UC Irvine

UC Irvine Electronic Theses and Dissertations

Title

Continuous EF Games between Asymptotic Cones

Permalink

https://escholarship.org/uc/item/8w25437h

Author

Schirle, Jessica Marie

Publication Date

2025

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial-ShareAlike License, available at https://creativecommons.org/licenses/by-nc-sa/4.0/

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

UNIVERSITY OF CALIFORNIA,
IRVINE

Continuous EF Games between Asymptotic Cones

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Mathematics

by

Jessica Schirle

Dissertation Committee:
Professor Isaac Goldbring, Chair
Professor Jesse Wolfson
Professor Martin Zeman

2025

© 2025 Jessica Schirle

DEDICATION

To Regina

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGMENTS

I would like to thank all of the UCI Math graduate students I have come to befriend over the years, as without their continued support, I surely would have left graduate school long ago.

I am also grateful to the organizers of the Young Women in Model Theory and Applications workshop for inviting me to present on my research in Bonn. It is always inspiring to see the incredible work of other logicians and to allow that to spur me into doing better work of my own.

I would also like to thank my church community in Orange County. I am indebted to those who have helped me find community outside academia and reminded me that the kingdom of God is something built, not found.

And last but not least, I thank my advisor, who for the past six years has been nothing but kind, patient, and gracious, even when I did not deserve it. His gentle encouragement and confidence in my own mathematical ability has carried me through a global pandemic and much more.

VITA

Jessica Schirle

EDUCATION

Doctor of Philosophy in Mathematics — **2025**
University of California, Irvine — *Irvine, California*

Master of Science in Mathematics — **2020**
University of California, Irvine — *Irvine, California*

Bachelor of Science in Mathematics — **2018**
University of Illinois at Urbana-Champaign — *Urbana, Illinois*

RESEARCH EXPERIENCE

Graduate Student Researcher — **2020–2024**
University of California, Irvine — *Irvine, California*

TEACHING EXPERIENCE

Teaching Assistant — **2018–2024**
University of California, Irvine — *Irvine, California*

ABSTRACT OF THE DISSERTATION

Continuous EF Games between Asymptotic Cones

By

Jessica Schirle

Doctor of Philosophy in Mathematics

University of California, Irvine, 2025

Professor Isaac Goldbring, Chair

We adapt a partial notion of EF games to a continuous logic for unbounded pointed metric structures, and use this to investigate the elementary equivalence of certain metric structures. Particular focus is placed upon EF games between asymptotic cones of symmetric spaces X arising from semisimple Lie groups. Kramer, et al. showed that, depending on the truth of CH, there is, up to homeomorphism, either 1 or $2^{2^{\aleph_0}}$-many asymptotic cones of X as one varies the choice of ultrafilter. This leaves open the possibility that all such asymptotic cones are elementarily equivalent.

Towards a proof of the elementary equivalence, we utilize the fact that these asymptotic cones are known to be isometric to the point spaces of certain nondiscrete affine $\mathbb{R}$-buildings. We investigate the building structure and demonstrate the elementarily equivalence of parallel classes of walls, which are fundamental to the classification of affine $\mathbb{R}$-buildings.

Chapter 1

Introduction

Geometric group theory is guided by the perspective that one may study a group G from the spaces upon which G acts. One may always equip G with some group presentation, which gives rise to a Cayley graph upon which G acts by isometries using left multiplication. The Cayley graph is unique to a given presentation, but not to a group itself, as one group may have many different presentations. However, when we restrict to the case of finitely generated presentations, it is unique up to quasi-isometry, and indeed up to bi-Lipschitz homeomorphism. In some sense, groups which are quasi-isometric "look the same from far away," and we may formalize this statement by means of an "asymptotic invariant" known as an asymptotic cone of the group.

Asymptotic cones arose from Gromov's theorem on groups of polynomial growth [16], as the proof requires a notion of a "limit" of metric spaces. Later, van den Dries and Wilkie gave a formalization of asymptotic cones using nonstandard analysis, which depends on a choice of ultrafilter on the naturals [40]. This suggests the following imprecise question: how much does an asymptotic cone depend on the choice of ultrafilter? More precisely, Gromov asked in [17] whether there exists a finitely presented group with at least two asymptotic cones up to homeomorphism, and even remarked that this should be studied via model theory. In [36], Thomas and Velickovic found an example of a finitely generated group with at least two asymptotic cones up to homeomorphism. Later, Kramer, Shelah, Tent, and Thomas found a finitely presented group such that, if the Continuum Hypothesis fails, then there are $2^{\mathfrak{c}}$-many asymptotic cones up to homeomorphism [26]; however, in the case that CH is true, their example has a unique asymptotic cone up to homeomorphism, and so it was still not known whether it was provable in ZFC that such a finitely presented group existed. Finally, in [33], Ol'shanskii and Sapir were able to find a finitely presented group with at least two non-homeomorphic asymptotic cones, regardless of the truth of CH.

The fact that asymptotic cones may be formalized as metric ultraproducts makes it natural to study them from the point of view of *continuous* model theory. In [29], Luther develops a version of continuous model theory for unbounded, pointed metric structures that allows for formulas equipped with enough free variables to allow for geometric objects, such as rays and geodesics, to be the zero sets of bona fide $\mathcal{L}$-formulae. This, along with characterizing definability within a class through zero sets being preserved under ultraproducts, allows for a study of asymptotic geometry through the definability of rays. As we mentioned, the asymptotic cones studied by KSTT may not be homeomorphic, but any good model theorist knows that elementary equivalence of structures is weaker than knowing that they are isomorphic. It is thus natural to refine our question: is there a finitely presented group such that the asymptotic cones have different $\mathcal{L}$-theories as one varies the choice of ultrafilter?

This question is more difficult to answer. As we are dealing with ultra*products* and not ultra*powers*, a direct application of Łoś' Theorem is not possible, and we ought to utilize the old-fashioned method of EF games, suitably adapted to our version of unbounded continuous model theory. Further, one may naïvely suspect that it is possible to strategize for an EF game by simply playing "index by index" in the ultraproduct. That is, in $\prod_{i\in I}\mathcal{M}_i/\mathcal{U}$, each $\mathcal{M}_i$ is elementarily equivalent to itself, hence if Player I plays some $[(a_i)]_{\mathcal{U}} \in \prod\mathcal{M}_i/\mathcal{U}$, then perhaps Player II could win by responding with the equivalence class of a sequence (b_i) where each b_i is determined according to a winning strategy in the structure $\mathcal{M}_i$. The caveat is that to utilize such a coordinate-wise strategy, Player II must pick a representative for each of Player I's moves, and it is not clear that she may always find a "good" representative for each of Player I's moves, even though the underlying structures are the same index-by-index.

Instead, it seems necessary to identify the resulting asymptotic cone and develop Player II's strategy given this information. In the case of KSTT, they investigated asymptotic cones of symmetric spaces arising from connected semisimple Lie groups. There, one can find a finitely presented subgroup in the form of a uniform lattice, which is quasi-isometric to the symmetric space, and hence has the same asymptotic cone as the symmetric space itself up to homeomorphism. But the other reason to work with asymptotic cones of such symmetric spaces is because they have been shown to be a particular kind of *building*, commonly referred to as a (nondiscrete) affine $\mathbb{R}$-building. Buildings themselves arose to help classify Lie groups, and they are typically understood as simplicial complexes. However, the asymptotic cones of KSTT cannot be equipped with the structure of a simplicial complex, hence the emphasis that some authors place on the "nondiscrete" portion of the terminology.

In this dissertation, we adapt the notion of EF games to test for elementary equivalence in a version of unbounded continuous logic that allows for arbitrarily indexed sequences of free variables. We use this to show the elementary equivalence of suitably nice $\mathbb{R}$-trees, which gives a proper class worth of non-homeomorphic but elementarily equivalent $\mathcal{L}$-structures. Tits identified that one may recover the isomorphism type of an affine $\mathbb{R}$-building[1] using projective valuations associated to a "wall tree" or "panel tree," which are themselves $\mathbb{R}$-trees. Towards a proof that the asymptotic cones of KSTT are elementarily equivalent, we investigate a few properties of nondiscrete affine $\mathbb{R}$-buildings and prove that in nice enough buildings (which generalize the case of the asymptotic cones from KSTT), any two

[1] His terminology, which has verbatim carried over into some English literature, is "système d'appartements": literally "system of apartments."

classes of parallel walls are elementarily equivalent. In particular, as the asymptotic cones are homogeneous and the buildings associated to them are thick, we see that any point is contained in some thick wall, and this suggests that one may be able to develop an EF game strategy "between" classes of parallel walls that demonstrates the elementary equivalence of the asymptotic cones.

Chapter 2

Continuous Model Theory

Continuous model theory is an adaptation of "classical" model theory to a different underlying logic that will allow us, as logicians, to work with important concepts from areas such as geometry and analysis. The basic ideas to transfer from classical model theory to continuous model theory are to swap the atomic formula "$x = y$" for the formula "$d(x, y)$," swap logical connectives (viewed as truth functions) for uniformly continuous functions, and swap the quantifiers $\forall x$ and $\exists y$ for $\sup_x$ and $\inf_y$, although slight variations may be made as we shall see. These changes are natural enough that we are still able to recover many of the key concepts grounding classical model theory. Several different versions of continuous model theory have been developed over the years, and [3] is a good resource as an introduction to *bounded* continuous model theory.

There are usually slight hurdles to working with *unbounded* structures in continuous model theory, namely in terms of quantification. The two fairly standard approaches are to (1) work in a many-sorted structure where sorts are taken to be balls of arbitrarily large radii or (2) only allow for bounded quantifiers. The version of continuous model theory that we work with in this dissertation was developed in [29]. Luther generally takes the second approach but adapts it to work with quantifiers that represent the *average* sup and inf, in a precise sense. The metric structures themselves will be pointed metric spaces, and as we shall typically be interested in the case of groups viewed as metric structures, this shall not impose too much structure, as the identity element is a clear choice of base point.

2.1 Syntactics

As we lay out Luther's version of continuous model theory, we take the approach of presenting the syntax first. This is simply a matter of personal pedagogical taste, as the author found it intuitive to learn first about purely symbolic logic, which one may then apply to various philosophical and mathematical settings. The main downside to presenting the syntactics first is that some of our choices may seem odd, and we must delay their explanations until after we have seen metric structures.

We first begin with the notion of *controlled functions*.

Definition 2.1. *Let I be an arbitrary indexing set, $(M_n, d_n, \star_n)_{n \in I}$ and $(M, d, \star)$ be pointed metric spaces. We say that a function $f : \prod_{n \in I} M_n \to M$ is* **controlled** *if there exist functions*

- $\lambda : \mathbb{R}_+^I \to \mathbb{R}_+$,
- $N : \mathbb{R}_+^I \times \mathbb{R}_+ \to \mathcal{P}_\omega(I)$, *and*
- $\delta : \mathbb{R}_+^I \times \mathbb{R}_+ \to \mathbb{R}_+$,

such that, for every $(r_i) \in \mathbb{R}_+^I$, $\epsilon > 0$, and $(x_i), (y_i) \in \prod_{n \in I} M_n$, we have

1. *if $\|x_n\| < r_n$ for all $n \in N((r_i), \epsilon)$, then $\|f((x_i))\| \leqslant \lambda((r_i))$, and*

2. if $\|x_n\|, \|y_n\| < r_n$ and $d(x_n, y_n) < \delta((r_i), \epsilon)$ for all $n \in N((r_i), \epsilon)$, then we have $d(f((x_i)), f((y_i))) \leqslant \epsilon$.

If f is a controlled function, we will say that it is **controlled by** (λ, N, δ), and we will refer to (λ, N, δ) as **controllers**. Intuitively, the controllers tell us that, if we want our function to be bounded and uniformly continuous on bounded inputs, then we need only look at finitely many of the coordinates. The controller N specifies exactly which coordinates we must look at, λ gives us our uniform bound on the function, and δ gives us the usual δ necessary for uniform continuity, each of these (except λ) dependent on the bounds (r_i) and $\epsilon > 0$ that we care about.

In fact, this intuitive explanation of controlled functions is captured in the following proposition.

Proposition 2.1. *Suppose $f : \prod_{i=0}^{n} M_i \to M$ is a function. Then f is controlled if and only if the restriction of f to any bounded set is bounded and uniformly continuous.*

Proof. The proof is straightforward. For details, see [29]. □

Remark 2.1. *One may reasonably wonder why we are defining controlled functions to be defined on an arbitrary product of spaces. Those who have seen continuous model theory may (rightfully) fear that controlled functions will become our* class *of logical connectives, and this is indeed the case. In principle, this is to allow ourselves to write down formulas indexed by sets such as $\mathbb{R}$, whose zero sets may be geodesics.*

The follow-up criticism would then ask why we do not simply define controlled functions on products whose indexing sets are no larger than $|\mathbb{R}|$, which avoids our logical connectives being a proper class in size. This is a valid criticism, and we remark that this is likely a more reasonable approach to take if one hoped to develop a proof theory for this flavor of continuous logic. However, our approach is, in general, not too dissimilar to the move from

$\mathcal{L}_{\omega,\omega}$ to $\mathcal{L}_{\infty,\omega}$ in First Order Logic, as was pointed out by a keen participant at the 2023 Graduate Student Conference in Logic.

We are now prepared to introduce the notion of a language for our logic. In classical model theory, one must specify the arity of functions and relations; similarly, we shall specify both the arity and the controllers for our functions and relations.

Definition 2.2. *A* **language** $\mathcal{L}$ *is a tuple* $(\mathcal{S}, \mathcal{F}, \mathcal{R}, \mathfrak{a}, \mathrm{dom}, \mathrm{rng}, \mathrm{ctrl})$ *such that*

- $\mathcal{S}, \mathcal{F}, \mathcal{R}$ *are disjoint sets of symbols, called* **sort***,* **function***, and* **relation symbols***, respectively,*
- $\mathfrak{a}$ *is a function assigning to each symbol from* $\mathcal{F} \cup \mathcal{R}$ *an indexing set* I*, which we call the* **arity** *of that symbol, or we might say, if* $\mathfrak{a}(f) = I$*, that* f *is an* I**-ary** *function,*
- dom *is a function assigning to each symbol* $l \in \mathcal{F} \cup \mathcal{R}$ *some sequence* $(S_i)_{i \in \mathfrak{a}(l)}$*, where each* $S_i \in \mathcal{S}$*,*[1]
- rng *is a function assigning to each symbol from* $\mathcal{F}$ *some* $S \in \mathcal{S}$*, and*
- ctrl *is a function assigning to each symbol* $l \in \mathcal{F} \cup \mathcal{R}$ *controllers* (λ, N, δ) *such that the indexing set is* $\mathfrak{a}(l)$*.*

As is somewhat customary in model theory, we choose not to distinguish constant symbols, but instead think of these as the functions of arity $\varnothing$.[2] We must also specify controllers for any constant symbol. The controller N can only be the constant $\varnothing$ map, and we may define δ as the constant map with value π (indeed, any choice works here). The more interesting part is λ. Note that in this case, λ is a constant function, and we may view this as giving

[1] The choice of using the letter l is because it is the midpoint between f and r in the standard English alphabet.

[2] It is more common to write "arity 0," but in our case, the arity of a function is an indexing set rather than a natural number. Regardless, the domain of any $\varnothing$-ary function will still be the empty product, which has a unique element, and so a $\varnothing$-ary function has as its range a unique element of its codomain.

us an upper bound on how far the interpretation of our constant symbol may lie from the base point of a pointed metric space.

Next, we build up our class of $\mathcal{L}$-terms. These are, intuitively, the formal expressions which *could* refer to an element of a structure, after interpreting any free variables. We assume that for each sort $S \in \mathcal{S}$, we have a proper class worth of variables, say, indexed by Ord, and a $\varnothing$-ary symbol $\star_S$. Usually, we will not write the sort to which each variable x is associated, and it should be understood from context.

Definition 2.3. *Let $\mathcal{L}$ be a language. Then we define the class of* **$\mathcal{L}$-terms** *inductively:*

- *For each $S \in \mathcal{S}$, every variable x of sort S and $\star_S$ are $\mathcal{L}$-terms.*
- *If $f \in \mathcal{F}$ is an I-ary function symbol and $(t_i)_{i \in I}$ are $\mathcal{L}$-terms such that $t_i \in \mathrm{dom}(f)_i$, then $f((t_i)_{i \in I})$ is a $\mathcal{L}$-term of sort $\mathrm{rng}(f)$.*[3]
- *All $\mathcal{L}$-terms are formed by finitely many applications of the above two steps.*

Definition 2.4. *Let $\mathcal{L}$ be a language. Then we define the class of* **$\mathcal{L}$-formulae** *inductively:*

- *If $R \in \mathcal{R}$ is an I-ary relation symbol and $(t_i)_{i \in I}$ are $\mathcal{L}$-terms such that $t_i \in \mathrm{dom}(R)_i$, then $R((t_i)_{i \in I})$ is an $\mathcal{L}$-formula. Likewise, if t_1, t_2 are $\mathcal{L}$-terms of a common sort S, then $d_S(t_1, t_2)$ is an $\mathcal{L}$-formula. Formulae of this kind are the* **atomic $\mathcal{L}$-formulae**.
- *If $(\varphi_i)_{i \in I}$ are $\mathcal{L}$-formulae and $c : \mathbb{R}^I \to \mathbb{R}$ is a controlled function, then $c((\varphi_i)_{i \in I})$ is an $\mathcal{L}$-formula.*
- *If φ is an $\mathcal{L}$-formula and $\bar{x}$ is a finite tuple of variables of length n in φ which have not been used in some previous application of this step, and if $\bar{r}, \bar{s} \in \mathbb{R}^n$ with each $r_i < s_i$, then*

$$\sup_{\bar{x}} \Big|_{\bar{r}}^{\bar{s}} \varphi \qquad \textit{and} \qquad \inf_{\bar{x}} \Big|_{\bar{r}}^{\bar{s}} \varphi$$

[3]N.B.: This includes the case of all constant symbols.

are $\mathcal{L}$-formulae.

- *All $\mathcal{L}$-formulae are constructed by finitely many applications of the above steps.*
- *For an $\mathcal{L}$-formula ψ, we consider each of the atomic $\mathcal{L}$-formulae used, any of the φ_i from an application of the second step, and any of the φ from an application of the third step to be the* **subformulae** *of ψ.*

Definition 2.5. *In a $\mathcal{L}$-formula φ, if x_i appears in some application of the $\sup_{\bar{x}}|_{\bar{r}}^{\bar{s}}$ or $\inf_{\bar{x}}|_{\bar{r}}^{\bar{s}}$, then x_i is said to be a* **bound variable**, *otherwise it is said to be a* **free variable**. *An $\mathcal{L}$-formula without free variables is called an* **$\mathcal{L}$-sentence**.

Remark 2.2. *By convention, if x_j is bound in some formula φ_k and $c((\varphi_i)_{i\in I})$ is a $\mathcal{L}$-formula where $k \in I$, then x_j is assumed not to be free in φ_i for* any *$i \in I$. This is possible without loss of generality by simply changing any instances of x_j occurring in the other formulae to a new variable that does not occur in any of the φ_i's. We note that with a proper class of free variables, we are always able to do so.*

Definition 2.6. *We say an $\mathcal{L}$-formula is* **quantifier free** *(abbreviated q.f.) if it is constructed without using $\sup_{\bar{x}}|_{\bar{r}}^{\bar{s}}$ or $\inf_{\bar{x}}|_{\bar{r}}^{\bar{s}}$.*

It will be occasionally helpful to have the notion of *quantifier depth* of an $\mathcal{L}$-formula φ. One may expect this to be the number of times one has utilized the sup or inf step to construct φ. However, we must define this as a class function, as a formula may have an arbitrary number of subformulae.

Definition 2.7. *We define the class function* dep *from the class of all $\mathcal{L}$-formulae to the class* Ord. *We shall say that* $\mathrm{dep}(\varphi)$ *is the* **quantifier depth of** *φ.*

1. $\mathrm{dep}(\varphi(\bar{x})) := 0$ *if φ is an atomic $\mathcal{L}$-formula.*

2. *If* $f : \mathbb{R}^I \to \mathbb{R}$ *is a controlled function and* $(\varphi_i)_{i \in I}$ *are* $\mathcal{L}$*-formulae, then*

$$\operatorname{dep}\left(f((\varphi_i)_{i \in I})\right) := \sup_{i \in I}\left(\operatorname{dep}\left(\varphi_i\right)\right).$$

3. *If* φ *is an* $\mathcal{L}$*-formula with a finite tuple* $\bar{x}$ *of free variables, and appropriate* $\bar{r}, \bar{s}$*, then*

$$\operatorname{dep}\left(\sup_{\bar{x}} \Big|_{\bar{r}}^{\bar{s}} \varphi\right) := \operatorname{dep}(\varphi) + 1.$$

4. *The* inf *case is defined similarly.*

We record a few easy consequences of this definition:

Lemma 2.1. 1. *If* φ *is a subformula of* ψ*, then* $\operatorname{dep}(\varphi) \leqslant \operatorname{dep}(\psi)$.

2. *For each* $\gamma \in \mathrm{Ord}$*, there is an* $\mathcal{L}$*-sentence of depth* γ*.*

Proof. 1. Obvious.

2. We can prove this by transfinite induction on ordinals.

   For $\gamma = 0$, the sentence $d(\star_S, \star_S)$ suffices.

   Let $\gamma = \alpha + 1$ be a successor ordinal, and suppose that we have an $\mathcal{L}$-sentence φ of depth α. Pick a variable x which does not appear in φ. Then $\sup_x \big|_0^1 (\varphi + d(\star_S, x))$ is an $\mathcal{L}$-sentence of depth $\alpha + 1 = \gamma$.

   Suppose that γ is a limit ordinal, and for each $\xi < \gamma$, there is an $\mathcal{L}$-sentence φ_ξ of depth ξ. Let $c : \mathbb{R}^\gamma :\to \mathbb{R}$ be any controlled function. Then $c((\varphi_\xi)_{\xi < \gamma})$ is clearly an $\mathcal{L}$-sentence of depth γ. □

This demonstrates that quantifier depth is indeed a function onto all of Ord, and further, once we define the notion of an $\mathcal{L}$-*theory*, this will imply that the theory of an $\mathcal{L}$-structure is indeed a proper class. We will not comment much further on the issue of proper classes.

We do note one important feature of $\mathcal{L}$-sentences though, which we will need when discussing EF games. Namely, we have the following:

Proposition 2.2. *If φ is an $\mathcal{L}$-sentence, then any ψ which is a subformula of φ must have* at most *finitely many free variables.*

Proof. This follows from the fact that formulae are formed using finitely many applications of controlled functions and applications of quantifiers. If ψ had infinitely many free variables, say $(x_1, x_2, \dots)$, then in finitely many steps, one is able to quantify over only finitely many of these variables. □

2.2 Semantics

We now introduce the semantics for our continuous logic. We present the semantics for a (potentially) multi-sorted structure, though for most of the dissertation, we shall only care about single-sorted structures. We begin with the notion of an $\mathcal{L}$-*structure*.

Definition 2.8. *Let $\mathcal{L}$ be a language. We say that an $\mathcal{L}$-**structure** $\mathcal{M}$ is given by the following data:*

- *a complete, pointed metric space $(M_S, d_S, \star_S)$ for each $S \in \mathcal{S}$,*
- *for each $f \in \mathcal{F}$ with $\mathfrak{a}(f) = I$, $\mathrm{dom}(f) = (S_i)_{i \in I}$, and $\mathrm{rng}(f) = S$, a controlled function*

$$f^{\mathcal{M}} : \prod_{i \in I} M_{S_i} \to M_S$$

  whose controllers are given by $\mathrm{ctrl}(f)$, and

- *for each* $R \in \mathcal{R}$ *with* $\mathfrak{a} = I(R)$ *and* $\mathrm{dom}(R) = (S_i)_{i\in I}$*, a controlled function*

$$R^{\mathcal{M}} : \prod_{i\in I} M_{S_i} \to \mathbb{R}_{\geqslant 0}$$

  whose controllers are given by $\mathrm{ctrl}(R)$*.*

A Comment on Philosophy

Before we move onto interpretations and satisfaction of $\mathcal{L}$-formulae, we comment briefly on the philosophical relationship between Continuous Logic and First Order Logic. This comment will be relatively informal in nature. Many authors have written about how a "classical" model theorist ought to think about and approach continuous model theory for the first time. We offer the following perspective, which is not novel, but we include it as a pedagogical framework for any learners who may find it odd to think of logical connectives as controlled functions.

In standard texts on logic for philosophy (e.g. [35]), one may take the approach of defining "truth" for Propositional Logic in terms of *valuations*. Intuitively, one begins with the concept of an *interpretation* $\mathscr{I}$, which maps from the set of propositions $\{P_1, P_2, \dots\}$ to the set $\{0, 1\}$. This then induces a *valuation* $V_{\mathscr{I}}$ which maps from the set of "well-formed formulas" to the set $\{0, 1\}$. We note a few important observations:

- Any philosophical notion of "truth" or "falsity" is something that we, as humans, ascribe to the values of 0 and 1. The definitions involved in Propositional Logic are only meant to formalize the approach we take to logic. Similarly, the symbols $\neg$ and $\to$ hold no inherent meaning, but the rules by which we define our valuation $V_{\mathscr{I}}$ suggest to us that these should be treated as "not" and "implies."

- The connectives themselves can be viewed as "truth functions," which is really a philosophical notion, but again may be formalized as functions $f : \{0,1\}^n \to \{0,1\}$. In this way, it's natural to think of formulae from Propositional Logic not as strings of symbols with any inherent "truth" or "falsity" attached to them, but rather as functions in their own right. Sider makes the definition that φ *symbolizes* a truth function $f : \{0,1\}^n \to \{0,1\}$ if φ consists of $P_1, \ldots, P_n$ and, for every interpretation $\mathscr{I}$, $V_{\mathscr{I}}(\varphi) = f(\mathscr{I}(P_1), \ldots, \mathscr{I}(P_n))$. It is possible to show that any truth function may be symbolized in Propositional Logic.

- Applying the perspective of valuations to First Order Logic, one must define how to evaluate formulae involving the quantifier $\forall x$. This is handled by introducing the notion of a variable assignment (which one should think of as a function $g : \{x_1, x_2, \ldots\} \to M$ where M is some fixed structure) and saying that $V_{M,g}(\forall x \varphi) = 1$ iff, for any $m \in M$ and g such that $g(x) = m$, we have $V_{M,g}(\varphi) = 1$. In this fashion, formulae may still be viewed as functions, and interpretations are now values within a model rather than just 0 or 1.

Continuous Logic is not much different. Our formulae are still functions, although they now take on values in $\mathbb{R}_{\geqslant 0}$. Given the above point of view, perhaps one should expect that our logical connectives should consist of (or should be able to express) *all* functions $f : \mathbb{R}^n_{\geqslant 0} \to \mathbb{R}_{\geqslant 0}$. This is far too general to develop a reasonable model theory for continuous logic, hence our choice to use controlled functions.

The other key difference (which we shall formalize shortly) is that our quantifiers are now going to be *average supremum and infimum*. While this is technically distinct from the quantifiers we see in First Order Logic, we shall see that they are not too dissimilar, and intuitively play much nicer with controlled functions on the reals.

But perhaps the most serious distinction of Continuous Logic is that we shall say that a structure satisfies a formula when its "valuation" is 0, rather than 1. In some sense, this is not too grave a sin, as the atomic formula "$x = y$" should be true in First Order Logic if and only if x and y have the same interpretation, and in any metric structure, x and y have the same interpretation if and only if the formula "$d(x, y)$" evaluates to 0.

With this in mind, we formalize the interpretations of $\mathcal{L}$-terms and $\mathcal{L}$-formulae relative to an $\mathcal{L}$-structure $\mathcal{M}$.

