot v_1 = v,$$

as in Proposition 5.9. Then the word $w \cdot v_1$ is commuting. Furthermore, if

$$F \mathop{\not\!\smile\!\!\!\!\!\mid}_{A\cap B} A,$$

then

$$|v'| \not\subseteq \tilde{u} \subsetneq \tilde{u}_A,$$

where $\tilde{u}$ and $\tilde{u}_A$ are the final segments of u and u_A, respectively.

Proof. By transitivity of non-forking, we have $F \mathop{\smile\!\!\!\!\!\mid}_{G_B} G_A$, so

$$u_A = [u \cdot v] = u_1 \cdot w \cdot v_1.$$

We will first show that $w \cdot v_1$ is commuting. Considering its final segment, write $w \cdot v_1 \approx a \cdot b$ where b is a commuting word. In order to show that a is trivial, we need only show that b and a commute.

Now, the word $u' \cdot w$ is left absorbed by $w \cdot v_1$, so write $u' \cdot w \approx u_a \cdot u_b$, by Lemma 5.14, where u_a is left-absorbed by a, and u_b commutes with a and is left-absorbed by b. Choose a flag path $P : G_B \xrightarrow{v' \cdot a} H \xrightarrow{b} G_A$, for some flag H. Choosing P independent from F over G_B, G_A, we may assume that $F \mathop{\smile\!\!\!\!\!\mid}_{G_B} P$, so G_B is a base-point of F over the nice set determined by P. Therefore

$$\mathrm{w}(F, H) = [u \cdot v' \cdot a] = u_1 \cdot u' \cdot w \cdot v' \cdot a = u_1 \cdot u_b \cdot u_a \cdot v' \cdot a = u_1 \cdot a \cdot u_b.$$

Choose now a flag K with $F \xrightarrow{u_1 \cdot a} K \xrightarrow{u_b} H$. Observe that $\mathrm{w}(K, G_A) \preceq [u_b \cdot b] = b$. On the other hand, $\mathrm{w}(F, G_A) = [u \cdot v] = u_1 \cdot w \cdot v_1 \approx u_1 \cdot a \cdot b$, so $\mathrm{w}(K, G_A) = b$.

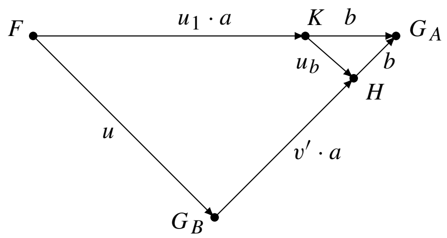

Let $\mathrm{C}(a)$ be the set of colours commuting with a. Clearly $\mathrm{C}(a) \cap \mathcal{S}_{\mathrm{R}}(a) = \emptyset$ and both $\mathrm{Wob}(v' \cdot a, b)$ and $\mathrm{Wob}(u_1 \cdot a, b)$ are subsets of $S = \mathcal{S}_{\mathrm{R}}(a) \cup \mathrm{C}(a)$. Lemma 6.19 implies that H/S lies in $\mathrm{acl}(AB)$ and K/S in $\mathrm{acl}(AF)$. Since u_b commutes with a and $K \xrightarrow{u_b} H$, we see that

$$K/S = H/S \in \mathrm{acl}(AB) \cap \mathrm{acl}(AF) = \mathrm{acl}(A).$$

Lemma 6.20 implies that $b \subset S$. However, no letter of b is contained in $\mathcal{S}_{\mathrm{R}}(a)$, since $a \cdot b$ is reduced and b is commuting, so b and a commute, as desired.

Let us now show that $\tilde{u} \subset \tilde{u}_A$. Since $a = 1$ and $b = w \cdot v_1$, the previous diagram yields the following picture:

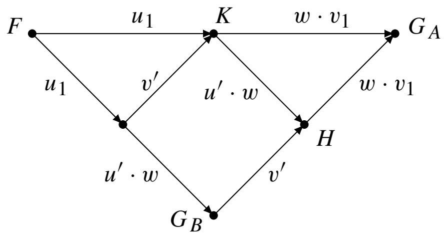

The final segment $\tilde{u}$ equals $\hat{u}_1 \cdot \tilde{u}' \cdot w$ for some commuting final subword $\hat{u}_1$ of u_1, which commutes with $u' \cdot w$. Since v' is (properly) right-absorbed by u_1, we have $\mathcal{S}_{\mathrm{R}}(v') \subset \mathcal{S}_{\mathrm{R}}(u_1) \subset T = \hat{u}_1 \cup \mathrm{C}(\hat{u}_1)$, so

$$\mathrm{Wob}(v', w \cdot v_1) \subset \mathrm{Wob}(u_1, w \cdot v_1) \subset T.$$