Definition 2.9. *Let $\mathcal{L}$ be a language and $\mathcal{M}$ an $\mathcal{L}$-structure. For each $\mathcal{L}$-term t, we define its* **interpretation** $t^{\mathcal{M}}$ *inductively as follows:*

- *If t is a variable x of sort S, then $t^{\mathcal{M}} : M_S \to M_S$ is the identity function.*
- *If t is the symbol $\star_S$ for some sort S, then $t^{\mathcal{M}} : M_S^{\varnothing} \to M_S$ is the $\varnothing$-ary function that maps the unique point of $M_S^{\varnothing}$ to the base point $\star_S \in M_S$.*
- *If f is a function symbol with $\mathrm{dom}(f) = (S_i)_{i\in I}$ and (t_i) are $\mathcal{L}$-terms such that, for each $i \in I$, t_i is an $\mathcal{L}$-term of sort S_i and $t_i^{\mathcal{M}}$ has already been defined, then*

$$[f((t_i)_{i\in I})]^{\mathcal{M}} = f^{\mathcal{M}}((t_i^{\mathcal{M}})_{i\in I}).$$

Definition 2.10. *Let $\mathcal{L}$ be a language, $\mathcal{M}$ an $\mathcal{L}$-structure, and $\varphi((x_i)_{i\in I})$ an $\mathcal{L}$-formula. We shall inductively define $\varphi^{\mathcal{M}}((x_i)_{i\in I})$, the* **interpretation** *of φ in $\mathcal{M}$:*

- *If $\varphi((x_i)_{i\in I}) = d_S(t_1, t_2)$ where t_i are $\mathcal{L}$-terms of sort S, then $\varphi^{\mathcal{M}}((x_i)_{i\in I}) = d_S^{\mathcal{M}}(t_1^{\mathcal{M}}, t_2^{\mathcal{M}})$.*[4]
- *If $\varphi((x_i)_{i\in I}) = R((t_i)_{i\in I})$ with $\mathrm{dom}(R) = (S_i)_{i\in I}$ and each t_i is an $\mathcal{L}$-term of sort S_i, then*

$$\varphi^{\mathcal{M}}((x_i)_{i\in I}) = R^{\mathcal{M}}((t_i^{\mathcal{M}})_{i\in I}).$$

[4]It is worth noting that the $\mathcal{L}$-term t_1 need not consist of all variables in I; however, if I_1 and I_2 are the variables appearing in the terms t_1 and t_2, then we must have $I_1 \cup I_2 = I$.

- *If* $\varphi((x_i)_{i\in I}) = c((\psi_i)_{i\in I})$ *for some controlled function* $c : \mathbb{R}^I_{\geqslant 0} \to \mathbb{R}_{\geqslant 0}$ *and each of the* $\psi_i^{\mathcal{M}}$ *have been defined, then*

$$\varphi^{\mathcal{M}}((x_i)_{i\in I}) = u((\psi_i^{\mathcal{M}})_{i\in I}).$$

- *If* $\varphi((x_i)_{i\in I})$ *is of the form* $\sup_{\bar{y}} \big|_{\bar{r}}^{\bar{s}} \psi((x_i)_{i\in I})$ *where* $\bar{y} = (y_1, \ldots, y_n)$*, each* y_k *is a variable of sort* S_k*,* $\bar{r}, \bar{s} \in \mathbb{R}^n_{\geqslant 0}$*, and* $\psi^{\mathcal{M}}$ *has been defined, then*

$$\varphi^{\mathcal{M}}((x_i)_{i\in I}) := \left(\prod_{k=1}^{n} \frac{1}{s_k - r_k}\right) \int_{r_n}^{s_n} \cdots \int_{r_1}^{s_1} \sup_{\bar{y}\in B_{\rho_1,\ldots,\rho_n}(\mathcal{M})} \psi^{\mathcal{M}}((x_i)_{i\in I})\, \mathrm{d}\rho_1 \cdots \mathrm{d}\rho_n$$

where $M_S = \prod_{k=1}^n M_{S_k}$ *and* $B_{\rho_1,\ldots,\rho_n}(\mathcal{M})$ *denotes* $\prod_{k=1}^n B_{\rho_k}(\star_{S_k})$*, i.e. the projection onto the* k^{th} *coordinate is simply the open ball of radius* ρ_k *in* M_{S_k} *centered at* $\star_{S_k}$*.*

- *The case of* $\inf_{\bar{y}} \big|_{\bar{r}}^{\bar{s}} \psi((x_i)_{i\in I})$ *is handled similarly.*

Even to those fairly familiar with continuous model theory, our quantifiers may seem a bit strange. It is in general possible to take the approach of using bounded quantifiers to allow for bona fide unbounded metric structures. The approach we take here is not only to bound our quantifiers, but to take the *average* as we vary the bound $\bar{\rho}$ over some bounded interval; this is beneficial for ensuring that the quantifiers are continuous with respect to changing bounds $\bar{r}$ and $\bar{s}$.

There's an obvious concern about integrability in this definition. Note that we are indeed working with a function of the form

$$\Psi(\bar{\rho}, (x_i)_{i\in I}) := \sup_{\bar{y}\in B_{\rho_1,\ldots,\rho_n}(M_S)} \psi^{\mathcal{M}}((x_i)_{i\in I})$$

where $\Psi : \mathbb{R}^n_{\geqslant 0} \times \prod_{i\in I} M_{S_i} \to \mathbb{R}_{\geqslant 0}$. That is, we are really integrating Ψ over the variables ρ_k which live in $\mathbb{R}$, and it turns out that this function is indeed Riemann integrable in a precise way. Namely, we can really only discuss Riemann integrability of the function Ψ after fixing $(x_i) = (b_i)$ relative to some model $\mathcal{M}$. However, the next proposition demonstrates that, given $\epsilon > 0$ and a bound on the (b_i), one can find a uniform partition of $[r_1, s_1] \times \cdots \times [r_n, s_n]$ such that for *any* choice of $\mathcal{L}$-structure $\mathcal{M}$ and (b_i) from $\mathcal{M}$, the Riemann sum is within ϵ of the actual integral. In this way, the Riemann integrability of the interpretation of quantifiers is really a feature of the logic itself and is not dependent on any semantic choice.

Proposition 2.3. *Let I be an indexing set, y a variable of sort S_0, $(x_i)_{i\in I}$ a sequence of variables where each x_i is a variable of sort S_i, $\psi(y, (x_i)_{i\in I})$ an $\mathcal{L}$-formula, $\mathcal{M}$ an $\mathcal{L}$-structure, and $r, s \in \mathbb{R}_{\geqslant 0}$ with $r < s$. Then, for $(b_i)_{i\in I} \in \prod_{i\in I} M_{S_i}$, we define*

$$s_b^{\mathcal{M}}(\rho) := \sup_{y\in B_\rho(M_{S_0})} \psi^{\mathcal{M}}(y, (b_i))$$

and $s_b^{\mathcal{M}}$ is Riemann integrable.

Further, given any sequence of positive reals $(r_i)_{i\in I}$ and $\epsilon > 0$, there is a finite $N_0 \subseteq I$, real number $\Delta > 0$, and uniform partition $\rho_0 < \cdots < \rho_K$ of $[r, s]$ of mesh Δ such that for any $\mathcal{L}$-structure $\mathcal{M}$, any $(b_i)_{i\in I}$ with $\|b_i\| < r_i$ for each $i \in N_0$, and any $s_0^*, \ldots, s_{K-1}^* \in \mathbb{R}$ satisfying

$$s_b^{\mathcal{M}}(\rho_k) \leqslant s_k \leqslant s_b^{\mathcal{M}}(\rho_{k+1})$$

for all $k < K$, we have

$$\left| \int_r^s s_b^{\mathcal{M}}(\rho)\, \mathrm{d}\rho - \sum_{k<K} s_k^* \Delta \right| < \epsilon.$$

A similar statement is true when we extend to appropriate finite tuples for y and change sup *to* inf.

Proof. This is Lemma 2.7.6 from [29]. The proof relies on the fact that ψ is controlled by some controllers (λ, N, δ) which do not depend on a choice of $\mathcal{M}$. □

Crucial to later discussions will be the fact that $\mathcal{L}$-formulae themselves are controlled functions. Technically, we should say that the *interpretation* of any $\mathcal{L}$-formula $\varphi^{\mathcal{M}}((x_i))$ is a controlled function, but as it turns out, the controllers for φ do not depend on our choice of $\mathcal{L}$-structure! This is captured in Lemma 2.7.4 of [29], which we record here:

Lemma 2.2. *For every $\mathcal{L}$-formula $\varphi((x_i))$, there are controllers (λ, N, δ) such that, for every $\mathcal{L}$-structure $\mathcal{M}$, the function $\varphi^{\mathcal{M}}((x_i))$ is controlled by (λ, N, δ).*

Before we move on to the more model theoretic notions for this version of continuous logic, we shall first introduce some terminology that will be helpful to us.

Definition 2.11. *Let (X, d) be a metric space. The* **character density of X***, denoted* $\text{char}(X)$*, is the smallest cardinality of a dense subset of X.*

In some sense, the character density is the correct analog to cardinality of subsets of universes of $\mathcal{L}$-structures. This will appear in our discussion of saturation.

2.3 Model Theory

We are now prepared to introduce various model theoretic terminology and results. For more details and results on this particular flavor of continuous model theory, one should consult [29].

Definition 2.12. *Let $\mathcal{M}$ be an $\mathcal{L}$-structure, $\varphi((x_i)_{i\in I})$ an $\mathcal{L}$-formula and $(a_i)_{i\in I} \in \mathcal{M}$. Then we say that $(a_i)_{i\in I}$* ***satisfies*** *φ* ***in*** *$\mathcal{M}$, which we denote as $\mathcal{M} \models \varphi((a_i)_{i\in I})$, if*

$$\varphi^{\mathcal{M}}((a_i)_{i\in I}) = 0.$$

Remark 2.3. *This is a slight deviation from the terminology used in [3] in that we do not introduce the notion of $\mathcal{L}$-conditions, but is more consistent with Luther's work. The only* *distinction between this and Luther's work is that he only defines satisfaction of $\mathcal{L}$-sentences, which is not much of an issue, as we briefly explain below.*

Definition 2.13. *Fix a language $\mathcal{L}$, $\mathcal{M}$ an $\mathcal{L}$-structure, and some set $A \subseteq \mathcal{M}$ taken in the sense that A consists of elements of (potentially) multiple sorts of $\mathcal{M}$. We say that $\mathcal{L}_A := \mathcal{L} \cup A$ is an* **expansion of $\mathcal{L}$ by constants** *if we view each $a \in A$ as a constant symbol corresponding to its respective sort of $\mathcal{M}$ with controller λ defined to the constant function with value $d^{\mathcal{M}}(\star_S, a)$ when a belongs to sort S.*[5]

Of course, if $\mathcal{M}$ is an $\mathcal{L}$-structure and $A \subseteq \mathcal{M}$, then there's an obvious way to create an $\mathcal{L}_A$-structure $\mathcal{M}_A$ which *extends* $\mathcal{M}$. In this way, Luther could easily recover the notion of (a_i) satisfying some $\mathcal{L}$-formula $\varphi((x_i))$ in $\mathcal{M}$ simply by requiring that, if one views $\varphi((a_i))$ as an $\mathcal{L}_A$-sentence, then $\varphi((a_i))$ is satisfied in $\mathcal{M}_A$.

Just as one may define First Order Logic using only the connectives $\neg$, $\rightarrow$, and $\forall$ while treating other common connectives as shorthand, we shall introduce some common shorthand below.

Definition 2.14. *Let φ, ψ be $\mathcal{L}$-formulae and $r \in \mathbb{R}_{\geqslant 0}$. Then we make the following conventions:*

- "$\varphi \dot{=} \psi$" $:= |\varphi - \psi|$,
- "$\varphi \dot{\leqslant} \psi$" $:= \max(0, \varphi - \psi)$,
- "$\varphi \dotdiv \psi$" $:= \max(0, \varphi - \psi)$, *and*[6]
- "r" *may be regarded as shorthand for the $\mathcal{L}$-formula given by the $\varnothing$-ary relation R_r which maps to the constant r.*

[5] It is good to remember that we are viewing constants as $\varnothing$-ary function symbols and so must specify controllers. This was discussed just after Definition 2.2.

[6] N.B.: The symbol $\dotdiv$ is *not* typeset using `\dot{-}`, as this produces $\dot{-}$, and the dot is clearly too high. Instead, one must import the package `mathabx` and use the command `\dotdiv`.

It may seem strange to see that $\varphi\dot{\leqslant}\psi$ and $\varphi \dot{-} \psi$ are both shorthand for the same $\mathcal{L}$-formula. The reason here is a psychological one; for example, we shall later define the *Gromov product* in a metric space, and while $\dot{\leqslant}$ suffices, it is easier to convince oneself that our formalization of this product is correct when written with $\dot{-}$. One should think of this logical connective as replacing usual subtraction, but as ordinary subtraction is not a controlled function to $\mathbb{R}_{\geqslant 0}$, we must "cut off" the subtraction before it drops into the negative reals.

We justify our choice of shorthand in the first two cases with the following easy fact:

Fact 2.1. *Let $\mathcal{M}$ be an $\mathcal{L}$-structure and φ, ψ $\mathcal{L}$-sentences. Then,*

- $\mathcal{M} \models \varphi\dot{=}\psi$ *if and only if* $\varphi^{\mathcal{M}} = \psi^{\mathcal{M}}$, *and*
- $\mathcal{M} \models \varphi\dot{\leqslant}\psi$ *if and only if* $\varphi^{\mathcal{M}} \leqslant \psi^{\mathcal{M}}$.

We also note that continuous logic is a *positive* logic, meaning that we do not have a connective which behaves like negation. This is why we have introduced shorthand for $\dot{\leqslant}$, but we do not have a shorthand of $\dot{<}$. However, we do have conjunction and disjunction. This is captured in the following:

Proposition 2.4. *Let $\mathcal{M}$ be an $\mathcal{L}$-structure and φ, ψ be $\mathcal{L}$-sentences. Then,*

- $\mathcal{M} \models \varphi + \psi$ *if and only if* $\mathcal{M} \models \varphi$ *and* $\mathcal{M} \models \psi$, *and*
- $\mathcal{M} \models \varphi \cdot \psi$ *if and only if* $\mathcal{M} \models \varphi$ *or* $\mathcal{M} \models \psi$.

Even further, we have countable conjunction! We note that the sum of two $\mathcal{L}$-formulae will be zero if and only if each interpretation is zero, hence addition behaves as conjunction; similarly, multiplication behaves as disjunction. We also have the following proposition:

Proposition 2.5. *Let $(\varphi_n((x_i))_{n\in\mathbb{N}}$ be a sequence of $\mathcal{L}$-formulae in free variables (x_i). Then there is an $\mathcal{L}$-formula which we denote as $\bigwedge_{n\in\mathbb{N}} \varphi_n((x_i))$ such that, for any $\mathcal{L}$-structure $\mathcal{M}$*

and any $(a_i) \in \mathcal{M}$, *we have that*

$$\mathcal{M} \models \bigwedge_{n\in\mathbb{N}} \varphi_n((a_i)) \iff \mathcal{M} \models \varphi_n((a_i)) \text{ for each } n \in \mathbb{N}.$$

Proof. We claim that $c((x_n)_{n\in\mathbb{N}}) := \sum_{n=0}^{\infty} 2^{-n} x_n/(1+x_n)$ is a controlled function, hence we may take $\bigwedge_{n\in\mathbb{N}}$ to be shorthand for this connective. We find controllers for c. Note that the image of c is bounded by 2, and for $\epsilon > 0$, one can pick N and $\delta > 0$ by letting N be large enough such that $2^{-N} < \epsilon/3$ and δ so that

$$\sum_{n=1}^{N} \frac{\delta}{2^n} < \epsilon/3.$$

Then, if $(y_n)_{n\in\mathbb{N}_+}$ is such that $|x_n - y_n| < \delta$ for each $n = 1, \ldots, N$, we have

$$\begin{aligned}
&\left| \sum_{n=1}^{\infty} \frac{1}{2^n} \cdot \frac{x_n}{1+x_n} - \sum_{n=1}^{\infty} \frac{1}{2^n} \cdot \frac{y_n}{1+y_n} \right| \\
\leqslant &\left| \sum_{n=1}^{N} \frac{1}{2^n} \left(\frac{x_n}{1+x_n} - \frac{y_n}{1+y_n} \right) \right| + \left| \sum_{n=N+1}^{\infty} \frac{1}{2^n} \cdot \frac{x_n}{1+x_n} \right| + \left| \sum_{n=N+1}^{\infty} \frac{1}{2^n} \cdot \frac{y_n}{1+y_n} \right| \\
\leqslant &\left| \sum_{n=1}^{N} \frac{1}{2^n} \cdot \frac{x_n - y_n}{(1+x_n)(1+y_n)} \right| + \left| \sum_{n=N+1}^{\infty} \frac{1}{2^n} \right| + \left| \sum_{n=N+1}^{\infty} \frac{1}{2^n} \right| \\
\leqslant &\sum_{n=1}^{N} \frac{1}{2^n} |x_n - y_n| + \left| \sum_{n=N+1}^{\infty} \frac{1}{2^n} \right| + \left| \sum_{n=N+1}^{\infty} \frac{1}{2^n} \right| \\
< &\epsilon/3 + \epsilon/3 + \epsilon/3 = \epsilon.
\end{aligned}$$

Thus, it follows that our shorthand is valid. It should be clear that $\mathcal{M} \models \bigwedge_{n\in\mathbb{N}} \varphi_n((a_i))$ if and only if $\mathcal{M} \models \varphi_n((a_i))$ for each $n \in \mathbb{N}$. This follows from noting that $\mathcal{M} \models \bigwedge_{n\in\mathbb{N}} \varphi_n((a_i))$ if and only if the interpretation of each summand is 0. ◻

As mentioned, the quantifiers in this setting are not quite the same as the quantifiers $\forall$ and $\exists$ one encounters in classical model theory. However, they are not too far off, as we see in the following proposition.

Proposition 2.6. *Let $\varphi(\bar{x};(y_i)_{i\in I})$ be an $\mathcal{L}$-formula, $\mathcal{M}$ an $\mathcal{L}$-structure, and $(a_i)_{i\in I} \in \mathcal{M}$.*

1. *$\mathcal{M} \models \sup_{\bar{x}} \big|_{\bar{r}}^{\bar{s}} \varphi(\bar{x};(a_i)_{i\in I})$ if and only if $\mathcal{M} \models \varphi(\bar{b};(a_i)_{i\in I})$ for all $\bar{b} \in B_{\bar{s}}(\mathcal{M})$.*

2. *$\mathcal{M} \models \inf_{\bar{x}} \big|_{\bar{r}}^{\bar{s}} \varphi(\bar{x};(a_i)_{i\in I})$ if and only if for each $\epsilon > 0$, there is $\bar{b} \in B_{\bar{r}}(\mathcal{M})$ such that $\mathcal{M} \models \varphi(\bar{b};(a_i)_{i\in I}) \dot{\leqslant} \epsilon$.*

Proof. 1. This follows from the fact that $\mathcal{L}$-formulae take on values in $\mathbb{R}_{\geqslant 0}$. Thus, the average sup of $\varphi(\bar{x};(a_i)_{i\in I})$ is 0 in $\mathcal{M}$ if and only if $\sup\limits_{\bar{x}\in B_{\bar{s}}} \varphi(\bar{x};(a_i)_{i\in I})$ is 0 in $\mathcal{M}$, which happens if and only if $\varphi(\bar{b};(a_i)_{i\in I})$ is 0 for all $\bar{b} \in B_{\bar{s}}(\mathcal{M})$.

2. ($\Rightarrow$): Suppose that $\mathcal{M} \models \inf_{\bar{x}} \big|_{\bar{s}}^{\bar{s}} \varphi(\bar{x};(a_i)_{i\in I})$. Again, the average inf being 0 implies that it is 0 for each $\bar{r}'$ where $\bar{r} \leqslant \bar{r}' \leqslant \bar{s}$. In particular, $\inf\limits_{\bar{x}\in B_{\bar{r}}(\mathcal{M})} \varphi(\bar{x};(a_i)_{i\in I})$ is 0 in $\mathcal{M}$, and we win.

   ($\Leftarrow$): Suppose that, for each $\epsilon > 0$, there is $\bar{b} \in B_{\bar{r}}(\mathcal{M})$ such that $\mathcal{M} \models \varphi(\bar{b};(a_i)_{i\in I}) \leqslant \epsilon$. Then in particular, we see that $\inf\limits_{\bar{x}\in B_{\bar{r}}(\mathcal{M})} \varphi(\bar{x};(a_i)_{i\in I}) = 0$, thus $\inf\limits_{\bar{x}\in B_{\bar{r}'}(\mathcal{M})} \varphi(\bar{x};(a_i)_{i\in I}) = 0$ for any $\bar{r}' \geqslant \bar{r}$, hence $\mathcal{M} \models \inf_{\bar{x}} \big|_{\bar{r}}^{\bar{s}} \varphi(\bar{x};(a_i)_{i\in I})$. □

Remark 2.4. *The previous proposition tells us, in some sense, that $\sup_x \big|_r^s$ acts like $\forall$ (with the main caveat being that we are only able to quantifier over bounded subsets), and that $\inf_x \big|_r^s$ acts only* approximately *like $\exists$. That is to say, although for each $\epsilon > 0$ we may find some $\bar{b}_\epsilon$ such that $\varphi^{\mathcal{M}}(\bar{b}_\epsilon;(a_i)) \leqslant \epsilon$, this is not enough to guarantee the existence of an element $\bar{b}_0$ such that $\varphi^{\mathcal{M}}(\bar{b}_0;(a_i)) = 0$. Of course, if $\mathcal{M}$ were compact, then this would be enough to guarantee the existence of such a $\bar{b}_0$. However, in the case that $\mathcal{M}$ is not compact, one may find relatively simple examples where such a $\bar{b}_0$ does not exist.*

The next is an example of a bounded, (necessarily) noncompact $\mathcal{L}$-structure in a language with only one function symbol that demonstrates how inf fails to act as $\exists$.

Example 2.1. *Let* (X, d_p) *be the closed unit ball in* $\ell^p(\mathbb{R})$ *with* $p > 0$,[7] *and let* $\mathcal{L} := \{f\}$ *where* f *is a unary function symbol.*[8] *We define an* $\mathcal{L}$*-structure* $\mathcal{M}$ *as follows:*

- *Our complete, pointed metric space will be* $(X, d_p, 0)$ *where* 0 *will be interpreted as the constant* 0 *sequence.*
- *We interpret* f *by* $f^{\mathcal{M}}(a_0, a_1, \ldots) := (0, a_0, a_1, \ldots)$.

To see why f *is controlled, we note that by Proposition 2.1, it suffices to show that* f *is bounded and uniformly continuous. Of course, as* f *maps into a bounded set, it is clearly bounded, and uniform continuity follows from the fact that* f *is an isometry. Now, let* $\varphi(x) := |d(\star, x) - 1| + d(x, f(x))$.[9] *I claim that* $\mathcal{M} \models \inf_x \big|_1^2 \varphi(x)$. *Indeed, for each* $\epsilon > 0$, *we may select* $n \in \mathbb{N}$ *such that* $(2/n)^{1/p} < \epsilon$, *and let* $a_n := (1/\sqrt[p]{n}, 1/\sqrt[p]{n}, \ldots, 1/\sqrt[p]{n}, 0, 0, \ldots)$ *where there are* n *instances of* $1/n$. *By construction,* $d^{\mathcal{M}}(\star, a_n) = 1$, *and we see that*

$$d^{\mathcal{M}}(a_n, f(a_n)) = \left(\left(\frac{1}{\sqrt[p]{n}}\right)^p + \left(\frac{1}{\sqrt[p]{n}}\right)^p\right)^{1/p} = \left(\frac{2}{n}\right)^{1/p} < \epsilon.$$

Thus, we see that $\varphi^{\mathcal{M}}(a_n) < \epsilon$, *and thus* $\mathcal{M} \models \inf_x \big|_1^2 \varphi(x)$. *However, suppose that there were* $a = (a_n)_{n \in \mathbb{N}} \in \mathcal{M}$ *such that* $\mathcal{M} \models \varphi(a)$. *Then in particular,* $d_p(a, f(a)) = 0$, *hence for each* $n \in \mathbb{N}$, *it must be the case that* $a_n = a_{n+1}$, *and in particular,* $a_0 = 0$, *hence* $a = (0, 0, \ldots)$. *However,* $\varphi^{\mathcal{M}}(0) = 1$, *a contradiction.*

It is easy to check that $\mathcal{M}$ in the previous example is not compact. As we have built up a continuous model theory for unbounded metric structures, we shall typically be interested in

[7]Recall that $\ell^p(\mathbb{R})$ is defined to be the set of sequences $(x_n)_{n \in \mathbb{N}}$ such that $\sum_{n=0}^{\infty} |x_n|^p < \infty$, with norm given by $\|x_n\|_p := \left(\sum_{n=0}^{\infty} |x_n|^p\right)^{1/p}$.

[8]We technically should also specify the controllers if we are including f in our language, but we forgo this as we are only working with this specific language once. We will explain why the interpretation of this symbol is controlled, and so one can imagine that we have specified those same controllers within our language.

[9]It is worthwhile to point out that there is no typographical error here, as d is our logical symbol and must be written as such in any $\mathcal{L}$-formula. It will be interpreted as d_p in the $\mathcal{L}$-structure we have described.

noncompact spaces. This, however, does not rule out the possibility of $\inf_x |_r^s$ behaving like $\exists$. Later, we shall discuss the notion of *types* and see that, in the case that an $\mathcal{L}$-structure $\mathcal{M}$ is sufficiently "saturated," then $\inf_x |_r^s$ indeed behaves as $\exists$ within $B_r(\mathcal{M})$.

Definition 2.15. *We say that any class*[10] T *of* $\mathcal{L}$*-sentences is an* **$\mathcal{L}$-theory**. *If there is some* $\mathcal{L}$*-structure* $\mathcal{M}$ *such that* $\mathcal{M} \models \varphi$ *for all* $\varphi \in T$*, we shall say that* T *is satisfiable, and that* $\mathcal{M}$ *is a model of* T*.*

Given $\mathcal{L}$*-structure* $\mathcal{M}$*, we say that the* **theory of** $\mathcal{M}$*, denoted* $\mathrm{Th}(\mathcal{M})$*, is the class of all* $\mathcal{L}$*-sentences* φ *such that* $\mathcal{M} \models \varphi$*.*

Definition 2.16. *Given a class* $\mathcal{K}$ *of* $\mathcal{L}$*-structures, we shall say that* $\mathcal{K}$ *is* **$\mathcal{L}$-axiomatizable** *if there is an* $\mathcal{L}$*-theory* T *such that*

$$\mathcal{K} = \{\mathcal{M} \text{ an } \mathcal{L}\text{-structure} : \mathcal{M} \models T\}.$$

Definition 2.17. *Let* $\mathcal{L}$ *be a language and* $\mathcal{M}, \mathcal{N}$ *be* $\mathcal{L}$*-structures. We shall say that* $\mathcal{M}$ *and* $\mathcal{N}$ *are* **elementarily equivalent***, denoted* $\mathcal{M} \equiv \mathcal{N}$*, if* $\mathrm{Th}(\mathcal{M}) = \mathrm{Th}(\mathcal{N})$*.*

In classical model theory, one typically defines the notion of an $\mathcal{L}$-embedding from one $\mathcal{L}$-structure to another. This generalizes the notion of, say, a group monomorphism by adding the requirements that each of our function, relation, and constants are preserved in a meaningful way.

We make a similar definition in continuous model theory.

Definition 2.18. *Let* $\mathcal{L}$ *be a language and* $\mathcal{M}, \mathcal{N}$ *be* $\mathcal{L}$*-structures. We say that* $h : \mathcal{M} \hookrightarrow \mathcal{N}$ *is an* **$\mathcal{L}$-embedding** *if* h *is a collection of functions* $h_S : M_S \to N_S$ *such that*

[10] Again, we must be cautious not to say "set."

- *for each distance symbol* d_S *and* $a, b \in M_S$

$$d_S^{\mathcal{N}}(h_S(a), h_S(b)) = d_S^{\mathcal{M}}(a, b),$$

- *for each function symbol* f *(including the* $\varnothing$*-ary functions* $\star_S$*) with* $\mathrm{dom}(f) = (S_i)_{i \in I}$, $\mathrm{rng}(f) = S$, *and* $a_i \in M_{S_i}$

$$f^{\mathcal{N}}((h_{S_i}(a_i))_{i \in I}) = h_S(f^{\mathcal{M}}((a_i)_{i \in I})), and$$

- *for each relation symbol* R *with* $\mathrm{dom}(R) = (S_i)_{i \in I}$ *and* $a_i \in M_{S_i}$

$$R^{\mathcal{N}}((h_{S_i}(a_i))_{i \in I}) = R^{\mathcal{M}}((a_i)_{i \in I}).$$

If the sorts are clear, we shall often simply write $h((a_i)_{i \in I})$ *rather than specifying which function from* h *acts on the individual elements from* (a_i)*.*

If, further, we have that, for any $\mathcal{L}$*-formula* $\varphi((x_i)_{i \in I})$ *and* $(a_i) \in \mathcal{M}$

$$\varphi^{\mathcal{N}}((h(a_i))_{i \in I}) = \varphi^{\mathcal{M}}((a_i)_{i \in I}),$$

then we say that h *is an* ***$\mathcal{L}$-elementary embedding*** *of* $\mathcal{M}$ *into* $\mathcal{N}$*. If* $\mathcal{M} \subseteq \mathcal{N}$ *and the inclusion map is an* $\mathcal{L}$*-elementary embedding, we shall write* $\mathcal{M} \preceq \mathcal{N}$ *and say that* $\mathcal{M}$ *is an elementary substructure of* $\mathcal{N}$*.*

Definition 2.19. *Let* $\mathcal{L}$ *be a language and* $\mathcal{M}, \mathcal{N}$ *be* $\mathcal{L}$*-structures. We say that* $h : \mathcal{M} \hookrightarrow \mathcal{N}$ *is an* **$\mathcal{L}$-isomorphism** *if* h *is a surjective* $\mathcal{L}$*-embedding. If there exists an* $\mathcal{L}$*-isomorphism between* $\mathcal{M}$ *and* $\mathcal{N}$*, we shall say that they are* $\mathcal{L}$*-isomorphic and write* $\mathcal{M} \cong \mathcal{N}$*.*

Remark 2.5. *Note that structures being isomorphic implies that they are in fact isometric as metric spaces. Further, any isomorphism must map base points to base points, as we* *always include $\star_S$ as a symbol within our logic. If our language includes additional function and relation symbols, we see that $\mathcal{L}$-structures being $\mathcal{L}$-isomorphic is an incredibly strong requirement, as these functions and relations must also be preserved under the isometry. As expected, even though we do not assume that an $\mathcal{L}$-isomorphism is an elementary embedding, isormorphic structures satisfy the same $\mathcal{L}$-formulae in the following sense.*

Proposition 2.7. *Let $\mathcal{M}$, $\mathcal{N}$ be $\mathcal{L}$-structures and $h : \mathcal{M} \hookrightarrow \mathcal{N}$ an $\mathcal{L}$-isomorphism. Then for any $\mathcal{L}$-formula $\varphi((x_i)_{i\in I})$ and $(a_i) \in \mathcal{M}$, we have*

$$\varphi^{\mathcal{N}}(h(a_i)_{i\in I}) = \varphi^{\mathcal{M}}((a_i)_{i\in I})$$

Proof. By induction on construction of $\mathcal{L}$-formulae. □

2.4 EF Games

Although structures being $\mathcal{L}$-isomorphic implies that they are elementarily equivalent, we note that the converse is not always true. We shall see later that certain $\mathbb{R}$-trees may be elementarily equivalent but yet not homeomorphic, let alone isometric!

In order to demonstrate this, we need a simple way to test for elementary equivalence between two $\mathcal{L}$-structures. The way this is handled in classical model theory is via *Ehrenfeucht-Fraïssé games*. In [20] and [21], Bradd Hart develops a version of EF games for a different version of continuous model theory, and we adapt this strategy to our setting. We comment later on some of the issues which arise in attempting to adapt this to our setting.

We shall begin by first defining EF-like games that require far too much data. This will be useful in proving that a winning strategy for EF-like games is equivalent to two structures being elementarily equivalent. We shall then refine our notion of EF-like game to something much easier to play.

Definition 2.20. *We give the data for an* **(atomic/general) CNA game** *between $\mathcal{L}$-structures $\mathcal{M}$ and $\mathcal{N}$ as follows:*[11]

- *a sequence of (atomic/general) $\mathcal{L}$-formulas $(\varphi_i(\bar{x}))_{i\in I}$ amongst them at most finitely many (say, n) free variables $\bar{x}$,*
- *a sequence $r_1, \ldots, r_n$ of positive reals,*
- *an $\epsilon > 0$,*
- *a finite set $\Gamma \subseteq I$, and*
- *finite ϵ-covers $\mathcal{C}_{\xi_1}, \ldots, \mathcal{C}_{\xi_k}$ of the images of $\varphi_{\xi_i}(B_{\bar{r}})$ for each $\xi_i \in \Gamma$. Each ϵ-cover will be made up of closed intervals with nonempty interior, each of length at most ϵ, such that no three intersect and no endpoint is an endpoint of two intervals.*

A CNA game is played as one would play an EF game in classical logic with a small additional rule. On turn i, Player I may choose play $a_i \in \mathcal{M}$ (or, resp. $b_i \in \mathcal{N}$) subject to the condition that $\|a_i\| < r_i$ (resp. $\|b_i\| < r_i$), and Player II then returns with either $b_i \in \mathcal{N}$ (resp. $a_i \in \mathcal{M}$) subject to the same constraint.

Player II wins if for each $k \in \Gamma$, there is some $C \in \mathcal{C}_{\xi_k}$ such that $\varphi^{\mathcal{M}}_{\xi_k}(\bar{a}), \varphi^{\mathcal{N}}_{\xi_k}(\bar{b}) \in C$.

We shall sometimes refer to the set of all valid moves as the **field of play**. It is worth noting that each of our $\mathcal{L}$-formulae in a CNA game is a controlled function from a finite product of sorts, hence is bounded and uniformly continuous on bounded subsets; thus our final condition requiring a finite ϵ-cover is well-defined.

We introduce the following notion which is helpful for quantifying when two $\mathcal{L}$-structures are "almost" equivalent.

[11] We have chosen the abbreviation CNA as this is the abbreviation for an overly convoluted board game that, arguably, would take longer to play than any finite length EF game.

Definition 2.21. *If F is a collection of $\mathcal{L}$-sentences and $\epsilon > 0$, then we shall say that $\mathcal{M}$ and $\mathcal{N}$ are* ϵ**-equivalent on** *F, denoted $\mathcal{M} \equiv_\epsilon^F \mathcal{N}$, if for each $\mathcal{L}$-sentence $\varphi \in F$, we have $|\varphi^{\mathcal{M}} - \varphi^{\mathcal{N}}| < \epsilon$. If F is the class of all $\mathcal{L}$-sentences, then we shall say that $\mathcal{M}$ and $\mathcal{N}$ are* ϵ**-equivalent** *and write $\mathcal{M} \equiv_\epsilon \mathcal{N}$.*

We used the letter F in the previous definition because, in nearly all instances, we shall care about showing ϵ-equivalence on only a finite set of $\mathcal{L}$-sentences.

Admittedly, the win condition for Player II in a CNA game is a bit odd. We clarify the win condition with the following lemma.

Lemma 2.3. *Fix a general CNA game as above, let $\mathcal{L}_{\bar{c}}$ by an expansion by constants where $\lambda_{c_i} := r_i$, and define $F := \{\varphi_{\xi_k}(\bar{c}) : k \in \Gamma\}$. Then there is an $\epsilon' > 0$ such that if Player II can guarantee $(\mathcal{M}, \bar{a}) \equiv_{\epsilon'}^F (\mathcal{N}, \bar{b})$, then Player II wins the CNA game.*

Proof. It should be clear that all plays respect the controllers for the new constants, hence it makes sense to view $(\mathcal{M}, \bar{a})$ and $(\mathcal{N}, \bar{b})$ as $\mathcal{L}_{\bar{c}}$-structures. Given that for each $k \in \Gamma$, we have a finite ϵ-cover $\mathcal{C}_{\xi_k} = \{C_1, \ldots, C_m\}$ consisting of closed intervals with nonempty interior of length at most ϵ such that no three intersect and no endpoint is the endpoint of two intervals.

Then, let $\epsilon' := \frac{1}{2}\min(\text{len}(C_i \cap C_j) : C_i \cap C_j \neq \varnothing)$. It then suffices to show that, for each $t \in \bigcup \mathcal{C}_k$, there is some $C \in \mathcal{C}_k$ such that $[t - \epsilon', t + \epsilon'] \subseteq C$. If there is a unique $C \in \mathcal{C}_k$ such that $t \in C$, then clearly $[t - \epsilon', t + \epsilon'] \subseteq C$. However, if there are $C, C' \in \mathcal{C}_k$ such that $t \in C, C'$ and $C \neq C'$, then we note that $C \cap C' = [c, c']$ some subinterval. WLOG, assume that c' is the maximal element of C. Then if $t \geqslant \frac{1}{2}(c + c')$, we see that $[t - \epsilon', t + \epsilon'] \subseteq C'$, and if $t \leqslant \frac{1}{2}(c + c')$, then $[t - \epsilon', t + \epsilon'] \subseteq C$. □

We now prove our result on CNA games in this version of continuous logic. Unfortunately, we have not been able to establish the full equivalence that one finds in First Order Logic or bounded continuous logic. Namely, we have not been able to show that $\mathcal{M} \equiv \mathcal{N}$ implies that Player II has a winning strategy in all CNA games between $\mathcal{M}$ and $\mathcal{N}$. The difficulty arises from the fact that we must restrict the field of play, and that $\inf_x \big|_r^s \varphi(x)$ does not come close to capturing existence of a witness to $\varphi(x)$ in $B_r(\mathcal{M})$ unless the interpretation of this $\mathcal{L}$-sentence is exactly 0. In [21], the crucial lemma for this portion of the proof relies on, in some sense, being able to guarantee a winning strategy for a length 1 game by guaranteeing ϵ-equivalence on some other (more complicated) finite set of $\mathcal{L}$-sentences.

In Appendix B, we elaborate on some proposed adaptations of this lemma and explain the technical difficulties in fully adapting this proof to our setting. However, even without this full equivalence, we are still able to recover a test for elementary equivalence of $\mathcal{L}$-structures.

Theorem 2.1. *Given two $\mathcal{L}$-structures $\mathcal{M}$ and $\mathcal{N}$, TFAE:*

1. *Player II has a winning strategy for all general CNA games.*

2. *Player II has a winning strategy for all atomic CNA games.*

Further, if either of the above is true, then $\mathcal{M} \equiv \mathcal{N}$.

Proof. We note that (1) implies $\mathcal{M} \equiv \mathcal{N}$ because Player II can win any general CNA game where each formula has at most 0 free variables. Because she can win for any $\epsilon > 0$, this implies that for any $\mathcal{L}$-sentence φ and $\epsilon > 0$, we have $|\varphi^{\mathcal{M}} - \varphi^{\mathcal{N}}| < \epsilon$, hence $\varphi^{\mathcal{M}} = \varphi^{\mathcal{N}}$. It is worthwhile to note that, by the data given in our CNA game, we have shown that we may inductively construct any $\mathcal{L}$-formulae φ where every subformula of φ contains at most finitely many free variables. By Proposition 2.2, this encompasses all $\mathcal{L}$-sentences.

Now, for the equivalence of (1) and (2). Of course, (1) $\Rightarrow$ (2) is obvious.

We'll prove (2) $\Rightarrow$ (1) by induction on the depth of formulas. The base case is trivial. Suppose that $\varphi(\bar{x})$ is some $\mathcal{L}$-formula and that Player II has a winning strategy for all CNA games where the depth of each formula is $< \operatorname{dep}(\varphi)$.

- Suppose that $\varphi(\bar{x}) = f(\psi_\xi(\bar{x}))$ where $f : \mathbb{R}^{\alpha}_{\geqslant 0} \to \mathbb{R}_{\geqslant 0}$ is some controlled function.[12] Now, because $\mathcal{L}$-formulae are controlled functions, we know that there are controllers $(\lambda_\xi, N_\xi, \delta_\xi)$ for each $\psi_\xi(\bar{x})$ where $\xi \in \alpha$. And, because all eventual plays $\bar{a}$ and $\bar{b}$ must satisfy that $\|a_n\|, \|b_n\| < r_n$, we know that each $\psi_\xi(\bar{x})$ must be bounded by $\lambda_\xi(\bar{r})$.

  Define a sequence $(\rho_\xi)_{\xi<\alpha}$ by $\rho_\xi := \lambda_\xi(\bar{r})$. Now, because f is controlled, we know that there are (λ, N, δ) that act as controllers for f with data (ρ_ξ) and ϵ. In particular, $N := N((\rho_\xi), \epsilon) \subseteq \alpha$ is a finite set of indices such that, if $\|\psi_\xi(\bar{x})\| < \rho_\xi$ for each $\xi \in N$ (which is necessarily the case for any valid play of this game), then if $|\psi_\xi(\bar{a}) - \psi_\xi(\bar{b})| < \delta_\xi$, we know that $|f((\psi_\xi(\bar{a}))) - f((\psi_\xi(\bar{b})))| < \epsilon$.

  Now, set $\delta := \min\{\delta_\xi : \xi \in N\}$, and fix finite δ-covers $\mathcal{D}_1, \ldots, \mathcal{D}_k$ of $[-\rho_{\xi_1}, \rho_{\xi_1}]$, $\ldots$, $[-\rho_{\xi_k}, \rho_{\xi_k}]$ such that, for $D_1 \in \mathcal{D}_1, \ldots, D_k \in \mathcal{D}_l$, we have

$$f(\mathbb{R} \times \cdots \times D_i \times \cdots \mathbb{R} \times \cdots) \subseteq C \quad \text{for some } C \in \mathcal{C}$$

  (here, we mean that we've replaced the $(\xi_i)^{th}$-coordinate of the domain of f with a corresponding interval, of length at most δ, coming from our δ-cover).

  Now, we simply play a CNA game with the finitely many ψ_{ξ_i}'s appearing in our list of formulae, each of their coordinates captured by some finite set N', using the same $\bar{r}$, using δ in the role of ϵ, and using the covers $\mathcal{D}_i$ in the place of the $\mathcal{C}_i$.

  By the inductive hypothesis, Player II has a winning strategy to guarantee that each of the $\psi_{\xi_i}(\bar{a}), \psi_{\xi_i}(\bar{b}) \in D_i$, and hence $\varphi(\bar{a}), \varphi(\bar{b}) \in C$.

- Suppose that $\varphi(\bar{x}) = \sup_y \big|_s^t \psi(\bar{x}, y)$. For the sake of simplifying notation, we assume that y is a single variable, but the case in which $\bar{y}$ is a finite tuple of variables is very

[12] For notational simplicity, we assume that each $\psi_\xi(\bar{x})$ contains, as free variables, exactly the finitely many variables $\bar{x}$.

similar.

Fix some partition $(s, s+\gamma, \ldots, s+(l-1)\gamma, t)$ of $[s,t]$ such that, for any $\mathcal{L}$-structure $\mathcal{M}$ and $\bar{a} \in B_{\bar{r}}(\mathcal{M})$, the upper Riemann sum approximates the integral to within $(t-s)\epsilon/3$. More precisely, for any choice of $\mathcal{L}$-structure $\mathcal{M}$, we have that for any $\bar{a} \in B_{\bar{r}}(\mathcal{M})$,

$$\left|\left(\sum_{i=1}^{l} \gamma \cdot \sup_{y \in B_{s+i\gamma}(\mathcal{M})} \psi(\bar{a}, y)\right) - \int_s^t \sup_{y \in B_\rho(\mathcal{M})} \psi(\bar{a}, y) d\rho\right| < \frac{(t-s)\epsilon}{3}.$$

Such a partition exists by Proposition 2.3.

Now, consider the CNA game with data given by $\psi(\bar{x}, y)$, positive reals $\bar{r}, s+k\delta$ (for some fixed k amongst $1, \ldots, l$), $\frac{(t-s)\epsilon}{3\gamma l}$ in place of ϵ, and a finite $\frac{(t-s)\epsilon}{3\gamma l}$ cover $\mathcal{C}'$ of $\psi(B_{\bar{r}, s+k\delta})$ that refines and expands on the cover of the original game. I claim that Player II has a strategy to guarantee that

$$\left|\gamma \cdot \sup_{y \in B_{s+k\gamma}} \psi(\bar{a}, y) - \gamma \cdot \sup_{y \in B_{s+k\gamma}} \psi(\bar{b}, y)\right| < \gamma \cdot \frac{(t-s)\epsilon}{3\gamma l} = \frac{(t-s)\epsilon}{3l}.$$

This is essentially an adaptation of Hart's argument in [20]. Because $\operatorname{dep}(\psi) < \operatorname{dep}(\varphi)$, then by the inductive hypothesis, we see that Player II has a winning strategy for any game with just $\psi(\bar{x}, y)$ as the $\mathcal{L}$-formula (even though this is a longer game).

Now, if $\bar{a}$ and $\bar{b}$ have been played according to Player II's winning strategy, so that before the last turn of the game, we are in the position of $\psi(\bar{a}, y)$ and $\psi(\bar{b}, y)$. Now, suppose that $\sup_{y \in B_{s+k\gamma}(\mathcal{M})} \psi(\bar{a}, y) < \sup_{y \in B_{s+k\gamma}(\mathcal{N})} \psi(\bar{b}, y)$. Then, using Player I's help, we note that we may fix a sequence (b_n) such that $\psi(\bar{b}, b_n)$ monotonically increases to $\sup_{y \in B_{s+k\gamma}(\mathcal{N})} \psi(\bar{b}, y)$. For each b_n, there is thus a corresponding a_n and $C_n \in \mathcal{C}'$ according to Player II's winning strategy which will guarantee that $\psi(\bar{a}, a_n), \psi(\bar{b}, b_n) \in C_n$. Since $\mathcal{C}'$ is finite and $(\psi(\bar{b}, b_n))_{n \in \mathbb{N}}$ is an increasing sequence, we see that there is a $C \in \mathcal{C}'$ containing all but finitely many of the $\psi(\bar{b}, b_n)$, hence $\sup_{y \in B_{s+k\gamma}(\mathcal{N})} \psi(\bar{b}, y) \in C$, and similarly for $\sup_{y \in B_{s+k\gamma}(\mathcal{M})} \psi(\bar{a}, y)$. The claim follows.

Now, to simplify notation, we let $\theta_i(\bar{x}) := \gamma \cdot \sup_{y \in B_{s+i\gamma}} \psi(\bar{x}, y)$. Note that $\theta_i(\bar{x})$ is *not* an $\mathcal{L}$-formula. However, Player II *does* have a "winning strategy" for each $\theta_i(\bar{x})$, hence has a "winning strategy" for all finitely many of them simultaneously; this is seen by going through the same argument as above and playing the game of smaller depth with more $\mathcal{L}$-formulae. This implies that Player II has a strategy that guarantees

$$\left|\sum_{i=1}^{l} \theta_i(\bar{a}) - \sum_{i=1}^{l} \theta_i(\bar{b})\right| < \frac{(t-s)\epsilon}{3}.$$

This, together with the fact that

$$\left|\sum_{i=1}^{l} \theta_i(\bar{a}) - \int_s^t \sup_{y \in B_\rho} \psi(\bar{a}, y) d\rho\right| < \frac{(t-s)\epsilon}{3}$$

and

$$\left|\sum_{i=1}^{l} \theta_i(\bar{b}) - \int_s^t \sup_{y \in B_\rho} \psi(\bar{b}, y) d\rho\right| < \frac{(t-s)\epsilon}{3},$$

guarantees that

$$\left|\int_s^t \sup_{y \in B_\rho} \psi(\bar{a}, y) d\rho - \int_s^t \sup_{y \in B_\rho} \psi(\bar{b}, y) d\rho\right| < (t-s)\epsilon,$$

hence

$$\begin{aligned} & \left|\frac{1}{t-s}\int_s^t \sup_{y \in B_\rho} \psi(\bar{a}, y) d\rho - \frac{1}{t-s}\int_s^t \sup_{y \in B_\rho} \psi(\bar{b}, y) d\rho\right| \\ = \quad & \frac{1}{t-s}\left|\int_s^t \sup_{y \in B_\rho} \psi(\bar{a}, y) d\rho - \int_s^t \sup_{y \in B_\rho} \psi(\bar{b}, y) d\rho\right| \\ < \quad & \epsilon. \end{aligned}$$

- The inf case is handled similarly to the sup case.

Lastly, we note that we've only demonstrated how Player II is to win for a single formula. Since she only needs a winning strategy for finitely many formulas, she can clearly refine her strategy finitely many times to be compatible with finitely many other $\mathcal{L}$-formulas. □

As we mentioned, the data for a CNA game is far too cumbersome. We simplify things in the following Definition and Corollary.

Definition 2.22. *We give the data for an* **EF game** *between two $\mathcal{L}$-structures $\mathcal{M}$ and $\mathcal{N}$:*

- *finitely many atomic $\mathcal{L}$-formulas $\Gamma := \{\varphi_1(\bar{x}), \dots \varphi_k(\bar{x})\}$, each of which only has finitely many free variables $x_1, \dots, x_n$, and*
- *real numbers $r_1, \dots, r_n, \epsilon > 0$.*

The plays in an EF game follow the same rules as a CNA game, and Player II wins if $|\varphi_i^{\mathcal{M}}(\bar{a}) - \varphi_i^{\mathcal{N}}(\bar{b})| < \epsilon$ for all $i = 1, \dots, k$.

We shall denote such an EF game as $\mathcal{G}(\mathcal{M}, \mathcal{N}, \Gamma, \bar{r}, \epsilon)$

Corollary 2.1. *For any two $\mathcal{L}$-structures $\mathcal{M}$ and $\mathcal{N}$, if Player II has a winning strategy for all atomic EF games between $\mathcal{M}$ and $\mathcal{N}$, then $\mathcal{M} \equiv \mathcal{N}$.*

Proof. This follows from noting that the previous data was only needed for the proof so that we may build up our formulae by induction. Hence, any strategy for an EF game translates to a strategy for a CNA game and vice versa. We have reduced our necessary data to the only pieces Player II needs in order to actually play and check her win condition. We have simplified our win condition according to Lemma 2.3. □

It is worth noting that we cannot reduce from finitely many r_i to a single r. Of course, having a strategy for an arbitrary CNA game guarantees a strategy for an EF game by picking the CNA strategy when each $r_i = r$. However, knowing that one has a strategy for every EF game, it is not clear how to strategize for every CNA game. In particular, it does *not* suffice to "open the field of play" by taking $r := \max_i(r_i)$; this is because if Player I plays according to the constraint for some $r_i < r$, then Player II still has the opportunity to play outside this constraint. Conversely, any smaller value for r restricts Player I's moves!

However, we may come across situations in which the choice of r_i and ϵ does not change Player II's strategy. More precisely, we shall introduce the following terminology:

Definition 2.23. *Suppose that Player II has a winning strategy that works for all EF games $\mathcal{G}(\mathcal{M},\mathcal{N},\Gamma,\bar{r},\epsilon)$ as one varies ϵ.*[13] *Then we shall say that Player II has a* **perfect strategy** *for the game $\mathcal{G}(\mathcal{M},\mathcal{N},\Gamma,\bar{r},\epsilon)$.*

Clearly, if two spaces are isometric, Player II will have a perfect strategy for any EF game. One may reasonably ask if this is the only case in which Player II may have a perfect strategy for all atomic EF games, but we shall see explicit examples later where Player II has a perfect strategy between spaces that are non-homeomorphic. In fact, in [15], it is shown that in a variant of continuous logic, if two structures are both $\aleph_1$-saturated (cf. Definition 2.40), then they are e.e. up to depth k if and only if Player II has a perfect strategy in all EF games of length k. It is important to note that "depth" in this case does not precisely coincide with our notion, although they are certainly related. The structures that we see later will indeed satisfy this condition, and thus, it should not be surprising that if Player II has a winning strategy, then she indeed has a perfect strategy.

We have not proven that an analogous result about $\aleph_1$-saturated structures is true in this version of continuous model theory, though we suspect it would not be difficult to do so.

2.5 Definability

We would be remiss to write a dissertation on model theory and not define what it means for a subset of a product of a structure to be definable. We shan't utilize this notion much in our work, but we will comment on a few related questions later on.

In contrast to classical model theory, we shall *not* say that $A \subseteq \mathcal{M}^n$ is definable when A is

[13] I.e., Player II's strategy will guarantee that $|\varphi_i^{\mathcal{M}}(\bar{a}) - \varphi_i^{\mathcal{N}}(\bar{b})| = 0$.

the zeroset of some formula $\varphi(\bar{x})$. The reason for this is that it does not lead to our true goal. The goal of definability is to be able to quantify over particular subsets of a structure (or a product of sorts from a structure), and while the "zerosets" of formulae in classical logic lead to the right notion of definability, the same is not the case in continuous model theory.

Continuous model theorists will know that a definable set $A \subseteq M^n$ is one in which there is an $\mathcal{L}$-formula $\varphi(\bar{x})$ that interprets as $d(\bar{x}, A)$. In our version of continuous model theory, we have intentionally built up our formulae so that they may have arbitrarily many free variables, and this was done specifically so that we might be able to have formulae that define, say, the set of geodesics through the base point. In this manner, we should want to view a geodesic as a sequence of points indexed by $\mathbb{R}$. This means, however, that we shall need to be more cautious in our notion of definability, as we are no longer talking about finite products of sorts from a structure, but indeed arbitrary products.

The definition we give below is not the definition that Luther gives. However, it is shown to be equivalent in his Corollary 3.9.4, and the below definition is almost natural to a continuous model theorist after the previous discussion.

Definition 2.24. *Let $\mathcal{M}$ be an $\mathcal{L}$-structure, $C \subseteq \mathcal{M}$, $A \subseteq \prod_{i \in I} \mathcal{M}$. We say that that A is* **definable in $\mathcal{M}$ over** *C if, for every finite projection $\pi_N(A)$ and every $\bar{r} \in \mathbb{R}^n_+$, there is an $\mathcal{L}_C$-formula $\varphi(x_1, \ldots, x_N)$ such that, when $\bar{x} \in \bar{B}_{\bar{r}}(M)$, we have*

$$\varphi^{\mathcal{M}}(x_1, \ldots, x_N) = \bar{d}(\bar{x}, \pi_N(A))$$

where $\bar{d}$ denotes the maximum of the distances $d_i(x_i, \pi_i(A))$.

We present the following Proposition and Corollary to demonstrate that we indeed have the "correct" definition of definability for this logic.

Proposition 2.8. *Suppose $A \subseteq \prod_{i=1}^{n} \mathcal{M}$ and $C \subseteq \mathcal{M}$. Then TFAE:*

1. *A is definable in $\mathcal{M}$ over C.*

2. *There is $(c_n)_{n\in\mathbb{N}} \in C$ such that for any $\mathcal{L}$-formula $\varphi(\bar{x}, (y_i)_{i\in I})$ (where $\bar{x} = x_1, \ldots, x_n$), and for any sequences $\bar{r} \in \mathbb{R}_{\geqslant 0}$, $\bar{r}', (R_i)_{i\in I} \in \mathbb{R}_+$ such that $\bar{r}$, $\bar{r}'$ are of length n and each $r_i < r_i'$, there is a finite $F \subseteq I$ and $\mathcal{L}$-formula $\psi((y_i)_{i\in I}; (z_n))$ such that $\psi^{\mathcal{M}}((y_i)_{i\in I}; (c_n))$ is the function*

$$\left(\prod_{i=1}^{n} \frac{1}{r_i' - r_i}\right) \int_{r_n}^{r_n'} \cdots \int_{r_1}^{r_1'} \sup_{\bar{x} \in B_{\rho_1,\ldots,\rho_n}(A)} \varphi^{\mathcal{M}}(\bar{x}, (y_i)_{i\in I})\, \mathrm{d}\rho_1 \cdots \mathrm{d}\rho_n$$

   when restricted to y satisfying $\|y_i\| \leqslant R_i$ for all $i \in F$.

   Similarly for inf *in place of* sup.

This seems like a very convoluted proposition, but note that the expression is essentially what one should expect for the interpretation of $\sup_{\bar{x}\in A} \Big|_{\bar{r}}^{\bar{r}'} \varphi(\bar{x}; (y_i))$! Indeed, if A is bounded, one may take $\bar{r}$ larger than the bound on A and get exact sup and inf, rather than the average:

Corollary 2.2. *Suppose $A \subseteq \prod_{i=1}^{n} \mathcal{M}$ is a bounded definable set. Then for any $\varphi(\bar{x}; (y_i)_{i\in I})$, given bounds on the y_i, then there are formulae equivalent to $\sup_{\bar{x}\in A} \varphi^{\mathcal{M}}(\bar{x}; (y_i)_{i\in I})$ and $\inf_{\bar{x}\in A}(\bar{x}; (y_i)_{i\in I})$.*

We note that this characterization of definability is subject to A being a subset of a finite product of sorts from $\mathcal{M}$. This may seem odd considering the earlier discussion that we expanded our logic in order to be able to discuss definability of, say, geodesic rays from a base point. However, this is in principal not so bad, as Luther shows that under certain assumptions, one need only index rays by $\mathbb{N}$, and he develops a method for "quantifying over countably indexed sequences." In this sense, one need only guarantee definability for the first finitely many unit steps along rays from a base point in order to guarantee genuine quantification over all rays emanating from a base point.