Note that $|u' \cdot w| \subset T$ and as above, the flag K/T lies in $\mathrm{acl}(A)$. Lemma 6.20 implies that $|w \cdot v_1| \subset T$. Since the word $\hat{u}_1 \cdot w \cdot v_1$ is reduced, no letter of the commuting word $w \cdot v_1$ is contained in $\hat{u}_1$. Therefore $w \cdot v_1$ and $\hat{u}_1$ commute, so

$$\tilde{u} = \hat{u}_1 \cdot w \cdot \tilde{u}' \subset \hat{u}_1 \cdot w \cdot v_1 = \tilde{u}_A.$$

To conclude, notice that if $|v'| \subset \mathcal{S}_{\mathrm{R}}(u_A)$, then $G_A/\mathcal{S}_{\mathrm{R}}(u_A) = G_B/\mathcal{S}_{\mathrm{R}}(u_A)$, which would imply $F \mathop{\smile\!\!\!\!\!\!|}_{A \cap B} A$ by Corollary 6.8. Therefore, if $F \not\mathop{\smile\!\!\!\!\!\!|}_{A \cap B} A$, then neither $\tilde{u} = \tilde{u}_A$ nor $|v'| \subset \tilde{u}$: For the former, if $\tilde{u} = \tilde{u}_A$, then $\tilde{u}' = v_1$ and hence $u' = v_1 = 1$, since $\tilde{u}'$ is properly absorbed by v_1. For the latter, if $|v'| \subset \tilde{u}$, then $|v'| \subset \tilde{u}_A \subset \mathcal{S}_{\mathrm{R}}(u_A)$. $\square$

If the graph Γ has no edges, PS_Γ is the theory of an infinite set M partitioned into $|\Gamma|$ infinite sets $\mathcal{A}_\gamma$. This is a trivial theory of Morley rank 1 (and degree $|\Gamma|$) which is easily seen not to be 1-ample.

For a graph with at least one edge, we define its *minimal valency* as the minimum of the valencies of non-isolated vertices. In particular, it is at least 1.

Theorem 8.6. *Let Γ be a graph with at least one edge. Let r be its minimal valency and n in $\mathbb{N}$ be maximal such that the graph* $[0, n]$*:*

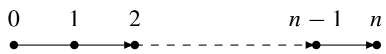

embeds as a full subgraph of Γ. Then the theory PS_Γ is n-ample but not $(|\Gamma| - r + 1)$-ample.

If Γ contains a full subgraph isomorphic to $[0, n]$ for some $n \in \mathbb{N}$, then $n \leq |\Gamma|$ and $r \leq |\Gamma| - n$, since the graph $[0, n]$ has minimal valency 1. Thus $|\Gamma| - r + 1$ is always bigger than n, as expected. For the graph $[0, n]$, the theorem says that its associated theory is n-ample but not $(n + 1)$-ample (cf. [15, Theorem 3.3], [2, Theorem 8.4]), hence the bounds are best possible; similarly for the graph consisting of $0, \dots, n + 1$ arranged in a circular way:

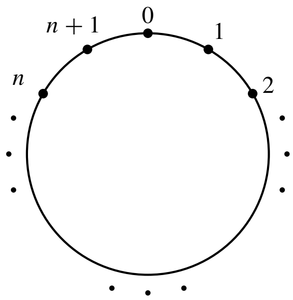

which has valency 2, so its theory is n-ample yet not $(n + 1)$-ample.

In particular, the theory of Γ is not 1-based if and only if Γ contains at least one edge. The complete graph $\mathbb{K}_n$ has minimal valency $n - 1$ and the theory $\mathrm{PS}_{\mathbb{K}_n}$ is CM-trivial for every n.

Proof of Theorem 8.6. Suppose that $[0, n]$ embeds as a full subgraph in Γ. Fix some flag F and consider the collection of flags $[0, n]$-equivalent to F. Corollary 8.4 and [2, Theorem 8.4] imply that PS_Γ is n-ample.