While definable sets are important, we are also interested in which predicates are definable, in some sense. This is typically done by defining a pseudometric on the space of all $\mathcal{L}$-formulae, perhaps subject to a given theory, and then completing the metric space that one obtains after modding out by the equivalence relation induced by the pseudometric. For our logic, we will not be able to obtain a pseudometric space, as many $\mathcal{L}$-formulae will have unbounded distance from each other. However, this does not stop us from working directly with formulae, as we see in the following definition.

Definition 2.25. *Fix an $\mathcal{L}$-theory T. Let $\mathcal{F}_{(x_i)_{i\in I}}$ denote the class of $\mathcal{L}$-formulae whose free variables consist of $(x_i)_{i\in I}$. Then we define the T-**separation** ρ_T on $\mathcal{F}_{(x_i)_{i\in I}}$ by*

$$\rho_T(\varphi,\psi) := \begin{cases} \sup\left\{\left|\varphi^{\mathcal{M}}((a_i)) - \psi^{\mathcal{M}}((a_i))\right| : \mathcal{M} \models T,\ (a_i) \in \mathcal{M}\right\} & \text{when this is bounded} \\ \infty & \text{else.} \end{cases}$$

It is easy to see that ρ_T is indeed an extended pseudometric; it is also not hard to find $\mathcal{L}$-formulae with unbounded separation. In fact, the $\mathcal{L}$-formulae $d(x,x)$ and $d(\star,x)$ suffice as simple examples. Regardless, the notion of a Cauchy sequence is still well-defined, even in an unbounded pseudometric space. This leads us to the following definition:

Definition 2.26. *We say that a T-**formula** is any equivalence class of Cauchy sequences of $\mathcal{L}$-formulae from some $\mathcal{F}_{(x_i)_{i\in I}}$, where two sequences are equivalent if for each $\epsilon > 0$, their formulae are eventually T-separated by at most ϵ.*

We shall not need these notions in our present work. However, we shall later encounter Robinson's asymptotic field ${}^{\rho}\mathbb{R}_{\mathcal{U}}$, and should one wish to show that ${}^{\rho}\mathbb{R}_{\mathcal{U}}$ lives in $\mathcal{K}^{eq}$ for a particular class of $\mathcal{L}$-structures, then one should develop a more robust notion of T-formulae in order to rigorously define $\mathcal{K}^{eq}$ within this logic.

2.6 Ultraproducts

An important tool from classical model theory is the notion of an *ultraproduct* of $\mathcal{L}$-structures. Intuitively, this is a way of taking an indexed set of $\mathcal{L}$-structures and creating a new $\mathcal{L}$-structure. This will be useful in defining asymptotic cones later. In order to give a precise definition, we first need to understand *filters* and *ultrafilters*.

Definition 2.27. *Let I be an indexing set. We say that $\mathcal{F} \subseteq \mathcal{P}(I)$ is a* **filter** *if each of the following conditions hold:*

1. *$\varnothing \notin \mathcal{F}$ and $I \in \mathcal{F}$.*
2. *If $A \in \mathcal{F}$ and $B \supseteq A$, then $B \in \mathcal{F}$ (we say that $\mathcal{F}$ is* **closed upwards***).*
3. *If $A, B \in \mathcal{F}$, then $A \cap B \in \mathcal{F}$.*

One should think of the sets in a filter as being the "large" sets. This is a helpful intuitive explanation, but once we introduce *ultrafilters*, we shall see that it is a very precise explanation as well. It is commonplace to provide a few guiding examples of filters, and so we do that here.

Example 2.2. *Fix I and some $a \in I$. Then $I_a := \{X \subseteq I : a \in X\}$ is a filter. Any filter of this form is called a* **principal filter***.*

In almost every case from here on, we shall ignore the principal filters.

Example 2.3. *Let I be a set of cardinality $\kappa \geqslant \aleph_0$. For any infinite $\mu \leqslant \kappa$, we define*

$$\mathcal{F}_{I,\mu} := \{X \subseteq I : |I \smallsetminus X| < \mu\}$$

and call this the μ-**Fréchet filter** *on* I.[14] *If* $\mu = \aleph_0$*, then we shall simply say "the Fréchet filter on* I*" and write* $\mathcal{F}_I$*. It is not hard to check that this indeed satisfies the conditions of being a filter.*

1. *Clearly,* $|I \smallsetminus \varnothing| = |I| \nless \mu$*, and* $|I \smallsetminus I| = 0 < \mu$*, hence the first condition holds.*
2. *If* $A \in \mathcal{F}_{I,\mu}$ *and* $B \supseteq A$*, then clearly* $|I \smallsetminus B| \leqslant |I \smallsetminus A| < \mu$*, hence the second condition holds.*
3. *If* $A, B \in \mathcal{F}_{I,\mu}$*, then because*

$$I \smallsetminus (A \cap B) = (I \smallsetminus A) \cup (I \smallsetminus B),$$

we see that $|I \smallsetminus (A \cap B)| = |(I \smallsetminus A) \cup (I \smallsetminus B)| \leqslant |I \smallsetminus A| + |I \smallsetminus B| < \mu + \mu = \mu$*, hence the third condition holds.*

Later in the dissertation, we shall primarily concern ourselves with filters on $\mathbb{N}$ which extend the Fréchet filter. As this filter consists exactly of the cofinite subsets of $\mathbb{N}$, it is almost natural to think that the sets $[n, \infty)$ "generate" the Fréchet filter in some sense, and this is indeed the case. We first need a bit of terminology in order to define this rigorously.

Definition 2.28. *Let* $\mathcal{F}$ *be a filter on* I*. We say that* $\mathcal{B} \subseteq \mathcal{F}$ *is a* **base for** $\mathcal{F}$ *if*

$$\mathcal{B}{\uparrow} := \{X \subseteq I : B \subseteq X \text{ for some } B \in \mathcal{B}\} = \mathcal{F}.$$

Intuitively, filter bases consist of the "smallest large sets." That is, if $\mathcal{B}$ is a base for $\mathcal{F}$, then for any $B \in \mathcal{B}$, if $A \in \mathcal{F}$ and $A \subseteq B$, then $A \in \mathcal{B}$. Of course, not every collection of subsets of I will form a filter base. In particular, if $\mathcal{B}$ contains sets $B_1, \ldots, B_n$ such that $\bigcap_{i=1}^n B_i = \varnothing$,

[14] N.B. The original Fréchet filter was defined just using $\mu = \aleph_0$, and so sometimes is referred to as the *cofinite filter*.

then because filters are closed under finite intersections and do not contain the empty set, this cannot be the base of some filter. In some sense, this is the only obstacle.

Definition 2.29. *We say that a collection $\mathcal{C}$ of sets has the* **finite intersection property** *(FIP) if, for any finite collection of sets $C_1, \ldots, C_n \in \mathcal{C}$,*

$$\bigcap_{i=1}^{n} C_i \neq \varnothing.$$

Proposition 2.9. *Suppose that $\mathcal{C}$ is a collection of subsets of I with the FIP. Then*

$$\langle \mathcal{C} \rangle := \{C_1 \cap \cdots \cap C_n : C_i \in \mathcal{C}, n \in \mathbb{N}_+\}\uparrow$$

is a filter, which we call the **filter generated by** *$\mathcal{C}$.*

Proof. This is just definition chasing.

1. $\varnothing \notin \langle \mathcal{C} \rangle$ follows from the assumption that $\bigcap_{i=1}^{n} C_i \neq \varnothing$. The fact that $I \in \langle \mathcal{C} \rangle$ follows from upward closure.

2. Follows from upwards closure.

3. Suppose $A, A' \in \langle \mathcal{C} \rangle$. Then there are $C_1, \ldots, C_n$ and $C'_1, \ldots C'_m$ such that

$$C_1 \cap \cdots \cap C_n \subseteq A, \quad C'_1 \cap \cdots \cap C'_m \subseteq A',$$

hence

$$C_1 \cap \cdots \cap C_n \cap C'_1 \cap \cdots \cap C'_m \subseteq A \cap A' \implies A \cap A' \in \langle \mathcal{C} \rangle$$

□

With this terminology, the following is an easy exercise:

Example 2.4. *Let $\mathcal{F}_{\mathbb{N}}$ be the Fréchet filter on $\mathbb{N}$, and let $X := \{[n, \infty) : n \in \mathbb{N}\}$. Then $\mathcal{F}_{\mathbb{N}} = \langle X \rangle$.*

We now introduce the notion of an *ultra*filter.

Definition 2.30. *Let $\mathcal{F}$ be a filter on I. We say that $\mathcal{F}$ is an* **ultrafilter** *if, additionally, for any $X \subseteq I$, either $X \in \mathcal{F}$ or $I \setminus X \in \mathcal{F}$.*

Notation. *We shall usually use the symbols $\mathcal{U}$ and $\mathcal{V}$ to denote ultrafilters.*

Of course, one cannot have both $X, I \setminus X \in \mathcal{F}$, as closing this under finite intersections would imply that $\varnothing \in \mathcal{F}$. It is natural to wonder whether or not ultrafilters exist. It is trivial to see that any principal filter is an ultrafilter. As we mentioned before, we shall primarily concern ourselves with *non*principal ultrafilters.

Definition 2.31. *An ultrafilter $\mathcal{U}$ on a set I is called* **nonprincipal** *if it is* not *of the form I_a for some $a \in I$.*

It is not obvious that nonprincipal ultrafilters on a set I exist. We shall prove shortly that they do, but we first characterize the nonprincipal ultrafilters in the following way:

Lemma 2.4. *Let I be an infinite set. Then an ultrafilter $\mathcal{U}$ on I is nonprincipal if and only if $\mathcal{U}$ extends the Fréchet filter on I.*

Proof. ($\Rightarrow$): Let $\mathcal{F}_I$ be the Fréchet filter on I. Suppose that $\mathcal{U}$ is nonprincipal on I, and for contradiction, suppose that $\mathcal{U} \not\supseteq \mathcal{F}_I$. Then there is some cofinite $X \subseteq I$ such that $X \notin \mathcal{U}$. Because $\mathcal{U}$ is an ultrafilter, then $I \setminus X \in \mathcal{U}$, which is a finite set. Because $\mathcal{U}$ is nonprincipal, we note that for each $x \in I \setminus X$, we must have that $\{x\} \notin \mathcal{U}$. Thus, $I \setminus \{x\} \in \mathcal{U}$, and

$$(I \setminus X) \cap \bigcap_{x \in I \setminus X} (I \setminus \{x\}) = \varnothing \in \mathcal{U},$$

a contradiction.

($\Leftarrow$): Suppose that $\mathcal{U}$ is an ultrafilter extending $\mathcal{F}_I$. Then clearly there is no $a \in I$ such that $\{a\} \in \mathcal{U}$, hence $\mathcal{U}$ is nonprincipal. □

This lemma demonstrates that, if $\mathcal{U}$ is known to be an ultrafilter, then it is enough to know that $\mathcal{U}$ contains *some* finite set to conclude that $\mathcal{U}$ is principal. This comports with the idea that filters ought to capture the "large" sets. If I is infinite, then it seems obvious that the finite sets ought to be small!

It should be clear that ultrafilters are maximal filters; that is, if $\mathcal{U}$ is a filter on I and $\mathcal{V} \supseteq \mathcal{U}$ is a filter, then $\mathcal{V} = \mathcal{U}$. The (somewhat) more surprising fact is that every maximal filter is an ultrafilter!

Proposition 2.10. *Let I be a set and $\mathcal{F}$ be a maximal filter on I. Then $\mathcal{F}$ is an ultrafilter.*

Proof. We prove this by contradiction. Suppose $\mathcal{F}$ is maximal on I and that $X \subseteq I$ is such that $X, I \smallsetminus X \notin \mathcal{F}$ (i.e. $\mathcal{F}$ is not ultra). Then I claim that $\mathcal{F} \cup \{I \smallsetminus X\}$ has the FIP. If not, then there are $F_1, \ldots, F_n$ such that

$$F_1 \cap \cdots \cap F_n \cap (I \smallsetminus X) = \varnothing \implies \bigcap_{i=1}^{n} F_i \subseteq X,$$

and because $\mathcal{F}$ is upwards closed, then $X \in \mathcal{F}$, a contradiction. Thus, $\mathcal{F} \cup \{I \smallsetminus X\}$ has the FIP, and $\langle \mathcal{F} \cup \{I \smallsetminus X\} \rangle$ is a filter properly extending $\mathcal{F}$, a contradiction. □

Proposition 2.11. *Let I be a set. Nonprincipal ultrafilters on I exist if and only if I is infinite.*

Proof. We note that a filter $\mathcal{U}$ is principal if and only if there is some $a \in I$ such that $\{a\} \in \mathcal{U}$. Thus, if I is finite and $\mathcal{U}$ were a nonprincipal ultrafilter on I, then $\{I \smallsetminus \{a\}\} \in \mathcal{U}$ for each

$a \in I$. But because $\mathcal{U}$ is closed under finite intersections, we'd have

$$\bigcap_{a \in I} (I \setminus \{a\}) = \varnothing \in \mathcal{U},$$

a contradiction.

Now, suppose that I is an infinite set. Let $\mathcal{F}$ be the Fréchet filter on I. By the argument in the proof of Proposition 2.10 and an application of Zorn's lemma, we can extend $\mathcal{F}$ to a maximal filter, hence an ultrafilter. The proof is complete by Lemma 2.4. □

One important way to view ultrafilters is by the following fact:

Proposition 2.12. *The ultrafilters on an infinite set I can be viewed as $\{0,1\}$-valued probability measures on $\mathcal{P}(I)$. More precisely, if $\mathcal{U}$ is an ultrafilter on I, then*

$$\mu(X) := \begin{cases} 1 & \text{if } X \in \mathcal{U} \\ 0 & \text{if } X \notin \mathcal{U} \end{cases}$$

is a probability measure on $\mathcal{P}(I)$.

We forgo a proof as this is merely to present an alternative viewpoint to those more comfortable with probability measures. For those familiar with probability measures, it is a routine exercise.

In certain instances, we shall require an additional assumption for our ultrafilters. Namely, all ultrafilters are closed under finite intersections, but it is not immediately clear if an ultrafilter could also be closed under countable intersections.

Definition 2.32. *We say that a nonprincipal ultrafilter $\mathcal{U}$ on a set I is* **countably complete** *if, for every countable collection $(A_n)_{n \in \mathbb{N}} \in \mathcal{U}$, we have that $\bigcap_{n \in \mathbb{N}} A_n \in \mathcal{U}$. If $\mathcal{U}$ is not countably complete, then we shall say that it is* **countably incomplete**.

It turns out that the existence of a countably complete nonprincipal ultrafilter is independent of ZFC and implies the existence of certain large cardinals; one may consult [14] for reference. So, if we come across a statement such as Proposition 2.15 that requires the assumption that $\mathcal{U}$ be countably incomplete, in several models of ZFC, this really is no extra assumption at all.

The following easy fact gives us an easy way to recharacterize the countably incomplete ultrafilters:

Fact 2.2. *An ultrafilter $\mathcal{U}$ on an indexing set I is countably incomplete if and only if there is a descending sequence of $A_n \in \mathcal{U}$ such that $\bigcap_{n\in\mathbb{N}} A_n \notin \mathcal{U}$. Here, by descending, we mean that $A_0 \supseteq A_1 \supseteq A_2 \supseteq \cdots$.*

Proof. $(\Rightarrow)$: If $\mathcal{U}$ is countably incomplete, then pick a countable collection of $(A_n)_{n\in\mathbb{N}} \in \mathcal{U}$ whose intersection is not in $\mathcal{U}$. We may then find a descending sequence by $A'_k := \bigcap_{n=0}^{k} A_n$, noting that each $A'_k \in \mathcal{U}$ by $\mathcal{U}$'s closure under countable intersections.

$(\Leftarrow)$: If such a sequence exists, then clearly it demonstrates the countable incompleteness of the ultrafilter. □

Ultrasection

We now are ready to introduce new constructions based on some choice of ultrafilter on an arbitrary infinite indexing set I. In classical model theory, the typical "ultra" constructions take on a flavor similar to what one sees in a measure theory course; namely, sequences indexed by I are identified if they agree "almost everywhere." The fact that one may view a ultrafilter as a $\{0,1\}$-valued probability measure actually makes one's use of "almost everywhere" precise. Whenever one defines a new object/idea dependent on an ultrafilter over an indexing set, it is common to prefix the terminology with "ultra."

In continuous model theory, we adapt things only slightly. As our structures are continuous in nature, we shall not merely identify sequences which *precisely* agree almost everywhere, but rather sequences which are ϵ-close almost everywhere for each $\epsilon > 0$.

Definition 2.33. *Let (X, d) be a metric space, I an indexing set, and $\mathcal{U}$ a nonprincipal ultrafilter on I. Then for a sequence $(x_i)_{i\in I} \in X$, we shall say that $x \in X$ is an* **ultralimit** *of (x_i) if, for each $\epsilon > 0$, $\{i \in I : d(x_i, x) < \epsilon\} \in \mathcal{U}$.*

Because metric spaces are Hausdorff, ultralimits are unique when they exist, and so we write $\lim_{\mathcal{U}} x_i = x$ to denote that *the* ultralimit of a sequence $(x_i)_{i\in I}$ is x. (The indexing set should always be understood when discussing the ultralimit of a sequence.) One should hope that our notion of ultralimit comports with the typical notion of limit in a metric space, and indeed this is true:

Proposition 2.13. *Let (X, d) be a metric space, $(x_n)_{n\in\mathbb{N}} \in X$ and $\mathcal{U}$ a nonprincipal ultrafilter on $\mathbb{N}$. Then if $\lim_{n\to\infty} x_n$ exists, $\lim_{\mathcal{U}} x_n$ exists and*

$$\lim_{n\to\infty} x_n = \lim_{\mathcal{U}} x_n.$$

Proof. Let $x = \lim_{n\to\infty} x_n$, and note that, for each $\epsilon > 0$, there is some $N \in \mathbb{N}$ such that for all $n \in [N, \infty)$, $d(x_n, x) < \epsilon$. Because $[N, \infty)$ is cofinite, $[N, \infty) \in \mathcal{U}$ and

$$[N, \infty) \subseteq \{n \in \mathbb{N} : d(x_n, x) < \epsilon\} \in \mathcal{U}.$$

Thus, $x = \lim_{\mathcal{U}} x_n$. □

It follows that ultralimits are a generalization of ordinary limits in a metric space. It turns out that ultralimits exist for unusual sequences as well:

Example 2.5. *Consider the sequence $(x_n) := (0, 1, 0, 1, \dots) \in \mathbb{R}^{\mathbb{N}}$. Then, for any ultrafilter*

$\mathcal{U}$ *on* $\mathbb{N}$ *we have*

$$\lim_{\mathcal{U}} x_n = \begin{cases} 0 & \textit{if } \{0, 2, 4, \dots\} \in \mathcal{U} \\ 1 & \textit{if } \{1, 3, 5, \dots\} \in \mathcal{U} \end{cases}$$

So, while this sequence does not converge as an ordinary limit, it does converge as an ultralimit. This is actually a special case of the following proposition. We also note that this demonstrates that the value to which a sequence "ultraconverges" is dependent on the choice of ultrafilter. This is important to keep in mind throughout the dissertation.

Proposition 2.14. *TFAE:*

1. (X, d) *is compact*
2. *For every sequence* $(x_n) \in X$ *and every ultrafilter* $\mathcal{U}$ *on* $\mathbb{N}$*,* $\lim_{\mathcal{U}} x_n$ *exists.*

Proof. For (1) $\Rightarrow$ (2), we prove the contrapositive. That is, suppose that there is a sequence $(x_n) \in X$ and $\mathcal{U}$ such that $\lim_{\mathcal{U}} x_n$ does not exist. Then for each $x \in X$, there is some neighborhood U_x such that $\{n \in \mathbb{N} : x_n \notin U_x\} \notin \mathcal{U}$. Clearly, the U_x form an open cover of X, and if we suppose for contradiction that there were a finite subcover U_{x_i} for $i = 1, \dots, k$, then

$$\varnothing = \left\{ n \in \mathbb{N} : x_n \notin \bigcap_{i=1}^{k} U_{x_i} \right\} = \bigcap_{i=1}^{k} \{n \in \mathbb{N} : x_n \in U_{x_i}\} \in \mathcal{U},$$

which is a contradiction.

For (2) $\Rightarrow$ (1), we prove sequential compactness. Namely, fix an arbitrary sequence $(x_n) \in X$. I claim that there is a convergent subsequence. Indeed, if we take arbitrary nonprincipal $\mathcal{U}$ on $\mathbb{N}$, then if we denote $x := \lim_{\mathcal{U}} x_n$, then we have

$$F_n := \{n \in \mathbb{N} : d(x_n, x) < 1/n\} \in \mathcal{U}.$$

Then, for each $m \in \mathbb{N}$, let $n_m := \min(i \in F_m : i \geqslant m)$. It follow that $(n_m)_{m \in \mathbb{N}}$ is strictly increasing, hence $(x_{n_m})_{m\in\mathbb{N}}$ is actually a subsequence of (x_n), and by construction, $d(x, x_{n_m}) \leqslant 1/m$, hence it is a convergent subsequence. □

Before we move into the construction of ultraproducts, we shall remind ourselves of the notion of a *pseudo*metric and see how it gives rise to a natural equivalence relation. This will be helpful in our definition, both in terms of stating things precisely, and also to simplify our notation a bit.

Definition 2.34. *Let X be a set. We say that $\rho : X \times X \to \mathbb{R}_{\geqslant 0}$ is a* **pseudometric** *if*

1. *$d(x, x) = 0$,*
2. *for all $x, y \in X$, $d(x, y) = d(y, x)$, and*
3. *for all $x, y, z \in X$, $d(x, z) \leqslant d(x, y) + d(y, z)$.*

Notably, the only requirement that would turn a pseudometric into a bona fide metric is requiring that $d(x, y) = 0 \implies x = y$. That is, pseudometrics allow for distinct points to be "distance" 0 from each other. This gives rise to the following equivalence relation:

Remark 2.6. *Let X be a set and ρ a pseudometric on X. Then for $a, b \in X$, we say that $a \sim_\rho b \iff \rho(a, b) = 0$. It is trivial to check that this is an equivalence relation, and we denote the $\sim_\rho$-equivalence class of a point $a \in X$ by $[a]_\rho$.*

If we define $d_\rho : (X/\sim_\rho)^2 \to \mathbb{R}_{\geqslant 0}$ by $d_\rho([a]_\rho, [b]_\rho) := \rho(a, b)$, then it is easy to see that d is a well-defined metric on $X/\sim_\rho$.

We are now prepared to introduce the notion of an *ultraproduct* relative to a sequence of $\mathcal{L}$-structures $(\mathcal{M}_i)_{i\in I}$ and ultrafilter $\mathcal{U}$ on I. In order to describe the sorts for the ultraproduct, we shall first introduce the notion of a *metric ultraproduct.* The idea is that we shall want to take an arbitrary product of metric spaces and once again recover a metric space by taking the ultralimit of distances between points. However, this does not work in general, as the distances between points may be unbounded on a $\mathcal{U}$-large set. The next definition explains the metric ultraproduct construction.

Definition 2.35. *Let $(X_i, d_i, \star_i)_{i \in I}$ be a sequence of pointed metric spaces indexed by a set I, and $\mathcal{U}$ an ultrafilter on I. We define the* **metric ultraproduct***, denoted $\prod_{i \in I}(X_i, d_i, \star_i)/\mathcal{U}$ as follows:*

1. *We define*

$$X := \left\{ (a_i) \in \prod_{i \in I} X_i : \lim_{\mathcal{U}} d_i(a_i, \star_i) < \infty \right\}.$$

That is, X consists of all sequences such that the ultralimit of the distances from the basepoints is finite. Note that this does not *mean that the sequences are bounded, but that they are bounded $\mathcal{U}$-almost everywhere.*

2. *We define a pseudometric ρ on X by*

$$\rho(a, b) := \lim_{\mathcal{U}} d_i(a_i, b_i)$$

where $a = (a_i)_{i \in I}$ and $b = (b_i)_{i \in I}$. We note that by restricting our focus to sequences whose ultralimit distance to the basepoint is bounded, by an application of the triangle inequality, ρ is indeed a pseudometric and does not take the value ∞.

3. *We let $\tilde{X} := X/\sim_\rho$, define $\star := [(\star_i)]_\rho$, and d be the resulting quotient metric on $\tilde{X}$.*

4. *The metric ultraproduct is then the pointed metric space $(\tilde{X}, d, \star)$.*

The above construction is crucial to understanding the notion of an ultraproduct. Indeed, if we are working in a language without function or relation symbols, then the ultraproduct of a sequence of $\mathcal{L}$-structures $(\mathcal{M}_i)_{i \in I}$ will simply consist of the metric ultraproducts of the sorts. It is important to point out that, in the above definition, we do *not* require the X_i to be complete metric spaces. In most cases, we shall concern ourselves with metric ultraproducts of complete metric spaces (and, in fact, ultraproducts of $\mathcal{L}$-structures), but it will be useful to have this more general notion. However, even if we begin with an arbitrary sequence of pointed metric spaces, none of which are complete, we still obtain a complete metric space when taking the metric ultraproduct with respect to a countably incomplete ultrafilter.

Proposition 2.15. *Suppose $(X_i, d_i, \star_i)_{i\in I}$ is a sequence of pointed metric spaces indexed by I. Then if $\mathcal{U}$ is a countably incomplete ultrafilter on I (in particular, $\mathcal{U}$ must be nonprincipal, hence I must be infinite), then $(X, d, \star) := \prod_{i\in I}(X_i, d_i, \star_i)/\mathcal{U}$ is complete.*

Proof. Fix a fast Cauchy sequence $(x_n)_{n\in\mathbb{N}}$ from X. That is, for each $n \in \mathbb{N}$, we have that $d(x_n, x_{n+1}) < 2^{-n}$. We shall assume that each x_n is given by some representative sequence $x_n := (x_n^i)_{i\in I}$. Note that this does *not* imply that, for each $i \in I$, the sequence $(x_n^i)_{n\in\mathbb{N}}$ is fast Cauchy, but if one fixes $F \subseteq \mathbb{N}$ a finite subset of indices, then for almost all $i \in I$ (with respect to $\mathcal{U}$), the sequence should be fast Cauchy "on F."

Now, because $\mathcal{U}$ is countably incomplete, by Fact 2.2, we may pick some descending sequence $(B_m)_{m\in\mathbb{N}}$ such that each $B_m \in \mathcal{U}$ yet $\bigcap_{m\in\mathbb{N}} B_m = \varnothing$. Now, define

$$A_m := \left\{i \in B_m : d(x_n^i, x_{n+1}^i) < 2^{-n} \text{ for } n = 0, \ldots, m\right\}.$$

It should be clear that each $A_m \in \mathcal{U}$, as it is the intersection of two sets belonging to $\mathcal{U}$: namely, the sets B_m and $\left\{i \in I : d(x_n^i, x_{n+1}^i) < 2^{-n} \text{ for } n = 0, \ldots, m\right\}$. We also have that $A_{m+1} \subseteq A_m$. We shall now construct a sequence $(y^i)_{i\in I}$ such that $\lim_{n\to\infty} x_n = [y^i]_{\mathcal{U}}$.

If $i \notin A_0$, then set $y^i := \star_i$. Otherwise, there is some maximal m such that $i \in A_m$ (because the countable intersection is empty), i.e. $i \in A_m \setminus A_{m+1}$. We may thus define $y^i := x_{m+1}^i$, hence $d(x_m^i, y^i) < 2^{-m}$.