Suppose now that PS_Γ is N-ample for some natural number N, and let $a_0, \dots, a_N$ be enumerations of small models as in Lemma 8.2. Total triviality of PS_Γ implies that we may replace a_N by a flag F. For $0 \leq i \leq N - 1$, let u^i be the reduced word connecting F to a base-point G_i in the nice set a_i. Lemma 8.5 applied to each triangle (F, a_i, a_{i+1}) implies that the final segment $\tilde{u}^{i+1}$ of u^{i+1} is properly contained in the final segment $\tilde{u}^i$ of u^i. In particular $|\tilde{u}^1| \geq N - 1$.

Let v be the reduced word connecting G_1 to G_0. Lemma 8.5 implies the existence of a word v' which is properly right-absorbed by u^1 and such that $|v'|$ is not contained in $\tilde{u}^1$. Let γ be in $|v'| \setminus \tilde{u}^1$. Since v' is absorbed by u^1, so is γ. Thus γ commutes with $\tilde{u}^1$ and must be properly right-absorbed by u^1. Hence γ is not isolated and has valency at most $|\Gamma| - |\tilde{u}^1| - 1$. So $r \leq |\Gamma| - N$, that is, $N < |\Gamma| - r + 1$. □

Acknowledgments. We thank the referee for a careful reading and for providing many valuable comments, which helped us improve the presentation of this article.

The second author conducted research with support of programme ANR-09-BLAN-0047 Modig, ANR-13-BS01-0006 Valcomo as well as an Alexander von Humboldt-Stiftung Forschungsstipendium für erfahrene Wissenschaftler 3.3-FRA/1137360 STP.

References

[1] Baudisch, A.: Kommutationsgleichungen in semifreien Gruppen. Acta Math. Acad. Sci. Hungar. **29**, 235–249 (1977) Zbl 0381.20027 MR 0463300

[2] Baudisch, A., Martin-Pizarro, A., Ziegler, M.: Ample hierarchy. Fund. Math. **224**, 97–153 (2014) Zbl 1345.03057 MR 3180584
[3] Baudisch, A., Pillay, A.: A free pseudospace. J Symbolic Logic **65**, 443–460 (2000) Zbl 0947.03050 MR 1782130
[4] Evans, D.: Ample dividing. J. Symbolic Logic **68**, 1385–1402 (2003) Zbl 1067.03045 MR 2017362
[5] Goode, J. B.: Some trivial considerations. J. Symbolic Logic **56**, 624–631 (1991) Zbl 0741.03015 MR 1133089
[6] Grundhöfer, T.: Basics on buildings. In: Tits Buildings and the Model Theory of Groups, London Math. Soc. Lecture Note Ser. 291, Cambridge Univ. Press, 1–21 (2002) Zbl 1035.51005 MR 2018380
[7] Haglund, F., Paulin, F.: Constructions arborescentes d'immeubles. Math. Ann. **325**, 137–164 (2003) Zbl 1025.51014 MR 1957268
[8] Hrushovski, E.: On stable non-equational theories. Unpublished notes (1991)
[9] Hrushovski, E.: A new strongly minimal set. Ann. Pure Appl. Logic **62**, 147–166 (1993) Zbl 0804.03020 MR 1226304
[10] Hrushovski, E., Pillay, A.: Weakly normal groups. In: Logic Colloquium '85 (Orsay, 1985), Stud. Logic Found. Math. 122, North-Holland, 233–244 (1987) Zbl 0636.03028 MR 0895647
[11] Junker, M.: A note on equational theories. J. Symbolic Logic **65**, 1705–1712 (2000) Zbl 0973.03046 MR 1812175
[12] Pillay, A.: The geometry of forking and groups of finite Morley rank. J. Symbolic Logic **60**, 1251–1259 (1995) Zbl 0845.03016 MR 1367208
[13] Pillay, A.: A note on CM-triviality and the geometry of forking. J. Symbolic Logic **65**, 474–480 (2000) Zbl 0945.03047 MR 1782132
[14] Pillay, A., Srour, G.: Closed sets and chain conditions in stable theories. J. Symbolic Logic **49**, 1350–1362 (1984) Zbl 0597.03018 MR 0771800
[15] Tent, K.: The free pseudospace is N-ample, but not $(N+1)$-ample. J. Symbolic Logic **79**, 410–428 (2014) Zbl 1337.03047 MR 3224974
[16] Tent, K., Ziegler, M.: A Course in Model Theory. Lecture Notes in Logic 40, Cambridge Univ. Press (2012) Zbl 1245.03002 MR 2908005
[17] Ziegler, M.: Strong Fraïssé limits. Preprint, http://home.mathematik.uni-freiburg.de/ziegler/preprints/starker˙fraisse.pdf (2011)