Then, we note that, if $y^i = \star_i$, then $d(\star_i, y^i) = 0$, and if $y^i = x^i_m$, then

$$d(\star_i, y^i) \leqslant d(\star_i, x^i_0) + \sum_{k=0}^{m} d(x^i_k, x^i_{k+1}) \leqslant d(\star_i, x^i_0) + 2,$$

hence $\lim_{\mathcal{U}} d(\star_i, y^i) < \infty$ as $\lim_{\mathcal{U}} d(\star_i, x^i_0) < \infty$. In other words, $[y^i]_{\mathcal{U}} \in X$, and $[y^i]_{\mathcal{U}}$ is the limit of the x_n by construction. □

Definition 2.36. *Fix $\mathcal{L}$, I, $\mathcal{U}$, and $(\mathcal{M}_i)_{i\in I}$ as in the previous definition. We define $\mathcal{M}$, the* **ultraproduct** *of the $\mathcal{M}_i$, as follows:*

1. *For each sort symbol $S \in \mathcal{L}$, the sort M_S is the ultraproduct sort $\left(\prod_{i\in I} \mathcal{M}_i/\mathcal{U}\right)_S$ defined above.*

2. *For each function symbol $f \in \mathcal{L}$ with domain* $\mathrm{dom}(f) = (S_j)_{j\in J}$ *and range* $\mathrm{rng}(f) = S$, *we define*

$$f^{\mathcal{M}} : \prod_{j\in J} M_{S_j} \to M_S$$

by

$$f^{\mathcal{M}}\Big(\big([a_j]_{\rho_{S_j}}\big)_{j\in J}\Big) := \big[f^{\mathcal{M}_i}\big((a_j)_{j\in J}\big)\big]_{\rho_S}.$$

3. *Similarly, for each relation symbol $R \in \mathcal{L}$ with domain* $\mathrm{dom}(R) = (S_j)_{j\in J}$, *we define*

$$R^{\mathcal{M}} : \prod_{j\in J} M_{S_j} \to \mathbb{R}_{\geqslant 0}$$

by

$$R^{\mathcal{M}}\Big(\big([a_j]_{\rho_{S_j}}\big)_{j\in J}\Big) := \lim_{\mathcal{U}} R^{\mathcal{M}_i}\big((a_j)_{j\in J}\big).$$

We denote the ultraproduct of the $\mathcal{M}_i$ as $\prod_{i\in I} \mathcal{M}_i/\mathcal{U}$. If each $\mathcal{M}_i$ is the same $\mathcal{L}$-structure $\mathcal{N}$, then we refer to the above construction as an **ultrapower** *of $\mathcal{N}$ and denote it as $\mathcal{N}^{\mathcal{U}}$.*[15]

[15] The terminology follows from powers being nothing but repeated products.

The above notation is, truthfully, a bit confusing to parse, notably due to the fact that both ultraproducts and controlled functions are indexed over separate, arbitrary sets I and J. For clarity, we opted to work specifically with equivalence classes, thus effectively removing mention of the indexing set I from being printed in our definition. To elucidate the notation just a bit more, note that inside $f^{\mathcal{M}}(-)$ is a sequence indexed by J of elements from the ultraproduct sort. Each of these elements comes from some sort S_j and is, in fact, an equivalence class induced by the pseudometric ρ_{S_j} on that sort. Perhaps more intuitively, if one thinks for a moment only of unary functions on a single sort, then we are defining our functions according to the principal that "a function sends the equivalence class of a sequence to the equivalence class of where the function maps each index of the representative."

One ought to be concerned about this previous definition. Namely, we need to know that our functions and relations are well-defined! Part of that includes showing that our relations actually take on values in $\mathbb{R}_{\geqslant 0}$ rather than diverging to infinity, and similarly that our functions send bounded sequences to bounded sequences. This is prima facie unclear, because even though, for each $j \in J$, the sequence $((a_j)_i)$ must be bounded $\mathcal{U}$-almost everywhere, it is not immediate that the interpretations of, say, $R^{\mathcal{M}_i}$ must also be. We must also check that the interpretations of relations and functions are controlled by the same controllers, as this is a piece of our language $\mathcal{L}$!

The proof of this fact is straightforward, but neither quick nor clean. Much of the difficulty in the proof comes only from keeping track of multiple indices. We comment that the notion of controlled functions was essentially defined so that this proof goes through. We record this below.

Proposition 2.16. *The interpretations in Definition 2.36 are well-defined, and* $\prod_{i \in I} \mathcal{M}_i / \mathcal{U}$ *is an* $\mathcal{L}$*-structure.*

Proof. This is Proposition 3.7.3 of [29]. □

The definition of ultraproducts, though somewhat convoluted, intuitively tells us that to understand the behavior of an ultraproduct, one need only understand how "$\mathcal{U}$-almost all" of the individual structures *almost* behave. More precisely, if for each $\epsilon > 0$, $\mathcal{U}$-almost all believe that $d(f(a_i), f(b_i)) < \epsilon$, then in the ultraproduct, $f([a]) = f([b])$. The next theorem, which is the analog of Łoś' Theorem[16] for this version of continuous logic, tells us that this intuition extends beyond just the basic structure of the ultraproduct. In some sense, if one wants to understand the "truth" of an $\mathcal{L}$-formulae φ in an ultraproduct, it is enough to know that, for each $\epsilon > 0$, φ is ϵ-close to being satisfied in $\mathcal{U}$-almost all structures in the construction of the ultraproduct.

Theorem 2.2 (Fundamental Theorem of Ultraproducts)**.** *Let $\mathcal{M} = \prod_{i\in I} \mathcal{M}_i/\mathcal{U}$ be an ultraproduct of $\mathcal{L}$-structures, and let $\varphi((x_j)_{j\in J})$ be an $\mathcal{L}$-formula. Then, if $a := (a_j)_{j\in J}$ and $a^i := (a^i_j)_{j\in J}$ is such that, for each $j \in J$, $\lim_{\mathcal{U}} a^i_j = a_j$, then*

$$\varphi^{\mathcal{M}}(a) = \lim_{\mathcal{U}} \varphi^{\mathcal{M}_i}(a^i).$$

In particular, we recall that if φ is an $\mathcal{L}$-sentence, then $\varphi^{\mathcal{M}}$ and $\varphi^{\mathcal{M}_i}$ are technically constant real-valued functions, and by a slight abuse of notation, we have $\varphi^{\mathcal{M}} = \lim_{\mathcal{U}} \varphi^{\mathcal{M}_i}$.[17]

Proof. This is by induction on the construction of $\mathcal{L}$-formulae. Details may be found in Theorem 3.7.4 of [29]. □

One important consequence of Fundamental Theorem of Ultraproducts is the Compactness Theorem. In classical model theory, this tells us that, in order to know if a given $\mathcal{L}$-theory has a model, it is enough to know that any finite subset of it has a model. In continuous model theory, we only need to know that any finite subset is "ϵ-close" to having a model.

[16]This is a Polish name, and is pronounced [ˈwɔɕ], approximately as the English "wash."

[17]We say that this is only a slight abuse of notation because we choose to view these as functions from the empty product to $\mathbb{R}$, and because the empty product has a unique element, one may simply read $\varphi^{\mathcal{M}}$ as the value $\varphi^{\mathcal{M}}(a)$ where a is the unique element of the empty product.

Theorem 2.3 (Compactness)**.** *Let T be an $\mathcal{L}$-theory where T is a set and not a proper class. Then TFAE:*

1. *There is a model $\mathcal{M} \models T$.*
2. *For each finite subset $F \subseteq T$, there is a model $\mathcal{M} \models F$.*
3. *For each finite subset $F \subseteq T$, there is an $\mathcal{L}$-structure $\mathcal{M}$ such that, for each $\varphi \in F$, we have $\varphi^{\mathcal{M}} \leqslant \epsilon$.*

The proof is very similar in flavor to how one proves Compactness from Łoś in classical model theory. As in the classical case, the proof is surprisingly short, so we include it here.[18]

Proof. It is clear that (1) ⇒ (2) and (2) ⇒ (3). We now prove that (3) ⇒ (1). Define

$$I := \{(F, \epsilon) : F \subseteq T \text{ is finite}, \epsilon > 0\}.$$

By assumption, for each $(F, \epsilon) \in I$, there is some $\mathcal{M}_{(F,\epsilon)}$ such that for each $\varphi \in F$, $\varphi^{\mathcal{M}_i} \leqslant \epsilon$. For each $(F, \epsilon) \in I$, define

$$\uparrow (F, \epsilon) \downarrow := \{(F', \epsilon') \in I : F \subseteq F' \text{ and } \epsilon' \leqslant \epsilon\}.$$

It is easy to check that $J : \{\uparrow (F, \epsilon) \downarrow : (F, \epsilon) \in I)\}$ has the FIP, hence may be extended to an ultrafilter $\mathcal{U}$ on I. Then I claim that $\mathcal{M} := \prod_{i \in I} \mathcal{M}_i / \mathcal{U} \models T$. Indeed, for each $\varphi \in T$, we note that $\{\varphi\}$ is a finite subset, hence for each $\epsilon > 0$, $\uparrow (\{\varphi\}, \epsilon) \downarrow \in \mathcal{U}$. By construction, we see that $\varphi^{\mathcal{M}_i} \leqslant \epsilon$ for each $i \in \uparrow (\{\varphi\}, \epsilon) \downarrow$, so by the definition of ultralimits, we have $\lim_{\mathcal{U}} \varphi^{\mathcal{M}_i} = 0$. Thus, by the Fundamental Theorem of Ultraproducts, $\mathcal{M} \models \varphi$, and because φ was arbitrary, we see that $\mathcal{M} \models T$. □

[18] For reference, Luther's proof of the Fundamental Theorem of Ultraproducts spans a bit over 4 pages, single-spaced!

2.7 Saturation

Łoś's Theorem gives us a powerful tool for building new $\mathcal{L}$-structures out of old ones. These new structures are, in many cases, "richer" than their base structures. For example, if $\varphi(x)$ is an $\mathcal{L}$-formula, $\mathcal{M}$ an $\mathcal{L}$-structure, and for each $\epsilon > 0$, there is some $b_\epsilon \in M$ such that $\varphi^{\mathcal{M}}(b_\epsilon) < \epsilon$, then although we are not guaranteed some $b_0 \in M$ such that $\varphi^{\mathcal{M}}(b_0) = 0$, there *is* such a $b_0 \in \mathcal{M}$ when $\mathcal{M} = \mathcal{N}^{\mathcal{U}}$ where $\mathcal{U}$ is a countably incomplete nonprincipal ultrafilter!

This is a special instance of a more general phenomenon known as *saturation*, which we shall now investigate.

Definition 2.37. *Suppose that $\mathcal{M}_A$ is some $\mathcal{L}_A$-structure, and I is some indexing set. We define*

$$\operatorname{tp}_{\mathcal{M}}((b_i)_{i\in I}/A) := \left\{\varphi(\bar{x}) \text{ an } \mathcal{L}_A\text{-formula with} < \omega \text{ free variables from } I : \mathcal{M}_A \models \varphi(\bar{b})\right\}.$$

*We then say that p is an I-**type over** A if there is some $\mathcal{L}_A$-structure $\mathcal{M}$ and $(b_i)_{i\in I} \in \mathcal{M}$ such that $p \subseteq \operatorname{tp}_{\mathcal{M}}((b_i)_{i\in I}/A)$.*

This is notably different than Luther's definition in that he had allowed for arbitrary $\mathcal{L}_A$ formulae and did not require that they mention only finitely many free variables. This additional assumption is necessary for the proof of Proposition 2.17 and will allow us to give a different description of types in Proposition 2.18.

Of course, our interest in types comes when an indexed sequence of elements *other than* the above (a_i) from $\mathcal{M}$ satisfies the type. We shall typically use the following terminology when discussing the satisfaction of $\mathcal{L}$-formulae in a type.

Definition 2.38. *If p is an I-type over A and $\mathcal{N}$ is an $\mathcal{L}_A$ structure with $(b_i) \in \mathcal{N}$ such that, for each $\varphi((x_i)) \in p$, $\mathcal{N} \models \varphi((b_i))$, we shall say that the indexed sequence (b_i) is a* **witness**

to p, and that p is **realized in $\mathcal{N}$**.

Definition 2.39. *Given a class $\mathcal{K}$ of $\mathcal{L}_A$-structures, we say that p is an* **I-type in $\mathcal{K}$ over A** *if there is some $\mathcal{L}_A$-structure $\mathcal{M} \in \mathcal{K}$ and $(a_i)_{i\in I} \in M$ such that $p \subseteq \mathrm{tp}_{\mathcal{M}}((a_i)_{i\in I}/A)$. We denote the space of all I-types in $\mathcal{K}$ over A by $S_I^{\mathcal{K}}(A)$. If $\mathcal{K} = T$ for some $\mathcal{L}_A$-theory T, then we shall write $S_I^T(A)$. Similarly, if $\mathcal{K} = \mathrm{Th}(\mathcal{M}_A)$ for some $\mathcal{L}_A$-structure $\mathcal{M}_A$, then we write $S_I^{\mathcal{M}}(A)$.*

Given a type p over some $A \subseteq \mathcal{M}$, it is natural to wonder if p has a realization in $\mathcal{M}$ itself.[19]

Definition 2.40. *We say that an $\mathcal{L}_A$-structure $\mathcal{M}$ is* **κ-saturated** *if, whenever $A \subseteq M$ is such that $\mathrm{char}(A) < \kappa$, I is an indexing set with $|I| \leqslant \kappa$, and $p \in S_I^M(A)$, then p is realized in M.*

If we would like to know whether an $\mathcal{L}$-structure is κ-saturated for some κ, then we have the following proposition which tells us that, under a certain assumption, we only need to check the 1-types.

Proposition 2.17. *An $\mathcal{L}$-structure $\mathcal{M}$ is κ-saturated if and only if, for each $A \subseteq M$ such that $\mathrm{char}(A) < \kappa$, all types from $S_1^M(A)$ are realized in M.*

Proof. This is proved in [29] by a simple transfinite induction argument. We comment that his argument relies on the assumption that each formula in a type $p \in S_\kappa(C)$ may only mention finitely many free variables. □

Our definition of I-types over A is a bit unintuitive. Currently, it depends on the knowledge of some nebulous $\mathcal{L}_A$ structure $\mathcal{M}$ and a sequence $(a_i)_{i\in I} \in \mathcal{M}$. However, using ultraproducts, we can get a slightly different characterization of I-types.

[19] N.B.: Although we are defining our type over a particular subset of $\mathcal{M}$, we are not requiring that $p \subseteq \mathrm{tp}_{\mathcal{M}}((b_i)/A)$. Indeed, there may be some other $\mathcal{L}_A$ structure $\mathcal{N}$ such that $p \subseteq \mathrm{tp}_{\mathcal{N}}((b_i)/A)$.

First, we note that there is something fundamentally different about types in this version of continuous model theory. It is **not** true that if p is a set of $\mathcal{L}$-formulae such that every finite subset $F \subseteq p$ has a model, then p has a model. This may fail for the following reason:

Nonexample 2.1. *Let $\mathcal{L}$ be the empty language and any unbounded $\mathcal{L}$-structure $\mathcal{M}$. Then the set*

$$p := \{d(\star, x) \dot{\geqslant} r : r \in \mathbb{R}_+\}$$

is is finitely satisfiable, but is not satisfiable. Indeed, because $\mathcal{M}$ is unbounded, we note that any finite subset $F \subseteq p$ has some maximal r_F which is mentioned in one of the $\mathcal{L}$-formulae, and we may select some $b \in \mathcal{M}$ such that $d^{\mathcal{M}}(\star, b) > r_F$. However, p cannot be satisfiable in any $\mathcal{L}$-structure, as any individual point must have finite distance to $\star$.

It is reasonable to ask why we do not have to make any comment about sentences of the form $d(\star, c) \geqslant r$ for some constant c when stating and proving the Compactness Theorem. More precisely, why is it that in a language $\mathcal{L}$ with some constant symbol c, we do not obtain a contradiction by applying the Compactness Theorem to the set of $\mathcal{L}$-sentences $T := \{d(\star, c) \geqslant r : r \in \mathbb{R}_+\}$? The reason is subtle, but important. The fact is that the Compactness Theorem does **not** apply for the simple reason that the controller λ_c associated to the constant symbol c implies that $d(\star, c) \leqslant r$ for some $r \in \mathbb{R}_+$, hence not every finite subset of T is satisfiable.

However, the previous example captures the only hurdle to viewing I-types as finitely satisfiable $\mathcal{L}$-formulae.

Proposition 2.18. *Suppose that p is a set of $\mathcal{L}_A$-formulae such that each formula consists of only finitely many free variables among $(x_i)_{i \in I}$ and such that, for every x_i, there is a formula of the form "$d(\star_S, x_i) \dot{\leqslant} r_i$" $\in p$. Then p is an I-type over A if and only if, for every finite subset $F \subseteq p$ and $\epsilon > 0$, there is some $\mathcal{M}_{(F,\epsilon)}$ and $(b_i)_{i \in I} \in \mathcal{M}_{(F,\epsilon)}$ such that for every $\varphi((x_i)) \in F$, $\varphi^{\mathcal{M}_{(F,\epsilon)}}((b_i)_{i \in I}) \leqslant \epsilon$.*

Proof. The ($\Rightarrow$) direction is clear. Indeed, if p is an I-type over A, then by definition, there is some $\mathcal{M}$ and $(b_i)_{i\in I} \in \mathcal{M}$ such that $p \subseteq \mathrm{tp}_{\mathcal{M}}((b_i)_{i\in I}/A)$. Hence for each $\varphi((x_i)_{i\in I}) \in p$, we have $\varphi^{\mathcal{M}}((b_i)_{i\in I}) = 0 \leqslant \epsilon$.

We now prove the ($\Leftarrow$) direction. We shall find $\mathcal{M}$ and $(b_i)_{i\in I}$ such that $\mathcal{M} \models \varphi((b_i)_{i\in I})$ for every $\varphi \in p$. Indeed, if B is a set of new constant symbols indexed by I, then we may expand to a language $\mathcal{L}_{A\cup B}$ as follows:

For each $i \in I$, we must specify a controller λ_i corresponding to the symbol b_i. By assumption, there is some formula of the form $d(\star_S, x_i)\dot{\leqslant}r_i$, hence we may specify that λ_i be the constant r_i function. We extend our $M_{(F,\epsilon)}$ to an $\mathcal{L}_{A\cup B}$-structure $M'_{(F,\epsilon)}$ as follows: for each $i \in I$

- If $d(\star_S, b_i) \leqslant r_i$, then we are free to interpret b_i as usual within $M_{(F,\epsilon)}$.
- If not, then we interpret b_i as the base point.

In this way, it's possible that many of our $M'_{(F,\epsilon)}$ simply have several constants naming the base point! However, it is important to note that if there is any formula of the form $d(\star_S, x_i)\dot{\leqslant}s \in F$ for some $s \leqslant r_i$, then b_i is interpreted as normal.

Now, consider the $\mathcal{L}_{A\cup B}$-theory T which is constructed by replacing each instance of x_i with the constant symbol b_i for every formula $\varphi \in p$. We note that if we are able to find some $\mathcal{L}_{A\cup B}$-structure $\mathcal{M} \models T$, then by viewing $\mathcal{M}$ as an $\mathcal{L}_A$-structure, we may take the interpretations of the b_i and see that $p \subseteq \mathrm{tp}_{\Updownarrow}((b_i)_{i\in I}/A)$.

Unfortunately, because we have (potentially) changed our b_i when moving from $\mathcal{L}_A$-structures to $\mathcal{L}_{A\cup B}$-structures, we *cannot* say that $M_{(F,\epsilon)} \models F'$ where F' denotes the subset of T one obtains by changing instances of x_i to b_i for each formula in F. However, this does pose much of a problem, as the same flavor of proof for the Compactness Theorem still goes through, as we shall now make explicit.

For any finite $F \subseteq T$, we define

$$\bar{F} := F \cup \{d(\star_S, b_i)\dot{\leqslant} r_i : b_i \text{ appearing in some formula in } F\}\,.$$

We then set $I := \left\{(\bar{F}, \epsilon) : F \subseteq T \text{ is finite}, \epsilon > 0\right\}$. We note that in *this* instance, because $\bar{F}$ is finite, each $\mathcal{M}_{(\bar{F},\epsilon)}$ satisfies $\varphi^{\mathcal{M}_{(\bar{F},\epsilon)}} \leqslant \epsilon$ for every $\varphi \in F$. Thus, by the same argument used in the proof of Compactness, we find an ultraproduct $\mathcal{M}$ indexed by I such that $\mathcal{M} \models T$, and we win. □

We can thus use the previous proposition to show that, for saturated structures, $\inf_x |_r^s$ behaves as $\exists$:

Corollary 2.3. *Suppose that $\varphi(\bar{x}; (y_i)_{i\in I})$ is an $\mathcal{L}$-formula, $\mathcal{M}$ a saturated $\mathcal{L}$-structure, and $(a_i)_{i\in I} \in \mathcal{M}$. Then $\mathcal{M} \models \inf_{\bar{x}} |_{\bar{r}}^{\bar{s}} \varphi(\bar{x}; (a_i)_{i\in I})$ if and only if there is $\bar{b} \in B_{\bar{r}}(\mathcal{M})$ such that $\mathcal{M} \models \varphi(\bar{b}; (a_i)_{i\in I})$.*

Proof. One direction is obvious. For the other direction, let p be the set

$$p := \{\varphi(\bar{x}; (a_i)_{i\in I})\dot{\leqslant} 1/n : n \in \mathbb{N}_+\} \cup \{d(\star_S, x_i) \leqslant r_i\}\,.$$

We note that each formula in p mentions only finitely many free variables. It is finitely satisfiable by Proposition 2.6, hence is an I-type over A by Proposition 2.18, and by saturation of $\mathcal{M}$, has a witness in $\mathcal{M}$. Clearly, any witness to p will be a tuple $\bar{b}$ such that $\mathcal{M} \models \varphi(\bar{b}; (a_i)_{i\in I})$. □

It follows that Example 2.1 is not saturated, hence is not expressible as an ultraproduct of $\mathcal{L}$-structures over some nonprincipal ultrafilter. We also note that the full saturation of $\mathcal{M}$ is only necessary to handle arbitrary sequences of parameters within our $\mathcal{L}$-formula.

For example, if I is countable, then the previous proof still goes through assuming only $\aleph_1$-saturation of $\mathcal{M}$.

2.8 Stability

Our last model theoretic notion we shall investigate has to do with the size of the type space. In classical model theory, we declare that a theory is κ-stable if, whenever $\mathcal{M} \models T$ and $A \subseteq \mathcal{M}$ is such that $|A| = \kappa$, then $|S_1^{\mathcal{M}}(A)| = \kappa$. In some sense, any model of the theory T cannot have "too many" types.

We cannot expect a perfect translation to the continuous setting. For example, one should not work with the cardinality of a subset of a given metric space, but rather with its character density. In this way, we might expect that a κ-stable theory T should state that, if $\mathcal{M} \models T$ and $A \subseteq M$ is such that $\mathrm{char}(A) < \kappa$, then $\mathrm{char}(S_1^{\mathcal{M}}(A)) < \kappa$. This, of course, means that we need to specify a metric on the type space $S_1^T(A)$.

Definition 2.41. *Let T be an $\mathcal{L}_A$-theory. We define the d-metric on $S_I^T(A)$ by*

$$d(p, q) := \inf\left\{d^{\mathcal{M}}((a_i), (b_i)) : \mathcal{M} \models T, (a_i) \textit{ realizes } p, (b_i) \textit{ realizes } q\right\}.$$

This leads to our continuous analog of stability:

Definition 2.42. *Let T be an $\mathcal{L}$-theory and fix κ an infinite cardinal. We say that T is* κ**-stable** *if for any $\mathcal{M} \models T$ and $A \subseteq \mathcal{M}$ such that* $\mathrm{char}(A) < \kappa$, *then* $\mathrm{char}(S_1^T(A)) < \kappa$. *If there is some κ for which T is κ-stable, we shall say that T is* **stable***; else it is* **unstable**.

As in classical model theory, being unstable is equivalent to the exhibiting the *order property.*

Definition 2.43. *Let T be an $\mathcal{L}$-theory. We say that T has the* **order property** *if there is an $\mathcal{L}$-formula $\varphi(\bar{x}, \bar{y})$ (each of the same finite length), a model $\mathcal{M} \models T$, and a bounded* *sequence of tuples* $(\bar{a}_n)_{n\in\mathbb{N}}$ *such that*

$$\varphi^{\mathcal{M}}(\bar{a}_i, \bar{a}_j) = \begin{cases} 0 & \textit{if } i < j \\ 1 & \textit{else.} \end{cases}$$

Proposition 2.19. *An* $\mathcal{L}$*-theory is unstable if and only if it has the order property.*

Proof. This is Proposition 4.0.5 of [29]. □

Recall that our notion of ultraproduct was dependent on a choice of nonprincipal ultrafilter $\mathcal{U}$ over I. It is natural to wonder if the ultraproduct depends on the choice of $\mathcal{U}$. In fact, stability theory shows us that, depending on the truth of the Continuum Hypothesis, one may have several different ultraproducts as we vary the choice of $\mathcal{U}$!

Indeed, under the failure of CH, the unstable structures give rise to $2^{2^{\aleph_0}}$-many ultrapowers as we vary the choice of ultrafilter $\mathcal{U}$. Something similar can be stated for ultraproducts, although it will benefit us to refine our notion "exhibiting the order property" when we are discussing a sequence $(\mathcal{M}_n)_{n\in\mathbb{N}}$ of $\mathcal{L}$-structures, rather fixing a single such structure. The following definition and theorem are restated from [13].

Definition 2.44. *We say that a sequence* $(\mathcal{M}_n)_{n\in\mathbb{N}}$ *of* $\mathcal{L}$*-structures* ***approaches the order property*** *if there is an* $\mathcal{L}$*-formula* $\varphi(\bar{x}, \bar{y})$ *(where* $\bar{x}$*,* $\bar{y}$ *are of the same finite length) and a real tuple* $\bar{r}$ *of the same length such that, for every* $\epsilon > 0$ *and* $k \in \mathbb{N}$*, for all but finitely many of the* $\mathcal{M}_n$*, there are* $\bar{a}_1, \ldots, \bar{a}_k \in B_{\bar{r}}(\mathcal{M}_n)$ *such that*

$$\varphi^{\mathcal{M}_n}(\bar{a}_i, \bar{a}_j) : \begin{cases} \leqslant \epsilon & \textit{when } i < j \leqslant k \\ \geqslant 1 - \epsilon & \textit{when } j \leqslant i \leqslant k. \end{cases}$$

Theorem 2.4. *Assume CH is false. Let* $(\mathcal{M}_n)_{n\in\mathbb{N}}$ *be a sequence of* $\mathcal{L}$*-structures which approaches the order property and such that* $|M_n| \leqslant 2^{\aleph_0}$*. Then there are* $2^{2^{\aleph_0}}$*-many mutually* *non-isometric ultraproducts* $\prod_{n\in\mathbb{N}} \mathcal{M}_n/\mathcal{U}$ *as one varies the choice of nonprincipal ultrafilter* $\mathcal{U}$ *on* $\mathbb{N}$.

Proof. This is Theorem 6.1 from [13]. □

Of course, even though the ultraproducts may be non-isometric, it does not follow that they are pairwise non elementarily equivalent. Indeed, we shall later see an example of a class of ultraproducts that, under the failure of CH, are pairwise non-homeomorphic, yet are all elementarily equivalent. We shall now shift our focus to the topic of geometric group theory, where we shall encounter natural occurrences of ultraproducts.

Chapter 3

Geometric Group Theory

We now seek to use the model theory we've seen in the previous sections to investigate questions from geometric group theory. The guiding principle of geometric group theory is to study groups by studying the spaces upon which they act. As mentioned in the introduction, if one has a particular group presentation for G, then one can define the *Cayley graph* with respect to this metric, which gives rise to a metric space upon which G acts by left multiplication. Typically, one searches for "asymptotic" invariants, as group presentations are not unique, but they are quasi-isometric when we restrict our focus to finitely generated groups. This makes the use of asymptotic cones and the study of properties preserved between asymptotic cones of interest.

3.1 The Cayley Graph

Given a group G with presentation $\langle S|R\rangle$, one can define the Cayley graph for the presentation. This is typically viewed as a directed graph upon which G acts in a very natural way, although we shall work with it as an undirected graph. Upon assigning a weight of 1 to each edge of the graph and taking the distance between vertices to be the length of a shortest path between them, we obtain a very natural choice of metric space to investigate. We first recall some basic definitions:

Definition 3.1. *Given a group G and a set $S \subseteq G$, we say that G is* **generated by** *S if every $g \in G$ can be written as a finite product $g = s_1 s_2 \cdots s_n$ where, for each $i = 1, \ldots, n$, $s_i \in S$ or $s_i^{-1} \in S$. We shall denote by F_S the free group generated by the elements of S.*

Definition 3.2. *Given a set S, we say that a* **word** *w* **from** *S is any expression of the form*

$$w = s_1^{\epsilon_1} \cdots s_k^{\epsilon_k}$$

where each $s_i \in S$ and $\epsilon_i \in \{\pm 1\}$.[1] *We say that the* **length** *of w is k. We allow for $k = 0$, and we refer to the unique word of length 0 as the "empty word."*

Definition 3.3. *We say that G has* **presentation** *$G = \langle S | R \rangle$ if G is generated by S, R is a set of words from S, and N is the smallest normal subgroup of F_S containing R, then*

$$G \cong F_S / N.$$

If G is already fixed and is generated by S, we shall write $G = \langle S \rangle$ to note that we are considering G as a group generated by S. We shall not specify R in this case.

The above definition is nice in that it is quick to state once one has an understanding of words and normal subgroups. However, it may seem rather abstruse to someone first learning group theory. We offer an alternate approach to group presentations that is more algorithmic in nature.

[1] Intuitively, we're looking at finite formal products of elements from S or their inverses.

Definition 3.4. *Given a set S and R a set of words from S, we say that two words w, w' from S are* R**-equivalent** *if we may transform w into w' using only the following moves:*

- *adding or removing ss^{-1} or $s^{-1}s$ for some $s \in S$, or*
- *adding or removing w for some word $w \in R$.*

Intuitively, the only allowable moves are to add or remove a formal expression which we should think of as equivalent to the identity element of the group. This is often how computations are in particular groups in, say, a first semester course on the subject.

Example 3.1. *The group D_n can be given the presentation $\langle s, r | s^2, r^n, (sr)^2 \rangle$. Given this presentation, we can show that, in D_8, the word s^3r^3sr is R-equivalent to r^2.*

$$s^3r^3sr \rightsquigarrow sr^3sr \rightsquigarrow sr^2s^2rsr \rightsquigarrow sr^2s \rightsquigarrow sr^2sr^4 \rightsquigarrow srs^2rsr^4 \rightsquigarrow srsr^3 \rightsquigarrow r^2$$

In fact, one can work purely with the group D_n using only this presentation and R-equivalence of words, in the following sense:

Proposition 3.1. *G has presentation $G = \langle S|R \rangle$ iff G is isomorphic to the group whose elements are R-equivalence classes of words from S, and whose group operation is concatenation.*

Given a group presentation, one can also put a metric on the R-equivalence classes of words using what's called the *word metric*.

Definition 3.5. *Let G be a group with presentation $\langle S|R \rangle$. We define the* **word norm relative to** $\langle S|R \rangle$ *of an element $g \in G$ to be the length of a shortest word w that is R-equivalent to g, and we denote this by $|g|$.*

We then define the **word metric relative to** $\langle S|R \rangle$ *to be $d_{\langle S|R \rangle}(g, h) := |gh^{-1}|$.*

It is a simple exercise to show that the word metric is indeed a metric. We shall now introduce the Cayley graph of a group presentation and see that there is a close relationship between this object and the word metric for the group.

Definition 3.6. *Given a group G with presentation $G = \langle S|R\rangle$, we define the* ***undirected*** **Cayley graph** *of $\langle S|R\rangle$ to be the graph whose vertices are elements of G, and two vertices $g, h \in G$ have an edge between them if and only if $gh^{-1} \in S \setminus \{e\}$* or *$hg^{-1} \in S \setminus \{e\}$.*

We draw a few examples for different group presentations.

Example 3.2. *The dihedral group $D_n = \langle s, r|s^2, r^n, (sr)^2\rangle$ has the following Cayley graph in the case $n = 6$:*

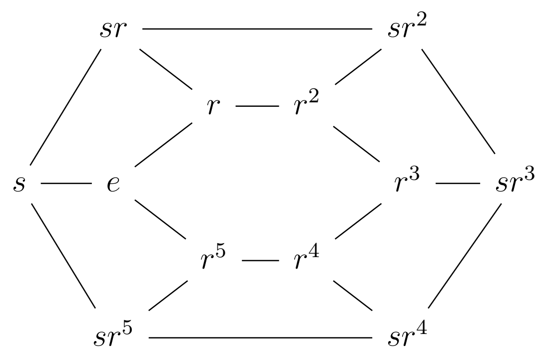

Example 3.3. *The group $\mathbb{Z}$ with presentation $\langle 1|\varnothing\rangle$ has the following Cayley graph:*

$$\cdots \;\text{——}\; -2 \;\text{——}\; -1 \;\text{——}\; 0 \;\text{——}\; 1 \;\text{——}\; 2 \;\text{——}\; \cdots$$

Example 3.4. *The group $\mathbb{Z}^2$ with presentation $\langle a, b|aba^{-1}b^{-1}\rangle$ has the following Cayley graph:*

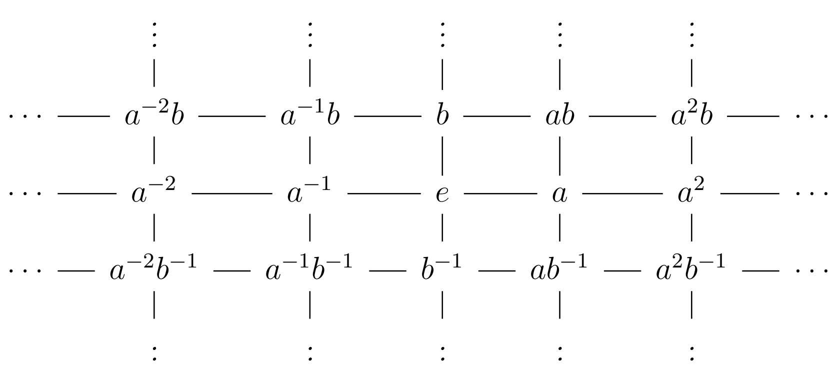

It's common to view a connected graph as a metric space, endowed with the so-called *path metric*; that is, the distance between vertices is given by the length of a shortest path between them. In fact, this is exactly the same as the word metric!

Proposition 3.2. *Let G be a group with presentation $\langle S|R\rangle$. Then the path metric on the Cayley graph coincides with the word metric.*

Proof. Consider a path $(v_1, v_2, \ldots, v_n)$ from g to h of minimal length. I.e. $v_1 = g$, $v_n = h$, and for each $i = 1, \ldots, n-1$, we have $v_i E v_{i+1}$. Then in particular $v_i v_{i+1}^{-1} \in S \setminus \{e\}$. Thus, we see that

$$gh^{-1} = (v_1 v_2^{-1})(v_2 v_3^{-1}) \cdots (v_{n-1} v_n^{-1})$$

gives us a word from S of length n. This shows that, if there is a path from g to h, then $|gh^{-1}| \leqslant d(g,h)$, where $d(g,h)$ denotes the length of the shortest path connecting g and h.

For the other direction, suppose that $gh^{-1} = s_1^{\epsilon_1} \cdots s_n^{\epsilon_n}$ is some word from S of length $n = |gh^{-1}|$. Then in particular, we see that

$$(g, s_1^{-\epsilon_1} g, s_2^{-\epsilon_2} s_1^{-\epsilon_1} g, \ \ldots \ , s_n^{-\epsilon_n} \cdots s_1^{-\epsilon_1} g)$$

is a path in the Cayley graph from g to h. This shows that $|gh^{-1}| \geqslant d(g,h)$, and thus we see that $d(g,h) = |gh^{-1}|$. □

It is worthwhile to note that the Cayley graph of any group presentation is path-connected. This is essentially by the previous proposition.

3.2 Quasi-Isometries

We now turn to the topic of quasi-isometries. These are meant to capture the notion of spaces having the same "large scale" geometry, in a way that is made precise later. One of first observations towards this end is that a finitely generated group G ought to have the same "large scale" geometry regardless of the choice of its generating set.

However, it is important to note that for a given group, the word metric is dependent on the choice of generating set. For example, we could have chosen a different generating set for $\mathbb{Z}$:

Example 3.5. *The group $\mathbb{Z}$ generated by $S = \{1, 2\}$ has the following Cayley graph:*

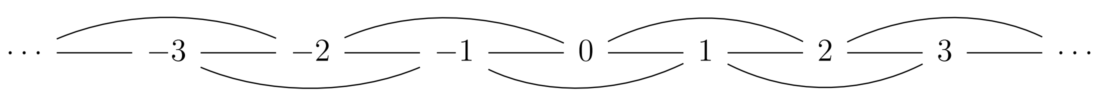

In this case, the word metric is given by

$$d_{\mathbb{Z},\{1,2\}}(n, m) = \begin{cases} \frac{1}{2}|n - m| & \text{if } n \equiv m \mod 2 \\ \frac{1}{2}|n - m - 1| + 1 & \text{if } n \not\equiv m \mod 2. \end{cases}$$

In the above example, we note that the word metric relative to this presentation is distinct from the metric in Example 3.3. However, it is not too far off. Namely, if we let $d_{\mathbb{Z},\{1\}}$ denote the word metric induced by the presentation $\mathbb{Z} = \langle 1|\varnothing\rangle$, then for any $n, m \in \mathbb{Z}$, we have the following:

$$\frac{1}{2}d_{\mathbb{Z},\{1\}}(n, m) - 1 \leqslant d_{\mathbb{Z},\{1,2\}}(n, m) \leqslant \frac{1}{2}d_{\mathbb{Z},\{1\}}(n, m) + 1.$$

In fact, this is not a coincidence; given any two finite presentations of a group, then they are *quasi-isometric* as metric spaces. This will be a useful fact later on when we investigate asymptotic cones.

Definition 3.7. *Let (X, d) and (Y, d') be metric spaces. We say that $f : X \to Y$ is an*

(A, B)**-quasi-isometry** *if* $A \geqslant 1$ *and* $B \geqslant 0$ *are such that*

- *for all* $x, y \in X$

$$\frac{1}{A}d(x,y) - B \leqslant d'(f(x), f(y)) \leqslant Ad(x,y) + B, \text{ and}$$

- *for each* $y \in Y$*, there is some* $x \in X$ *such that* $d(f(x), y) \leqslant B$*.*

If there exists a quasi-isometry $f : (X, d) \to (Y, d')$*, we shall say that they are* **quasi-isometric** *and shall write* $X \sim_{QI} Y$*.*

Some authors define a quasi-isometry in terms of a tuple (A, B, C), where $C \geqslant 0$ replaces B in the second condition. Of course, if one has an (A, B, C)-quasi-isometry in this sense, it will also be an $(A, \max\{B, C\})$-quasi-isometry in our sense, so not much is lost.

Proposition 3.3. $\sim_{QI}$ *is an equivalence relation.*

Proof. **(Reflexive)**: Obvious.

(Symmetric): Suppose $f : X \to Y$ is an (A, B)-quasi-isometry. We need to find some $g : Y \to X$ that is a quasi-isometry (though it need not be an (A, B)-quasi-isometry). For each $y \in Y$, choose some $x \in X$ such that $d'(f(x), y) \leqslant B$ and define $g(y) := x$. I claim that g is a quasi-isometry.

Fix $y, y' \in Y$, and note that

$$d'(f(g(y)), f(g(y'))) \leqslant d'(y, y') + 2B$$

by triangle inequality, and also

$$\frac{1}{A}d(g(y), g(y')) - B \leqslant d'(f(g(y)), f(g(y')))$$

because f is an (A, B)-quasi-isometry. Together, this implies

$$d(g(y), g(y')) \leqslant Ad'(y, y') + 3AB.$$

Similarly, because

$$d'(y, y') \leqslant d'(f(g(y)), f(g(y'))) + 2B$$

and

$$d'(f(g(y)), f(g(y'))) \leqslant Ad(g(y), g(y')) + B$$

we see that

$$\frac{1}{A}d'(y, y') - \frac{3B}{A} \leqslant d(g(y), g(y')).$$

Lastly, we note that if $x \in X$, then there is some $y \in Y$ such that $d'(f(x), y) \leqslant B$, and thus

$$\begin{aligned} d(x, g(y)) \leqslant & \quad d(x, g(f(x))) + d(g(f(x)), g(y)) \\ \leqslant & \quad B + Ad'(f(x), y) + 2AB + B \\ \leqslant & \quad 2B + 3AB. \end{aligned}$$

Thus, g is an (A, B')-quasi-isometry where $B' = 2B + 3AB$.

(Transitive): Finally, we prove that the composition of two quasi-isometries is a quasi-isometry, which will demonstrate transitivity. Let $f : (X, d) \to (Y, d')$ and $g : (Y, d') \to (Z, d'')$ be (A, B) and (A', B')-quasi-isometries. Then for any $x, y \in X$, we have

$$d''(g(f(x)), g(f(y))) \leqslant A'd'(f(x), f(y)) + B' \leqslant AA'd(x, y) + A'B + B'.$$

Similarly, we have

$$d''(g(f(x)), g(f(y))) \geqslant \frac{1}{A'}d'(f(x), f(y)) - B' \geqslant \frac{1}{AA'}d(x, y) - \frac{B'}{A} - B.$$

Finally, if $z \in Z$, then there is some $y \in Y$ such that $d''(g(y), z) \leqslant B'$ and some $x \in X$ such that $d'(f(x), y) \leqslant B$, hence

$$\begin{aligned} d''(g(f(x)), z) &\leqslant d''(g(f(x)), g(y)) + d''(g(y), z) \\ &\leqslant (A'd'(f(x), y) + A'B') + B' \\ &\leqslant A'B + A'B' + B'. \end{aligned}$$

□

We now record the following useful fact.

Proposition 3.4. *Let G be a group with presentations $\langle S|R \rangle$ and $\langle S'|R' \rangle$ where S, S' are finite. Denote the respective word metrics by $d_S, d_{S'}$. Then the identity map $i : (G, d_S) \to (G, d_{S'})$ is bi-Lipschitz.*[2]

The proof is relatively straightforward, and we note that we may find a sufficient A by $A := \max\{d_{S'}(e, s) | s \in S \cup S^{-1}\}$ where $S^{-1} := \{s^{-1} : s \in S\}$. Details may be found in [8].

We note, however, that the requirement that the presentations are finitely generated is a necessary one. A simple counterexample is the following:

Example 3.6. *Consider the group $\mathbb{Z}$ with generating sets $S = \{1\}$, $S' = \mathbb{Z} \smallsetminus \{0\}$, and induced word metrics $d_S, d_{S'}$. Clearly, $d_{S'}$ is the discrete metric, and $d_S(x, y) = |x - y|$. Suppose that $f : (\mathbb{Z}, d_S) \to (\mathbb{Z}, d_{S'})$ were an (A, B)-quasi-isometry. Then for any $x, y \in \mathbb{Z}$, we have*

$$\frac{1}{A}|x - y| - B \leqslant d_{S'}(f(x), f(y)) \leqslant A|x - y| + B$$

But, if we then take $y = 0$ and $x = A(B + 2)$, this would imply that $2 \leqslant d_{S'}(f(x), f(y))$,

[2] N.B.: A map $f : X \to Y$ is said to be A-bi-Lipschitz if for every $x, x' \in X$, we have that $\frac{1}{A} d_Y(f(x), f(x')) \leqslant d_X(x, y) \leqslant A d_Y(f(x), f(x'))$. This is *almost* an $(A, 0)$-quasi-isometry except that one need not require the images of bi-Lipschitz maps to be B-dense in their codomain. For example, an isometric embedding of $\mathbb{R}$ into $\mathbb{R}^2$ would be bi-Lipschitz but not quasi-isometric.

which is clearly impossible.

As geometric group theorists, our goal is to understand what information we may learn from how a group G acts on a space. In particular, if we are concerned with the large-scale geometry of G, then how must G act on a space X in order to conclude that G and X are quasi-isometric?

This question leads us to what is sometimes referred to as the "Fundamental Observation of Geometric Group Theory," or the Milnor-Schwarz Lemma. In order to state this precisely, we need a few definitions to understand how a group G may act on a particular metric space.

Definition 3.8. *Let G be a group which acts on a proper geodesic metric space (X, d) by isometries. We say that this action is* **properly discontinuous** *if, for each compact $K \subseteq X$, we have that $\{g \in G : gK \cap K \neq \varnothing\}$ is finite.*

Properly discontinuous actions tell us that, for any compact set K, most elements of G move K away from where it began. A simple example of a properly discontinuous action would be to consider $\mathbb{Z}$ acting on $\mathbb{R}$ by addition. Of course, compact subsets $K \subseteq \mathbb{R}$ are bounded, hence contained in some interval $K \subseteq [-R, R]$. But this implies that for all integers $n \in \mathbb{Z}$ such that $n > 2R$ or $n < -2R$, we have $nK \cap K = \varnothing$, hence this action is properly discontinuous. Not only that, but the orbit of any point (nearly) covers $\mathbb{R}$ in the following sense:

Definition 3.9. *Let G act on a proper geodesic metric space (X, d) by isometries. We say that this action is* **cocompact** *if, for any $x_0 \in X$, there is some $R > 0$ such that, for any $x \in X$, there is a $g \in G$ such that $x \in B_R(gx_0)$.*

Put another way, an action of G on X by isometries is cocompact if and only if, for any x_0, there is some $R > 0$ such that the orbit of x_0 is R-dense in X. In the case of $\mathbb{Z}$ acting on $\mathbb{R}$, this can be seen by taking $R = 1$ regardless of x_0, as for any real $x \in \mathbb{R}$, there is some $n \in \mathbb{Z}$

such that $|n + x_0 - x| < 1$. If G acts cocompactly on a proper geodesic metric space (X, d), then although we have not specified a generating set for G (hence we have not specified a word metric), one should maybe guess that "G is R-dense in X," whatever that might mean. Indeed, this is captured by fundamental observation of geometric group theory!

Theorem 3.1 (Milnor-Schwarz Lemma)**.** *Let G be a group and (X, d) a proper geodesic metric space on which G acts properly discontinuously and cocompactly.*[3] *Then G is finitely generated and is quasi-isometric to X.*

To be precise, the lemma proves that G has some finite presentation, and relative to this presentation, one we equip G with its word metric, we have $G \sim_{QI} X$. The proof itself is quite clever, and we present a sketch of it below. For a reference (including wonderful visuals), one may read [8].

Proof Sketch. We first need to find a generating set for G. We fix a base point x_0 and choose some R such that the orbit of x_0 is R-dense in X. Then, because X is proper, every closed ball $B_R(gx_0)$ is compact, and by the fact that G acts properly discontinuously on X, there are finitely many $g \in G \setminus \{e\}$ such that $B_R(gx_0) \cap B_R(x_0)$ is nonempty. We let S consist of these finitely many elements.

One can prove that S indeed generates G as follows:

We set $c := \inf\{d(x, gy) : x, y \in B, g \in G \setminus (S \cup \{e\})\}$ where $B := B_R(x_0)$ and one can see that $c > 0$, lest there be infinitely many g such that $d(B, gB)$ is bounded, which will contradict proper discontinuity.[4] We then pick any $g \in G \setminus (S \cup \{e\})$, and we note that $d(x_0, gx_0) \geqslant R + c$ because intuitively, one must first travel distance R from x_0 to reach the

[3] N.B.: In particular, G acts on X by isometries.

[4] Note that c is different than taking the minimal Hausdorff distance between B and gB, as the Hausdorff distance tells us the farthest one may be forced to travel from one set to another. Instead, we are looking for the smallest distance that one must travel from B to get to some gB.

boundary of $B_R(x_0)$, and then still must travel distance c to reach the boundary of $B_R(gx_0)$. It follows that there is $k \geqslant 2$ such that $R + (k-1)c \leqslant d(x_0, gx_0) < R + kc$.

Then, one picks a sequence of points $x_1, \ldots, x_{k+1}$ along a geodesic from x_0 to $x_{k+1} = gx_0$ such that $d(x_0, x_1) \leqslant R$ and $d(x_i, x_{i+1}) < c$, and in a similar vein to the first part of the proof of Proposition 3.2, each x_{i+1} corresponds to some $g_i \in G$ where $x_{i+1} \in g_i B$, and we'll have that the product $(g_0^{-1}g_1)(g_1^{-1}g_2)\cdots(g_{k-1}^{-1}g_k) = g$. It follows that each $g_i^{-1}g_{i+1} \in S$ because the distance between consecutive x_i was chosen to be smaller than c.

Finally, one can show that the map $g \mapsto gx_0$ is a quasi-isometry. In this way, the comment before the theorem about G that acts cocompactly on X being "R-dense in X" indeed follows! This does not immediately tell us that this map is a quasi-isometry, but one may consult the full proof to find explicit constants C, D such that

$$\frac{1}{C} d_S(g,h) - D \leqslant d(gx_0, hx_0) \leqslant \frac{1}{C} d_S(g,h) + D.$$

□

3.3 Hyperbolicity

One way of classifying groups is with the notion of δ-*hyperbolicity*. This has a few different characterizations, and we present a few of them here. Intuitively, hyperbolic spaces should be "negatively curved," although the characterization here does not rely on Riemannian curvature, and indeed may be stated for general metric spaces, rather than manifolds. One of the most common characterizations is through the notion of a δ-*thin* triangle:

Definition 3.10. *Let (X, d) be a geodesic metric space. We say that a triangle ABC is* ***δ-thin*** *for some $\delta \geqslant 0$ if the union of the δ-neighborhoods*[5] *of any two sides contains the*

[5] That is, the points which are distance at most δ from a point in the given side.

third side.

Of course, for any triangle in any geodesic space, one may find a δ such that the given triangle is δ-thin. Thus, our interest lies with those spaces for which one can find a $\delta \geqslant 0$ such that *all* triangles are δ-thin.

Definition 3.11. *Fix $\delta \geqslant 0$. We say that a geodesic metric space (X, d) is δ-**hyperbolic** if all triangles are δ-thin.*

The above definition has a nice visualization, as in Figure 3.1. There, the bubbles around the left and right sides are the δ-neighborhoods, and one can see that the third side on the bottom is fully contained in the union of these neighborhoods.

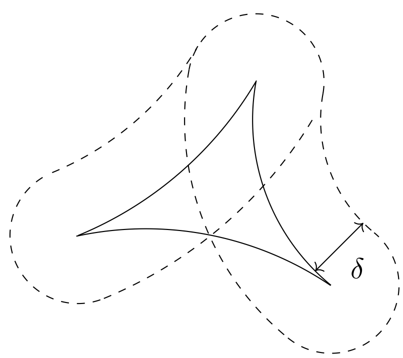

Figure 3.1: A sketch of a δ-thin triangle

It is worth noting that $\mathbb{R}^2$ is not δ-hyperbolic for *any* $\delta \geqslant 0$. If one fixes such a δ, we may simply consider an equilateral triangle with lengths $> \dfrac{4\delta}{\sqrt{3}}$. It is then worth asking if there are any spaces which are δ-hyperbolic for some $\delta \geqslant 0$. One should not be surprised to find that the hyperbolic plane is δ-hyperbolic.

Example 3.7. *The hyperbolic plane is δ-hyperbolic*

The above is not an immediate fact, but may be deduced by noting that all triangles in the hyperbolic plane live in some "ideal" triangle, and ideal triangles are all isometric and δ-thin for $\delta = \log(1+\sqrt{2})$. See [8] for details.

Our next example shall be of crucial importance later to our main results. Indeed, in our study of *buildings*, these shall appear as the underlying metric spaces to certain "parallel classes" of objects found within a building.

Definition 3.12. *We say that a metric space (X, d) is an $\mathbb{R}$-**tree** if there is a unique arc between any two points $x, y \in X$, and moreover, that unique arc is a geodesic. Formally, an arc from x to y is the image of an isometric embedding $\sigma : [a, b] \hookrightarrow X$ such that $\sigma(a) = x$ and $\sigma(b) = y$.*

To be clear, an $\mathbb{R}$-tree is more than a unique geodesic space. The fact that there is a unique *arc* between any two points means that there is only one *path* between two points that does not retrace itself. $\mathbb{R}^2$ equipped with the Paris metric[6] is an example of an $\mathbb{R}$-tree. We shall return to $\mathbb{R}$-trees a bit later.

While our definition of hyperbolicity is somewhat geometrically intuitive, it is perhaps harder to formalize in continuous logic, as the definition relies on some notion of quantifying over triangles. This is not impossible, but the following characterization of δ-hyperbolicity makes our lives as model theorists a bit easier.

Definition 3.13. *Let x, y, w be points in a metric space (X, d). Then we define the **Gromov product** $(x \cdot y)_w$ by*

$$(x \cdot y)_w := \frac{1}{2}(d(x, w) + d(w, y) - d(x, y)).$$

Intuitively, the Gromov product is half of how close the triangle inequality is to being an equality. In an $\mathbb{R}$-tree, the Gromov product will be the distance from w to the segment $[x, y]$. This allows us to give the following alternate characterization of δ-hyperbolic spaces.

[6] The author has also heard this referred to as the "Chicago suburb railway metric," as to get between any two suburbs on different lines, one must first take the train downtown to Union Station.

Definition 3.14. *We say that a metric space (X, d) satisfies the δ-**four point condition** if, for every $w, x, y, z \in X$, we have*

$$(x \cdot y)_w \geqslant \min\{(y \cdot z)_w, (x \cdot z)_w\} - \delta.$$

In the case where $\delta = 0$, the above condition is telling us that, for any quadruple of points w, x, y, z, there must be a "tie" for the smallest Gromov product with respect to w. This follows from noting that, if w is not in the triangle T formed by x, y, z, then there is a unique point of T which w is closest to, and this point must lie on at least two sides of T (by 0-hyperbolicity). We shall prove something slightly stronger in a moment which will help us understand when two $\mathbb{R}$-trees are elementarily equivalent.

We have the following well-known fact:

Fact 3.1. *(X, d) is δ-hyperbolic if and only if it satisfies the δ'-four point condition. More precisely, if (X, d) is δ-hyperbolic, then there is some $\delta' \geqslant 0$ such that (X, d) satisfies the δ'-four point condition, and if (X, d) satisfies the δ'-four point condition, there is some $\delta \geqslant 0$ such that (X, d) is δ-hyperbolic.*

To be extra clear, the δ and δ' above do not need to be the same. Of course, they certainly do not need to be the same for the simple reason that any δ-hyperbolic space is δ'-hyperbolic for $\delta' > \delta$, and similarly for the δ-four point condition. However, even if one only considers optimal δ, the two notions need not coincide. Thus, one is typically only concerned that a space is δ-hyperbolic for *some* $\delta > 0$. The case of $\delta = 0$ is privileged, as in this case, the two notions coincide, and, as we shall see in the next section, the 0-hyperbolic spaces arise from "zooming out" on any δ-hyperbolic space.

As mentioned, this characterization of δ-hyperbolicity is easier to work with as a model theorist, in the sense that we may axiomatize δ-hyperbolic spaces for a fixed δ.

Proposition 3.5. *Fix a language $\mathcal{L}$ and $\delta \geqslant 0$. The class of $\mathcal{L}$-structures satisfying the δ-four point condition is axiomatizable.*

Proof. We let $\psi(x,y,w)$ be the $\mathcal{L}$-formula given by the Gromov product; more precisely, we define $\psi(x,y,w) := \frac{1}{2}(d(x,w) + d(w,y) \dot{-} d(x,y))$. Then, for each $n \in \mathbb{N}_+$, let $\varphi_{n,\delta}$ be defined by

$$\varphi_{n,\delta} := \sup_{x,y,z,w} \Big|_{\bar{0}}^{\bar{n}} \psi(x,y,w) \dot{\geqslant} \min\left(\psi(y,z,w), \psi(x,z,w)\right) \dot{-} \delta.$$

Here, $\bar{0} = (0,0,0,0)$ and $\bar{n} = (n,n,n,n)$. We then define $T_\delta := \{\varphi_{n,\delta} : n \in \mathbb{N}_+\}$. Then, if $\mathcal{M}$ is an $\mathcal{L}$-structure, we see that if $\mathcal{M} \models T_\delta$, then by Proposition 2.6 and the fact that any quadruple of points must be contained in some ball of finite radius, $\mathcal{M}$ clearly satisfies the δ-four point condition. It should be clear that if $\mathcal{M}$ satisfies the δ-four point condition, then $\mathcal{M} \models T_\delta$. □

The caveat, of course, is that we may only axiomatize δ-hyperbolicity for a particular δ. In particular, if one fixes $\delta \geqslant 0$, then for any $\delta' > \delta$, one can find some (X,d) which is δ'-hyperbolic but not δ-hyperbolic.

$\mathbb{R}$-trees

We present a few basic facts about $\mathbb{R}$-trees, and we shall see an important example of a class of elementarily equivalent $\mathbb{R}$-trees as we mentioned before.

There are several equivalent definitions for $\mathbb{R}$-trees. We state a few of them here.

Fact 3.2. *Let (X,d) be a metric space. Then TFAE:*

1. *(X,d) is an $\mathbb{R}$-tree.*

2. *(X,d) is a 0-hyperbolic, connected geodesic space.*

3. (X, d) *is a geodesic space in which every triangle is a tripod (meaning that for any* $x, y, z \in X$*, the geodesics* $[x, y]$ *and* $[x, z]$ *intersect in a geodesic* $[x, c]$ *where* $c \in [y, z]$*).*

4. (X, d) *is a geodesic space with no subset homeomorphic to a circle.*

Proof. The equivalence of (1), (3), and (4) may be found in [7]. His definition of an $\mathbb{R}$-tree is close to (3) and is easily seen to be equivalent to it. The equivalence of (2) and (3) may be found in [12] where again, the definition of an $\mathbb{R}$-tree is equivalent to the definition used in [7]. □

Perhaps the last characterization gives the clearest reason as to why we refer to these as "trees," as this is the topological version of the graph theoretic characterization of trees. This also suggests the following standard terminology:

Definition 3.15. *Suppose that* (X, d) *is an* $\mathbb{R}$*-tree. For any point* $x \in X$*, we say that the* ***valency of*** x *is the number of connected components in* $X \setminus \{x\}$*. Each of these connected components is called a* ***branch at*** x*.*

The *bounded* continuous model theory of $\mathbb{R}$-trees has been studied in [6]. In particular, they axiomatize $\mathbb{R}$-trees of a bounded radius and identify its model companion. We do not give a full treatment of unbounded $\mathbb{R}$-trees in our continuous model theory, but we present a few choice results which will be of use to use in the next chapter.

Lemma 3.1. *Let* $\mathcal{L}$ *be the empty language, and suppose that* $\mathcal{L}$*-structures* $(X, d_X, \star_X)$ *and* $(Y, d_Y, \star_Y)$ *are* $\mathbb{R}$*-trees such that every* $x \in X$ *and* $y \in Y$ *have infinite valency, and each branch at* x *and* y *contains a ray. Then Player II has a perfect strategy for any atomic EF game between* X *and* Y*. In particular,* $X \equiv Y$*.*

Proof. We comment that we shall not make any mention of $r > 0$ which bounds the field of play. Indeed, this will not change Player II's strategy, as at each play, she will play exactly the same distance from $\star$ as Player I does. We shall prove a stronger statement by induction. Namely, for any points $a_1, \ldots, a_k \in X$ and $b_1, \ldots, b_k \in Y$, we shall define

$$X_k := \bigcup_{1 \leqslant i < j \leqslant k} [a_i, a_j]_X \cup [\star_X, a_i]_X \qquad \text{and} \qquad Y_k := \bigcup_{1 \leqslant i < j \leqslant k} [b_i, b_j]_Y \cup [\star, b_i]_Y.$$

Then we shall prove by induction that if there is a partial isometry $\sigma_k : X_k \to Y_k$ such that $\sigma_k(a_i) = b_i$ for each $i = 1, \ldots, k$, then for any $a_{k+1} \in X$, there is a $b_{k+1} \in Y$ and $\sigma_{k+1} : X_{k+1} \to Y_{k+1}$ extending σ_k such that $\sigma_{k+1}(a_{k+1}) = b_{k+1}$ (and vice versa if one considers b_{k+1} chosen from Y first).

On turn 1 of the EF game, if Player I plays, say, $a_1 \in X$, then it is trivial for Player II to respond with $b_1 \in Y$ such that $d_X(\star_X, a_1) = d_Y(\star_Y, b_1)$, and the isometry $\sigma_1 : X_1 \to Y_1$ is the obvious one.

Suppose that the induction hypothesis is true up to some k, and fix $a_{k+1} \in X$. The opposite case shall follow from a similar argument.

Let $\rho := d(a_{k+1}, X_k)$. I claim that there is a unique $\gamma \in X_k$ such that $d_X(\gamma, a_{k+1}) = \rho$. The fact that any such γ exists follows from the fact that X_k is compact (which can be easily seen using sequential compactness). Now suppose that there is some $\gamma' \in X_k$ such that $d_X(\gamma', a_{k+1}) = \rho$. Then by Fact 3.2, the triangle determined by γ, γ', a_{k+1} is a tripod, hence $[a_{k+1}, \gamma] \cap [a_{k+1}, \gamma'] = [a_{k+1}, c]$ for some $c \in [\gamma, \gamma']$. It is easy to see that X_k is path-connected, hence $[\gamma, \gamma'] \subseteq X_k$ and $c \in X_k$. Thus, since $c \in X_k \cap [a_{k+1}, \gamma]$, it follows that $c = \gamma$ and similarly $c = \gamma'$.

So, given γ such that $d_X(\gamma, a_{k+1}) = \rho$, we note that there are infinitely many branches at $\gamma' := \sigma_k(\gamma) \in Y$. Player II then selects b_{k+1} along any new branch such that $d_Y(b_{k+1}, \gamma') = \rho$, which is possible as each branch contains a ray. Then, if $\sigma' : [a_{k+1}, \gamma] \to [b_{k+1}, \gamma']$ is the isometry that sends $\gamma \mapsto \gamma'$, we may define

$$\sigma_{k+1} := \sigma_k \cup \{(x, \sigma'(x)) : x \in [a_{k+1}, \gamma]\}.$$

The fact that σ_{k+1} is an isometry follows from the fact that, for any $x \in [a_{k+1}, \gamma]$ and $y \in X_k$, we must have $\gamma \in [x, y]$ lest there be a loop, and similarly for the new branch in Y. Thus, $[x, y] = [x, \gamma] \cup [\gamma, y]$, and as each of these are geodesics, we have

$$d_X(x, y) = d_X(x, \gamma) + d_X(\gamma, y) = d_Y(\sigma_k(x), \sigma_k(\gamma)) + d_Y(\sigma_{k+1}(\gamma), \sigma_{k+1}(y))$$

$$= d_Y(\sigma_{k+1}(x), \sigma_{k+1}(y)).$$

Thus, because X_{k+1} and Y_{k+1} are isometric, it is trivial to see that Player II has a perfect strategy for any EF game. □

One could also have required that each branch at x contains points of arbitrary distance from x. Clearly, each branch containing a ray would imply that there are points of arbitrary distance from x, but the converse is not true in general. For example, one may have a connected component consisting only of a line segment which then branches into infinitely many branches each of length n. However, this violates the requirement that *every* point has infinite valency, and it is not hard to see that if every point $x \in X$ has infinite valency and each branch at x contains points of arbitrary distance from x, then each branch contains a ray.

The upshot of the previous lemma is that we did not specify the valency at each point, beyond requiring that it be infinite, and hence we have just demonstrated the elementary equivalence of a proper class worth of $\mathcal{L}$-structures! To be more precise, given some infinite cardinal κ, one may construct a *κ-universal $\mathbb{R}$-tree* T_κ which will satisfy the hypotheses of the previous lemma (one may consult [10] for details on the construction and how these specific

$\mathbb{R}$-trees fit into the next section). Clearly, if κ and λ are different infinite cardinals, then T_κ and T_λ must not be homeomorphic, yet are elementarily equivalent!

3.4 Asymptotic Cones

We now present one of the fundamental tools for geometric group theory. Asymptotic cones give a precise way to "zoom out" on pointed metric spaces, allowing one to study the "large scale behavior" of groups equipped with the word metric, for example. We present asymptotic cones as a particular kind of ultraproduct of metric spaces; this will allow us to use a few model theoretic tools to better understand some of the structure of asymptotic cones.

Definition 3.16. *Let* (X, d) *be a metric space and* $(\star_n)_{n\in\mathbb{N}_+}$ *a sequence of points from* X*,* $\mathcal{U}$ *a countably incomplete ultrafilter on* $\mathbb{N}_+$*, and* $(r_n)_{n\in\mathbb{N}}$ *a sequence of positive reals such that* $\lim_{\mathcal{U}} r_n = \infty$*. Then we say that the* **asymptotic cone of** X *with respect to this data is the metric ultraproduct*

$$\prod_{n\in\mathbb{N}_+} (X, \frac{d}{r_n}, \star_n)/\mathcal{U},$$

and we denote it as $\mathrm{Cone}_{\mathcal{U}}(X, (\star_n), (r_n))$*. If* $(\star_n)$ *is a constant sequence determined by some* $\star \in X$*, we shall write* $\mathrm{Cone}_{\mathcal{U}}(X, \star, (r_n))$*.*

Of course, as an asymptotic cone is a metric ultraproduct with respect to a countably incomplete ultrafilter, it must be complete as a metric space by Proposition 2.15. Indeed, one may always view an asymptotic cone as an $\mathcal{L}$-structure for $\mathcal{L}$ the empty language. We shall not require that the pointed metric spaces $(X, d/r_n, \star_n)$ to be $\mathcal{L}$-structures in general, but in most cases they will be. In particular, if our metric space X arises from a group G with its word metric, then X will be a discrete space, hence complete, and there's an obvious choice to let $\star := e$ the identity element from the group.

The next proposition, though, shows us that one need not focus on the choice of base point, as long as we select the same base point in each index.

Proposition 3.6. *For any $\star, \star' \in X$, we have that* $\mathrm{Cone}_{\mathcal{U}}(X, \star, (r_n)) = \mathrm{Cone}_{\mathcal{U}}(X, \star', (r_n))$. *In particular, these spaces are isometric.*

Proof. For the simplicity of notation, let $Y = \mathrm{Cone}_{\mathcal{U}}(X, \star, (r_n))$ and $Y' = \mathrm{Cone}_{\mathcal{U}}(X, \star', (r_n))$. Fix $[(a_n)], [(b_n)] \in Y$. We shall show that $\sigma : Y \to Y'$ defined by $[(a_n)]_Y \mapsto [(a_n)]_{Y'}$ is a well-defined isometry.

Indeed, the fact that this is an isometry is obvious. The only non-obvious part is that it is well-defined, and in particular that if $[(a_n)]_Y \in Y$, then the sequence $d(a_n, \star')/r_n$ has bounded ultralimit, hence $[(a_n)]_{Y'} \in Y'$. This follows from noting that

$$\lim_{\mathcal{U}} \frac{d(a_n, \star')}{r_n} \leqslant \lim_{\mathcal{U}} \frac{d(a_n, \star) + d(\star, \star')}{r_n} = \lim_{\mathcal{U}} \frac{d(a_n, \star)}{r_n}$$

as $r_n \to \infty$ in the ultralimit, hence $d(\star, \star')/r_n \to 0$. □

It's important to point out that, in the previous proposition, we needed to assume that our base points were unchanging in the construction of the asymptotic cone. However, if we know that our space is homogeneous, then one need not care about base points at all!

Proposition 3.7. *Suppose that X is homogeneous; i.e. for any $x, y \in X$, there is an isometry $f : X \to X$ such that $f(x) = y$. Then for any $(\star_n), (\star'_n) \in X$ and appropriate $\mathcal{U}$, $(r_n)_{n \in \mathbb{N}_+}$, we have that* $\mathrm{Cone}_{\mathcal{U}}(X, (\star_n), (r_n))$ *is isometric to* $\mathrm{Cone}_{\mathcal{U}}(X, (\star'_n), (r_n))$ *as pointed metric spaces.*

Proof. It is not hard to see that if (X, d) is homogeneous, then so is $(X, d/r)$ for any $r > 0$. Thus, for each $n \in \mathbb{N}_+$, we may find some isometry $f_n : (X, d/r_n) \to (X, d/r_n)$ such that

$f_n(\star_n) = \star'_n$. We then define $\tilde{f} : \text{Cone}_{\mathcal{U}}(X, (\star_n), (r_n)) \to \text{Cone}_{\mathcal{U}}(X, (\star'_n), (r_n))$ by

$$\tilde{f}([(a_n)]_{\mathcal{U}}) := [f_n(a_n)]_{\mathcal{U}}.$$

We show that $\tilde{f}$ is a well-defined isometry. Suppose that $[(a_n)]_{\mathcal{U}}, [(b_n)]_{\mathcal{U}} \in \text{Cone}(X, (\star_n), r_n)$. Then

$$d([(a_n)]_{\mathcal{U}}, [(b_n)]_{\mathcal{U}}) = \lim_{\mathcal{U}} \frac{d(a_n, b_n)}{r_n} = \lim_{\mathcal{U}} \frac{d(f_n(a_n), f_n(b_n))}{r_n} = d(\tilde{f}([(a_n)]_{\mathcal{U}}), \tilde{f}([(b_n)]_{\mathcal{U}})).$$

This not only shows that $\tilde{f}$ is an isometry, but also proves that $\tilde{f}$ is well-defined, because $[(a_n)]_{\mathcal{U}} = [(b_n)]_{\mathcal{U}}$ if and only if $d([(a_n)]_{\mathcal{U}}, [(b_n)]_{\mathcal{U}}) = 0$. □

Proposition 3.8. *Suppose that $(X, \star, d_X)$ and $(Y, \star_Y, d_Y)$ are pointed metric spaces, and fix appropriate $\mathcal{U}$ and (r_n). Then*

$$\text{Cone}_{\mathcal{U}}(X, \star_X, (r_n)) \times \text{Cone}_{\mathcal{U}}(Y, \star_Y, (r_n)) \cong \text{Cone}_{\mathcal{U}}(X \times Y, (\star_X, \star_Y), (r_n))$$

where the product metrics are induced by the p-norm.

Proof. Fix $[x_n], [x'_n] \in \text{Cone}_{\mathcal{U}}(X, \star_X, (r_n))$ and $[y_n], [y'_n] \in \text{Cone}_{\mathcal{U}}(Y, \star_Y, (r_n))$. Note that the metric on the LHS is given by

$$\left(\tilde{d}_X([x_n], [x'_n])^p + \tilde{d}_Y([y_n], [y'_n])^p\right)^{1/p}$$

where $\tilde{d}_X$ denotes the metric for the asymptotic cone of X, and similarly for Y. By expanding the ultralimit definition, we see that

$$\begin{aligned} \left(\tilde{d}_X([x_n], [x'_n])^p + \tilde{d}_Y([y_n], [y'_n])^p\right)^{1/p} &= \left(\left(\lim_{\mathcal{U}} \frac{d_X(x_n, x'_n)}{r_n}\right)^p + \left(\lim_{\mathcal{U}} \frac{d_Y(y_n, y'_n)}{r_n}\right)^p\right)^{1/p} \\ &= \lim_{\mathcal{U}} \left(\left(\frac{d_X(x_n, x'_n)}{r_n}\right)^p + \left(\frac{d_Y(y_n, y'_n)}{r_n}\right)^p\right)^{1/p} \end{aligned}$$

where the second equality follows from the fact that ultralimits commute with continuous functions. The result follows from the fact that this last expression is exactly the ultralimit of the product metric induced by the p-norm on $(X, d/r_n)$ and $(Y, d/r_n)$, hence is the metric for $\mathrm{Cone}_{\mathcal{U}}(X \times Y, (\star_X, \star_Y), (r_n))$. □

Thus, to understand an asymptotic cone of a finite product of spaces, it suffices to understand the cones of the factors. However, this does not translate to infinite products. Firstly, uncountable products of metric spaces need not even be metrizable. The following example demonstrates that, even in the case of a countable product of metric spaces, though one may still metrize the product topology, this need not play nicely with the asymptotic cone.

Example 3.8. *Let $(X_n, d_n, \star_n)_{n \in \mathbb{N}}$ be a sequence of metric spaces. We may metrize the product topology on $\prod X_n$ by, for any sequences $(x_n), (y_n) \in \prod X_n$, defining*

$$\tilde{d}((x_n), (y_n)) := \sum_{n=0}^{\infty} \frac{1}{2^n} \cdot \frac{d_n(x_n, y_n)}{1 + d(x_n, y_n)}.$$

However, one does not have an analogous result to Proposition 3.8 in this case. In fact, this makes the infinite product into a bounded metric space, hence its asymptotic cone is a point, as we shall see later.

We also shall introduce a few toy examples of asymptotic cones.

Example 3.9. *Consider the group $\mathbb{Z}^n$ with the taxicab metric d for some arbitrary $n \in \mathbb{N}_+$. Then for any ultrafilter $\mathcal{U}$ on $\mathbb{N}_+$ and $(r_n)_{n \in \mathbb{N}_+}$ of positive reals such that $\lim_{\mathcal{U}} r_n = \infty$,* $\mathrm{Cone}(\mathbb{Z}^n, 0, (r_n))$ *is isometric to $\mathbb{R}^n$ with the taxicab metric.*

Proof. By Proposition 3.8 with $p = 1$, it suffices to show that $\mathrm{Cone}_{\mathcal{U}}(\mathbb{Z}, 0, (r_n)) \cong \mathbb{R}$. Further, because $\mathbb{Z} \sim_{QI} \mathbb{R}$ (the inclusion map is a $(1, 1)$-quasi-isometry), it suffices to show that

$\text{Cone}_{\mathcal{U}}(\mathbb{R}, 0, (r_n)) \cong \mathbb{R}$. This is easy to see, as for any $\lambda \in \mathbb{R}$, we have that

$$\lim_{\mathcal{U}} \frac{d(\lambda r_n, 0)}{r_n} = \lambda,$$

hence we may define a map $\sigma : \mathbb{R} \to \text{Cone}_{\mathcal{U}}(\mathbb{R}, 0, (r_n))$ by $\sigma(\lambda) = [\lambda \cdot r_n]$. Then, we note that σ is an isometric embedding as

$$d(\sigma(\lambda), \sigma(\lambda')) = \lim_{\mathcal{U}} \frac{|\lambda r_n - \lambda' r_n|}{r_n} = |\lambda - \lambda'|$$

and it is also surjective, as for any $[x_n] \in \text{Cone}_{\mathcal{U}}(\mathbb{R}, 0, (r_n))$, we note that $\lambda := \lim_{\mathcal{U}} x_n \in \mathbb{R}$, hence $d(\sigma(\lambda), [x_n]) = 0$. □

Example 3.10. *If X is a bounded metric space, then its cone is a metric space with a single point. In particular, the cone of any compact space or space endowed with the discrete metric is a point.*

Proof. This follows directly from noting that, if X is bounded by M, then for any sequences $(x_n)_{n \in \mathbb{N}_+}$ and $(y_n)_{n \in \mathbb{N}_+}$, we have

$$\lim_{\mathcal{U}} d_n(x_n, y_n) \leqslant \lim_{\mathcal{U}} \frac{M}{r_n} \to 0.$$

□

We shall concern ourselves with asymptotic cones of groups. As asymptotic cones are meant to capture the "large-scale" geometry of a space, we should hope that this is an invariant of a f.g. group regardless of its presentation. Indeed, we saw in Proposition 3.4 that the identity map on a group with respect to two f.g. presentations is a quasi-isometry (and, in particular, is bi-Lipschitz). Indeed, two quasi-isometric spaces will always be bi-Lipschitz homeomorphic, as we see in the following proposition:

Proposition 3.9. *Let (X, d) and (Y, d') be metric spaces, $\star \in X$ some basepoint, $\mathcal{U}$ a nonprincipal ultrafilter on $\mathbb{N}_+$, and $(r_n)_{n\in\mathbb{N}_+}$ a sequence of positive reals such that $\lim_{\mathcal{U}} r_n = \infty$. Suppose that $f : (X, d) \to (Y, d')$ is an (A, B)-quasi-isometry. Then there exists $\tilde{f}$ an A-bi-Lipschitz homeomorphism*

$$\tilde{f} : \mathrm{Cone}_{\mathcal{U}}(X, \star, (r_n)) \to \mathrm{Cone}_{\mathcal{U}}(Y, f(\star), (r_n)).$$

Proof. We take a similar approach as in the proof of Proposition 3.7 and define our function $\tilde{f} : \mathrm{Cone}(X, \star, (r_n)) \to \mathrm{Cone}(Y, f(\star), (r_n))$ by

$$\tilde{f}([(a_n)]_{\mathcal{U}}) = [(f(a_n))]_{\mathcal{U}}.$$

To see that $\tilde{f}$ is a well-defined A-bi-Lipschitz homeomorphism, note that for any $[(a_n)], [(b_n)] \in \mathrm{Cone}_{\mathcal{U}}(X, \star, (r_n))$, we have

$$\frac{d(a_n, b_n)}{Ar_n} - \frac{B}{r_n} \leqslant \frac{d(f(a_n), f(b_n))}{r_n} \leqslant \frac{Ad(a_n, b_n)}{r_n} + \frac{B}{r_n}$$

by the fact that f is an (A, B)-quasi-isometry and dividing all sides by r_n. Hence, by taking ultralimits, we see that

$$\begin{aligned} \frac{1}{A} d([(a_n)], [(b_n)]) = \quad \lim_{\mathcal{U}} \frac{d(a_n, b_n)}{Ar_n} &= \lim_{\mathcal{U}} \frac{d(a_n, b_n)}{Ar_n} - \frac{B}{r_n} \\ &\leqslant \lim_{\mathcal{U}} \frac{d(f(a_n), f(b_n))}{r_n} \quad = d(\tilde{f}([a_n]), \tilde{f}([b_n])) \\ &\leqslant Ad([a_n], [b_n]). \end{aligned}$$

Thus, we see that $\tilde{f}$ is bi-Lipschitz; the fact that $\tilde{f}$ is well-defined now follows again from noting that if two sequences have distance 0 in the ultralimit, then so do their images under $\tilde{f}$ by the fact that $\tilde{f}$ is bi-Lipschitz.

As any bi-Lipschitz map is homeomorphic with its image, it now suffices to prove that $\tilde{f}$ is surjective. This follows from noting that, if $[y_n] \in \mathrm{Cone}_{\mathcal{U}}(Y, f(\star), (r_n))$, then there exists a sequence (x_n) such that $d(x_n, y_n) \leqslant B$ for all $n \in \mathbb{N}_+$ because f is an (A, B)-quasi-isometry. Thus

$$d([f(x_n)], [y_n]) = \lim_{\mathcal{U}} \frac{d(f(x_n), y_n)}{r_n} \leqslant \lim_{\mathcal{U}} \frac{B}{r_n} = 0,$$

hence $\tilde{f}([x_n]) = [y_n]$.

□

However, as continuous model theorists, this is less than ideal. One does not merely want a homeomorphism, but rather an $\mathcal{L}$-isomorphism, which is necessarily an isometry.

However, should we be able to demonstrate the missing direction from the EF game proof, then the underlying logic of the spaces becomes clearer. Namely, we would be able to show that two elementarily equivalent groups would give rise to elementarily equivalent asymptotic cones. Details may be found in Appendix B.

It is reasonable to ask what might happen if one were to work instead with G and H viewed as groups with the discrete metric, as this is one technique for viewing “classical” structures as metric structures. However, the next proposition demonstrates that not much is learned by doing this.

Proposition 3.10. *If X is a bounded metric space, then its cone is a metric space with a single point. In particular, the cone of any compact space or space endowed with the discrete metric is a point.*

Proof. This follows directly from noting that, if X is bounded by M, then for any sequences (x_n) and (y_n), we have

$$\lim_{\mathcal{U}} d_n(x_n, y_n) \leqslant \lim_{\mathcal{U}} \frac{M}{r_n} \to 0.$$

□

We also note that one should not expect the converse of this theorem to be true, especially in light of the fact that we are using continuous logic. One can find a few simple counterexamples:

Example 3.11. *Consider* $(\mathbb{Z}, d, 0)$ *and* $(\mathbb{R}, d, 0)$ *as* $\mathcal{L}$*-structures where* d *is the usual metric and* $\mathcal{L}$ *is the empty language. Then the inclusion map* $i : \mathbb{Z} \hookrightarrow \mathbb{R}$ *is clearly a* $(1,1)$*-quasi-isometry, and this induces a bona fide isometry* $\tilde{i} : \mathrm{Cone}(\mathbb{Z}, 0, (r_n)) \to \mathrm{Cone}(\mathbb{R}, 0, (r_n))$*. In particular, they are both* $\mathcal{L}$*-isomorphic to* $(\mathbb{R}, d, 0)$*, hence e.e. Yet* $(\mathbb{Z}, d, 0) \not\equiv (\mathbb{R}, d, 0)$*. Note that*

$$\mathbb{Z} \not\models \inf_x \Big|_1^2 |x - 0.5|,$$

yet

$$\mathbb{R} \models \inf_x \Big|_1^2 |x - 0.5|.$$

Even if we restricted our focus to cones of finitely generated groups, the converse does not hold:

Example 3.12. *Consider* $\mathcal{Z}_1 := (\mathbb{Z}, d, 0)$ *with the usual metric (i.e. the word metric when* $\mathbb{Z}$ *is generated by* $\{1\}$*) and* $\mathcal{Z}_2 := (\mathbb{Z}, d', 0)$ *where* d *is the word metric when* $\mathbb{Z}$ *is generated by* $\{1, 2\}$ *(c.f. Example 3.5). Each of these may be viewed as an* $\mathcal{L}$*-structure where* $\mathcal{L}$ *is the empty language simply because they are each complete pointed metric spaces. Now, let* φ *be the following* $\mathcal{L}$*-sentence:*

$$\varphi := \inf_{x_1, \ldots, x_4} \Big|_{\bar{1}}^{\bar{2}} \left(\sum_{i=1}^{4} \left[d(0, x_i) \dot{-} 1 \right] + \sum_{1 \leqslant i < j \leqslant 4} [d(x_i, x_j) \dot{-} 1] \cdot [d(x_i, x_j) \dot{-} 2] \right).$$

I claim that $\mathcal{Z}_1 \not\models \varphi$ *yet* $\mathcal{Z}_2 \models \varphi$*. Recall by Proposition 2.6 that* $\mathcal{Z}_i \models \varphi$ *if and only if for every* $\epsilon > 0$*, there are* $b_1, \ldots, b_4 \in B_1(\mathcal{Z}_i)$ *such that* $\psi(\bar{b})^{\mathcal{Z}_i} \leqslant \epsilon$ *where* ψ *is the* $\mathcal{L}$*-subformula without the* inf *quantifier.*

Let's parse this sentence. Note that the first sum $\sum_{i=1}^{4} d(0, x_i) \dot{-} 1$ *implies that each of the* x_i *must be* ϵ*-close to distance* 1 *from the basepoint* 0. *The second sum says that the distances between the interpretations of any two distinct* x_i *and* x_j *must be* ϵ*-close to either* 1 *or* 2. *In other words, this sentence is "true" if we are able to find* 4 *elements that are almost a distance* 1 *from* 0 *and yet pairwise are distance either* 1 *or* 2 *from each other. This is meant to formalize an existential statement, which is typically a routine task in classical model theory, but becomes a slight challenge in continuous model theory due to the fact that we no longer have a logical connective for negation.*

Note that in $\mathcal{Z}_2$, *the elements* $\pm 1, \pm 2$ *suffice as interpretations for* $x_1, \ldots, x_4$. *It should be clear that* $d'(\pm 1, 0) = 1$ *and* $d'(\pm 2, 0) = 1$. *Likewise, we see that*

$$d'(1,-1) = 1, \;\; d'(1,2) = 1, \;\; d'(1,-2) = 2, \;\; d'(-1,2) = 2, \;\; d'(-1,-2) = 1, \;\; d'(2,-2) = 2.$$

Thus, $\mathcal{Z}_2 \models \varphi$, *as this interpretation for* $x_1, \ldots, x_4$ *works for all* $\epsilon > 0$.

Now, fix $\epsilon = \frac{1}{2}$. *Then we note that, for each* $b_i \in B_1(\mathcal{Z}_1)$, *we must have* $(d(0, b_i) \dot{-} 1)^{\mathcal{Z}_1} \leqslant \frac{1}{2}$, *but this is only possible if* $b_i = \pm 1$. *Thus, by Pigeonhole, there must exist some* $b_i = b_j$ *with* $i \neq j$, *and in this case,*

$$\Big([d(b_i, b_j) \dot{-} 1] \cdot [d(b_i, b_j) \dot{-} 2]\Big)^{\mathcal{Z}_1} = |0 - 1| \cdot |0 - 2| = 2 \nleqslant \frac{1}{2}.$$

Thus, we see that $\mathcal{Z}_1 \not\models \varphi$.

Finally, it remains to show that $\mathrm{Cone}(\mathcal{Z}_1, 0, (r_n)) \equiv \mathrm{Cone}(\mathcal{Z}_2, 0, (r_n))$. *In fact, these cones are isometric, and in particular* $\mathcal{L}$*-isomorphic (the only additional requirement being that the isometry must map* $0 \mapsto 0$*). This is clear by noting that* $f : \mathcal{Z}_1 \hookrightarrow \mathcal{Z}_2$ *given by* $f(n) = 2n$ *is a* $(1,1)$*-quasi isometry, hence* $\tilde{f}$ *is a bona fide isometry and sends* $[0]_{\mathcal{U}} \mapsto [0]_{\mathcal{U}}$.

In general, there is much that is preserved as one passes from a metric space to an asymptotic cone. For example, hyperbolicity is preserved, although not the particular δ [17]. It is perhaps unsurprising that the asymptotic cone of a δ-hyperbolic space will be 0-hyperbolic; what is more surprising is that there is a partial converse to this fact! Namely, if one can find an ultrafilter such that all asymptotic cones of a metric space are 0-hyperbolic, then they necessarily arise as the cone of a hyperbolic space. This is presented in [9], and we record the statement here for reference.

Proposition 3.11. *If $(X, d, \star)$ is δ-hyperbolic, then every asymptotic cone is 0-hyperbolic, and is an $\mathbb{R}$-tree if (X, d) is geodesic.*

Moreover, if (X, d) is a geodesic metric space and there exists $\mathcal{U}$ such that for all $(\star_n)$ and every appropriate (r_n), $\mathrm{Cone}_{\mathcal{U}}(X, (\star_n), (r_n))$ is an $\mathbb{R}$-tree, then (X, d) is hyperbolic.

The first part of this proposition is easy to see with the use of continuous model theory and Łoś' Theorem. Indeed, we note that if $(X, d, \star)$ satisfies the δ-four point condition, then $(X, \frac{d}{n}, \star)$ satisfies the δ/n-four point condition. It follows that an asymptotic cone must satisfy the δ-four point condition for all $\delta \geqslant 0$, hence must satisfy the 0-four point condition and be 0-hyperbolic.

The proposition does leave open the possibility of a non-hyperbolic space with at least one, but not necessarily all, of its asymptotic cones being $\mathbb{R}$-trees. These are called *lacunary hyperbolic groups*, and have been studied in [32].

Chapter 4

Buildings

The asymptotic cones we study, while they are metric spaces in their own right, may be equipped with additional natural structure that makes them a certain kind of *building*. Intuitively, similarly to how one may think of simplicial complexes as higher-dimensional analogues of graphs, buildings may be thought of as higher-dimensional analogues of trees. Indeed, the dimension 1 buildings we examine turn out to be $\mathbb{R}$-trees that we saw earlier, once we view them as pure metric spaces. We shall make precise later how one may recover tree-like structure in the higher dimensional cases as well.

Firstly, there are several different kinds of buildings that exist within the literature. Much of building theory was first being developed in the late 1960s to 1970s in an attempt to give a "systematic procedure for the geometric interpretation of the semisimple Lie groups." In [38], Tits worked with the "spherical buildings." In our context, we will be mostly concerned with what are sometimes referred to as *nondiscrete affine* $\mathbb{R}$*-buildings*. These were first studied by Tits in [39] under the name "système d'appartements" and are a generalization of the so-called *Euclidean buildings*, which have been used in the study of algebraic groups over p-adic fields. We shall mention towards the end of the dissertation how one may recover a certain "spherical building at infinity" when working in a nondiscrete affine $\mathbb{R}$-building, though this notion will not be necessary to state our main result.

Each particular type of building comes with its own nuances, but they all broadly share the same terminology. For our purposes, we will only focus on the *nondiscrete affine* $\mathbb{R}$*-buildings*, pulling largely from [5] as reference material. In [5], they work with a generalization of the *nondiscrete affine* $\mathbb{R}$*-buildings* by considering *nondiscrete affine* Λ*-buildings* for an arbitrary ordered value group Λ. They note that there is a distinction between the metric induced by the root data from $\Lambda = \mathbb{R}$ and the Euclidean metric; however, this distinction *only* affects the metric one obtains in the so-called *model space*, and what we present from Bennett et al. shall remain true by utilizing the Euclidean metric as normal.

We comment here that much of our presentation of building theory is biased towards the nondiscrete case. As Schwer (née Hitzelberger) has laid out in her thesis [23], there is the following inclusion of building "types":

$$\{\text{Euclidean}\} \subseteq \{\text{nondiscrete affine-}\mathbb{R}\} \subseteq \{\text{nondiscrete affine-}\Lambda\}$$

Further, much of the terminology and notation is as of yet not completely standardized, as the theory continues to develop. Schwer has included in her thesis a "dictionary" to help a new learner navigate the oft-confusing landscape of varying terminology and notation.

4.1 Nondiscrete Affine $\mathbb{R}$-Buildings

A nondiscrete affine $\mathbb{R}$-building requires a few pieces of data. These involve a point space X, a model space $\mathbb{A}$ (which includes a group W_T that acts on the space underlying $\mathbb{A}$ by isometries), and a collection $\mathcal{F}$ (called an atlas) of isometric embeddings (called charts) $f : \mathbb{A} \hookrightarrow X$ subject to certain constraints. Intuitively, the idea is that a nondiscrete affine

$\mathbb{R}$-building is a space in which we embed flat space of some maximal dimension, and the arrangement of these maximal flat spaces must adhere to certain guidelines. In some sense, it is not dissimilar to the study of manifolds from the point of view of charts and atlases.

Given the data of $(X, \mathbb{A}, \mathcal{F})$, the continuous model theorist may be tempted to think of a building as an unbounded metric structure with a separate sort for the space underlying $\mathbb{A}$, the group W_T, and a collection of function symbols $\mathcal{F}$. While it is not a priori impossible to view this as a metric structure, our building "axioms" involve existential statements which quantify over the charts in our atlas, and it's not obvious that this is easily axiomatizable in our framework for continuous logic, as one does not allow for quantification over symbols from our language.

The Model Space $\mathbb{A}$

The model space $\mathbb{A}$ of a nondiscrete affine $\mathbb{R}$-building should be thought of intuitively as an n-dimensional vector space over $\mathbb{R}$ together with a group generated by reflections across hyperplanes. Here, we introduce the formal definition for the model space and establish a glossary of definitions that we'll be using in our discussion of buildings. Much of the terminology is completely standard, and we comment on the slight deviations that we make.

Definition 4.1. *We say that a set Φ of nonzero vectors in $\mathbb{R}^n$ is a* **root system** *if*

- *Φ is finite and spans $\mathbb{R}^n$,*
- *for all $\alpha \in \Phi$, $\Phi \cap \mathbb{R}\alpha = \{\alpha, -\alpha\}$,*
- *for all $\alpha \in \Phi$, $s_\alpha \Phi = \Phi$,*

where $\mathbb{R}\alpha := \{\lambda\alpha : \lambda \in \mathbb{R}\}$ and s_α is the action on $\mathbb{R}^n$ determined by reflection across the unique hyperplane orthogonal to α, which we denote by $H_{\alpha,0}$.

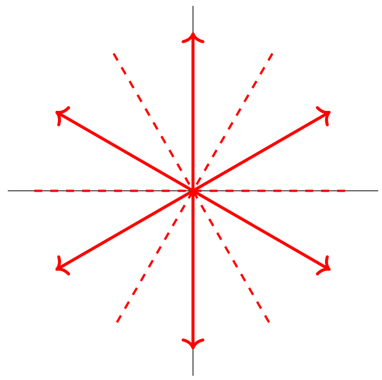

Figure 4.1: A root system in $\mathbb{R}^2$ and the hyperplanes orthogonal to the roots

One should think of a root system as the normal vectors which determine hyperplanes one will reflect over. In this way, we may refer to the the group of these reflections as a *reflection group*. It is known that finite reflection groups have been classified up to isomorphism; for a list of the isomorphism types and how (most) of these arise as finite reflection groups from a particular root system, see [1]. We will not be terribly interested in the specifics of these classifications, but it will be useful for us to see how root systems give rise to reflection groups which affect the geometry of our model space. Specifically, we shall later discuss algebraic groups, and it will be important that the root system is, in some sense, inherent to the algebraic group defined over some base field.

It is often helpful to have at least one picture of a root system and the corresponding reflection hyperplanes. We shall refer back to this picture later as we introduce more terminology.

In Figure 4.1, the vectors are the roots, and the dashed lines are the corresponding reflection hyperplanes. The reflection group that one obtains in isomorphic to D_3. It is not a requirement in general that all roots have the same magnitude, but it is not hard to convince oneself that this must be the case in Figure 4.1.

Definition 4.2. *Given a root system* Φ*, we define the* **spherical Weyl group** $\bar{W}$ *to be the group* $\langle s_\alpha : \alpha \in \Phi \rangle$*. Further, if we are given some* $\bar{W}$*-invariant translation group*

$T \subseteq \mathrm{Isom}(\mathbb{R}^n)$,[1] *then we define* $W_T := \bar{W} \ltimes T$*. In the case that* T *is the full translation group* $T = \mathbb{R}^n$,[2] *we shall simply write* W *instead of* W_T*, and refer to this as "the"* **full affine Weyl group**.

It is important to note our use of quotations around "the" in the previous definition, as "the" full affine Weyl group clearly depends on its spherical part. One should think of W_T as being the group where one is allowed not only to reflect along hyperplanes through the origin, but also along translates of these hyperplanes.

Definition 4.3. *We say that the* **model space** $\mathbb{A}$ *for an affine* $\mathbb{R}$*-building is the pair* $\mathbb{A} = (\mathbb{R}^n, W_T)$*.*

In most instances, we shall concern ourselves with T some nondiscrete subgroup of $\mathrm{Isom}(\mathbb{R}^n)$, and in such an instance, we shall use the terminology *nondiscrete* affine $\mathbb{R}$-building. However, it is worth noting that, in the case T is trivial, this reduces the model space $\mathbb{A}$ to a finite reflection group in the sense of [1], and we shall borrow a few of their results in order to describe the geometry of the model space.

First, we introduce several terms which are commonly used within building theory. However, in the literature, it is common to draw no distinction between terms such as *Weyl chamber* within the model space and the image of one within the building itself. We choose to distinguish between these concepts, and so we opt to append the prefix "pre-" to some of our terminology.

In [4], they use the root data to give a precise definition of a hyperplane in their more general model space by using the notion of co-roots. As we mentioned, we shall not need a robust notion of root systems in order to present our findings, and so we present a simplified (but precise) outline of our terminology below.

[1] Here, $\bar{W}$-invariant simply means that for any $w \in \bar{W}$, we have that $wTw^{-1} = T$.

[2] By abuse of notation; we are identifying $\mathbb{R}^n$ with the set of vectors we translate by.

Definition 4.4. *Given an affine Weyl group W_T, we say that a* **pre-wall** *is a subset $H_{\alpha,t} := t + H_{\alpha,0}$ where $t \in T$ and $H_{\alpha,0}$ is the hyperplane through the origin with normal vector given by the root $\alpha \in \Phi$ as before. If W is the full affine Weyl group, we shall by convention assume that t is a translation in a direction orthogonal to $H_{\alpha,0}$.*

The above convention will be useful when playing EF games between so-called "wall trees" in a building. We shall mention it again for clarity, but it is good to record now.

Definition 4.5. *We say that a* **pre-half apartment** *$H^+_{\alpha,t}$ is the subset of $\mathbb{R}^n$ determined by the connected component of $\mathbb{R}^n \setminus H_{\alpha,t}$ containing α, together with its boundary wall $H_{\alpha,t}$.*

Remark 4.1. *We note that the pre-half apartment "opposite" $H^+_{\alpha,t}$ is simply given by $H^+_{-\alpha,t}$, as $\alpha \in \Phi$ implies $-\alpha \in \Phi$.*

Example 4.1. *In Figure 4.1, one may find* 6 *pre-half apartments determined by $y \geqslant 0$, $y \geqslant \sqrt{3}x$, $y \geqslant -\sqrt{3}x$, and the same expressions where one changes $\geqslant$ to $\leqslant$. If T is nontrivial, then one may also find pre-half apartments which are, in effect, translations of these* 6*.*

Before we dive into the rest of our definitions of common objects found within a model space (and later, a building), we define the notion of a *simplicial cone*, because in all of the instances we shall care about, our pre-walls shall "cut" the model space into isomorphic simplicial cones.

Definition 4.6. *Suppose that V is a finite dimensional vector space. Then $C \subseteq V$ is a* ***simplicial cone*** *if there is some basis $b_1, \ldots, b_n$ of V such that C consists of all linear combinations $\sum \lambda_i b_i$ where all $\lambda_i > 0$.*

In other words, C is the interior of the cone over the simplex with vertices given by the b_i.

If one assumes that our root system is *essential* (which will always be the case for us; cf. §1.4 of [1]), then the hyperplanes $H_{\alpha,0}$ through the origin "cut" our space $\mathbb{R}^n$ into disconnected, isometric regions, with isometries induced by actions by $\bar{W}$. These regions turn out to be simplicial cones, and we may choose any of these simplicial cones, together with its boundary, to be a fundamental domain for the action of W_T.

Definition 4.7. *Given an essential root system, we say that the closure of a choice of simplicial cone "cut" by the hyperplanes fixed by $\bar{W}$ is the* **fundamental pre-Weyl chamber**, *and we denote it as C_f.*

In Figure 4.1e, we see that the hyperplanes "cut" $\mathbb{R}^2$ into 6 isometric pieces. If we consider the hyperplanes $y = 0$ and $y = \sqrt{3}x$ in the picture, we may choose the region defined by $y \geqslant 0$ and $y \leqslant \sqrt{3}x$ to be our fundamental pre-Weyl chamber. This is the simplicial cone (together with its boundary) which is completely contained within the first quadrant.

It may seem odd to choose a "fundamental" domain so arbitrarily, but the idea is simply to pick some domain with which to define the notion of a (not necessarily fundamental) pre-Weyl chamber. Figure 4.1 makes it perhaps too seemingly obvious that any other "chamber" is the image of the fundamental pre-Weyl chamber under an action by some element of $\bar{W}$, and fortunately, this is indeed the case in general.

Our use of the word "simplicial" suggests that one may find faces in a simplicial cone. One may intuitively guess that a 0-dimensional face is a point and a 1-dimensional face is a ray. This is indeed the case, but is imprecise for the same reason that a 1-dimensional face of a simplex is not *just* a line segment, but is the convex hull of two points determining the simplex. To give a more precise view of simplicial faces, we follow §1.4 of [1].

First, let $\mathcal{H} := \{H_1, \ldots, H_k\}$ be the finite set of hyperplanes arising from the root system and choose a linear $f_i : \mathbb{R}^n \to \mathbb{R}$ such that H_i is defined by $f_i = 0$. Then, one may define *cells* in $\mathbb{R}^n$ to be nonempty subsets determined by, for each i, an expression of the form $f_i = 0$, $f_i > 0$, or $f_i < 0$. The authors write this as a sign sequence according to $f_i = \sigma_i$ where $\sigma_i \in \{0, +-\}$, taking the convention that "$f_i = +$" denotes $f_i > 0$, and similarly for

"$f_i = -$".[3]

Then, if one has cells $A, B \subseteq \mathbb{R}^n$, then by abuse of notation, we let $\sigma_i(A)$ denote the sign corresponding to f_i, and we say that B is a *face* of A (denoted $B \leqslant A$) if for each i, either $\sigma_i(B) = 0$ or $\sigma_i(B) = \sigma_i(A)$. Intuitively, we are saying that one may use the hyperplanes cutting up our space $\mathbb{R}^n$ to determine our open chambers, and one may find a face of an open chamber by taking its intersection with one or more of the hyperplanes bounding it.

Definition 4.8.

- *The hyperplanes $H_{\alpha,t}$ fixed by some reflection from the affine Weyl group W_T are called* **pre-walls**.
- *A* **pre-Weyl chamber** *is the image of C_f under an action by W_T.*
- *The face of a pre-Weyl chamber is called a* **pre-Weyl simplex**; *the pre-Weyl simplices of co-dimension 1 are called* **pre-panels**. *The (unique) face of dimension 0 is the* **base** *of the pre-Weyl chamber/simplex. If x is the base of a pre-Weyl chamber/simplex, then we shall say that the chamber/simplex is* ***based at*** *x.*
- *We say that a set $S \subseteq \mathbb{A}$ is W_T-***convex** *if S is the intersection of finitely many pre-half apartments.*

Let us return to the case where T is trivial. Once one fixes a fundamental pre-Weyl chamber C_f, we may define S to be the set of reflections which fix a pre-panel of C_f. It turns out that $\bar{W}$ is generated by S; even further, there is a one-to-one correspondence between $\bar{W}$ and the pre-Weyl chambers which are the images of C under the action from $\bar{W}$, and the correspondence is given by $w \longleftrightarrow wC_f$.

Now, let $-C_f$ denote the pre-Weyl chamber "opposite" C_f, in the sense that, if C_f is the

[3] N.B. It is worth pointing out that the "nonempty" requirement is necessary. Indeed, it should not be hard to convince oneself that any sign sequence without 0 must either be empty or have nonempty interior. In Figure 4.1, this means that any sign sequence without 0 is either one of the open chambers or is empty, and clearly there are 8 such sign sequences and only 6 open chambers, hence some sign sequence will correspond to an empty subset of $\mathbb{R}^n$.

closure of the simplicial cone given by basis $\mathcal{B} = \{b_1, \ldots, b_n\}$, then $-C_f$ is the closure of the simplicial cone given by basis $-\mathcal{B} = \{-b_1, \ldots, -b_n\}$. Then it turns out that $\bar{W}$ has a unique longest element w_0 (with respect to the word metric induced by S), and w_0 may be characterized by the fact that $w_0 C_f = -C_f$.

Nondiscrete Affine ℝ-Buildings

Given a model space $\mathbb{A}$, we are now prepared to define nondiscrete affine ℝ-buildings. In the following defintion, we shall abuse our notation slightly and write $f : \mathbb{A} \hookrightarrow X$ to denote that $f : \mathbb{R}^n \to X$ is an injection; the use of the model space in our notation serves as a reminder that there is additional structure to be preserved.

Definition 4.9. *Given a model space $\mathbb{A}$, we say that $(X, \mathbb{A}, \mathcal{F})$ is a* **nondiscrete affine ℝ-building** *where X is a set, and $\mathcal{F}$ is a collection of injections $f : \mathbb{A} \hookrightarrow X$ (whose images are called apartments, and the images of pre-*"whatever" *are simply* "whatever"*) subject to the following requirements:*[4]

(A1) Given any $f \in \mathcal{F}$ and $w \in W_T$, $f \circ w \in \mathcal{F}$.

(A2) Given $f, g \in \mathcal{F}$ such that $f(\mathbb{A}) \cap g(\mathbb{A}) \neq \varnothing$, $f^{-1}(g(\mathbb{A}))$ is a closed, convex subset of $\mathbb{A}$, and there exists some $w \in W_T$ such that $f \restriction_{f^{-1}(g(\mathbb{A}))} = (g \circ w) \restriction_{f^{-1}(g(\mathbb{A}))}$.

(A3) For any $x, y \in X$, there is some $f \in \mathcal{F}$ such that $x, y \in f(\mathbb{A})$.

We note that we can define a function $d_X : X^2 \to \mathbb{R}_{\geqslant 0}$ between pairs of points $x, y \in X$ by finding some $f \in \mathcal{F}$ such that $x, y \in f(\mathbb{A})$ and setting $d_X(x, y) := d_{\mathbb{A}}(f^{-1}(x), f^{-1}(y))$. The fact that such an f exists follows from (A3), and the fact that this function is well-defined

[4] We choose to refer to these as "requirements" rather than the more psychologically natural "axioms" so as not to reinforce the false assumption that these are axiomatizable in continuous model theory.

follows from (A2). Requirement (A5) will imply that d_X also satisfies the triangle inequality, hence is a metric on X.

(A4) Given two Weyl chambers C_1 and C_2 in X, there are sub-Weyl chambers $C_i' \subseteq C_i$ (for $i = 1, 2$) that are contained in a common apartment.

(A5) For any apartment A in X and any $x \in A$, there is a retraction map $r_{A,x} : X \to A$ such that $r_{A,x}$ does not increase distances and $r_{A,x}^{-1}(x) = \{x\}$.

(A6) If f, g, h are charts whose apartments intersect pairwise in half-apartments, then $f(\mathbb{A}) \cap g(\mathbb{A}) \cap h(\mathbb{A}) \neq \varnothing$.

If $\mathbb{R}^n$ is the underlying space for $\mathbb{A}$, then we shall say that the **dimension of** X*, denoted* $\dim(X)$*, is n.*

Our comment about d_X is telling us that, intuitively, one may define a metric on a building by finding any apartment containing two points and measuring the distance between these points from the point of view of the model space $\mathbb{A}$. One should note that the typical definition of a nondiscrete affine $\mathbb{R}$-building does not formally require X to be a metric space, but as we shall see later, such a building equipped with this metric will be isometric to a certain asymptotic cone, so for our purposes, we should always think of such a building as a metric space with *additional* building structure. In light of this, one should think of the charts given by the atlas not just as injections, but as isometric embeddings of the model space $\mathbb{A}$ into our building.

Aside from this, there are several non-intuitive facts one may deduce from the above requirements, of which we mention only a few. In order to state some of these additional requirements, we introduce more terminology regarding *germs* of of Weyl simplices. This will be an important notion for us as we prove our main theorem.

Definition 4.10. *Let X be a nondiscrete affine $\mathbb{R}$-building. We say that two Weyl simplices S and T* **share the same germ** *if S and T are both based at x and if $S \cap T$ is a neighborhood of x in S and in T. We let $\Delta_x S$ denote the equivalence class of an x-based Weyl simplex S, and this is called the* **germ of** S **at** x*.*

If S is a Weyl chamber, then we say that $\Delta_x S$ is **contained** *in a set Y if there is some $\epsilon > 0$ such that $S \cap B_\epsilon(x) \subseteq Y$.*

Definition 4.11. *We say that two x-based germs are* **opposite** *if they are contained in a common apartment and are images of each other under the longest element of the spherical Weyl group $\bar{W}$. Two Weyl chambers are* **opposite at** x *if their germs are opposite.*

Note that two chambers being "opposite" each other is described using the spherical part of the affine Weyl group. But what if we were to consider the translates of Weyl chambers? This brings us to our notion of *parallel* Weyl chambers, and we will extend our definition to define the notion of parallel Weyl *simplices* as well.

Definition 4.12. *We use the word* **parallel** *to describe any of the following scenarios. In each case, being parallel is an equivalence relation, and we denote it by $\sim$.*

- *Two pre-walls are* **parallel** *if they are translates of one another. I.e., there are $\alpha \in \Phi$ and $t, t' \in \mathbb{R}$ such that $H \sim H' \implies H = H_{\alpha,t}$ and $H' = H_{\alpha,t'}$. That is, they may be determined by the same root.*
- *Two walls are* **parallel** *if the Hausdorff distance between them is finite.*
- *Two pre-Weyl chambers S and T are* **parallel** *if $S \cap T$ contains a pre-Weyl chamber.*
- *Two Weyl chambers S and T are* **parallel** *if $S \cap T$ contains a Weyl chamber.*
- *Two Weyl simplices S and T are* **parallel** *if there is a sequence of Weyl simplices $S = S_0, S_1, S_2, \ldots, S_n = T$ such that for each $k = 0, 1, \ldots, n-1$, the Weyl simplices*

S_k and S_{k+1} are translates of each other in a common apartment. That is, if $f \in \mathcal{F}$ is a chart whose image contains both S_k and S_{k+1}, then the pre-images of S_k and S_{k+1} are translates of each other.

Bennett proved in [4] that two hyperplanes in an affine Λ-building are parallel if and only if they are bounded distance from each other. Indeed, a similar statement can be made about Weyl chambers and Weyl simplices in a nondiscrete affine $\mathbb{R}$-building. Namely:

Proposition 4.1. *Two Weyl simplices contained in a common apartment are parallel if and only if they are bounded distance from each other.*

This gives us good justification to define parallelism for walls as we did above.

It is worth noting that nothing about our building requirements dictates that any subset which is isometric to $\mathbb{R}^n$ is realized as the image of some chart. Stated another way, it is not a requirement that a building be equipped with all charts that are consistent with the building requirements. A cute example to this effect is laid out in Remark 5.29 of [23]. This leads us to the following definition and observation:

Definition 4.13. *We say that a nondiscrete affine $\mathbb{R}$-building is* **atlas complete** *if every chart compatible with $(A1) - (A6)$ is contained in the atlas $\mathcal{F}$.*

It is known that every such building has a unique extension to a complete atlas. [39] Having a complete atlas will be a necessary assumption for a few of our minor results, so it is worthwhile to note that, because our buildings shall arise from ultraproducts (hence are $\aleph_1$-saturated), they shall be atlas complete [27].

We are now ready to state a few facts which follow from the building axioms $(A1) - (A6)$.

Proposition 4.2. *Suppose that X is a nondiscrete affine $\mathbb{R}$-building. Then the following are true:*

- *(EC)* *(Exchange Condition) Given two charts* $f, g \in \mathcal{A}$ *such that* $f(\mathbb{A}) \cap g(\mathbb{A})$ *is a half-apartment, there exists a third chart* $h \in \mathcal{A}$ *such that* $h(\mathbb{A}) \cap f(\mathbb{A})$ *and* $h(\mathbb{A}) \cap g(\mathbb{A})$ *are both half-apartments. In particular,* $h(\mathbb{A})$ *is the symmetric difference of* $f(\mathbb{A})$ *and* $g(\mathbb{A})$ *union the boundary wall of* $f(\mathbb{A}) \cap g(\mathbb{A})$*.*

- *(LA)* *(Large Atlas) Any two germs of Weyl chambers are contained in a common apartment.*

- *(CO)* *(Opposite Chambers) Any two Weyl chambers* S *and* T *which are opposite at* x *are contained in a unique common apartment.*

- *(TI)* *(Triangle Inequality) The function* d_X *defined using (A1) - (A3) satisfies the triangle inequality, hence is indeed a metric.*

- *(SC)* *(Sundial Configuration) Suppose* $f_1 \in \mathcal{F}$ *and* C *is a Weyl chamber in* $(X, \mathbb{A}, \mathcal{F})$ *such that* $C \cap f_1(\mathbb{A}) = P$ *is a panel, and* $M \subseteq f_1(\mathbb{A})$ *is the wall containing* P*. Then there exist* $f_2, f_3 \in \mathcal{F}$ *such that* $f_1(\mathbb{A}) \cap f_i(\mathbb{A})$ *is a half-apartment and* $(M \cup S) \subseteq f_i(\mathbb{A})$ *for* $i = 2, 3$.

Alternatively, one may instead use these deductions to form separate, equivalent lists of requirements, as is laid out in [5]. We state a few of these equivalences here:

Proposition 4.3. *Suppose that* $(X, \mathbb{A}, \mathcal{F})$ *satisfies conditions* $(A1)-(A3)$*. Then the following are equivalent:*

- $(X, \mathbb{A}, \mathcal{F})$ *satisfies (A4), (A5), and (A6).*
- $(X, \mathbb{A}, \mathcal{F})$ *satisfies (A4), (A5), and (EC).*
- $(X, \mathbb{A}, \mathcal{F})$ *satisfies (A4), (A5), and (SC).*
- $(X, \mathbb{A}, \mathcal{F})$ *satisfies (A4), (TI), and (A6).*
- $(X, \mathbb{A}, \mathcal{F})$ *satisfies (A4), (BI), and (A6).*

- *$(X, \mathbb{A}, \mathcal{F})$ satisfies (A4), (TI), and (EC).*

- *$(X, \mathbb{A}, \mathcal{F})$ satisfies (LA) and (CO).*

It may be difficult to gain an intuition for these axioms without a few examples and nonexamples. We shall present a few examples of nondiscrete affine $\mathbb{R}$-buildings of various dimensions. For the dimension 1 and 2 cases, we offer some sketches to help with building[5] intuition.

Example 4.2. *One could view $\mathbb{R}^n$ as a building itself. That is, $X = \mathbb{R}^n$, $\mathbb{A} = (\mathbb{R}^n, W_T)$, and $\mathcal{F}$ may be described by*

$$\mathcal{F} := \{f \circ w : w \in W_T\}$$

where f is any chart $f \in \mathcal{F}$. Indeed, as there is only one apartment (namely, X itself), checking axioms (A3), (LA), and (CO) is completely trivial. We note that (A1) follows from the fact that if $g = f \circ w \in \mathcal{F}$, then $g \circ w' = f \circ (w \circ w') \in \mathcal{F}$. And for (A2), clearly if $g_1, g_2 \in \mathcal{F}$ with $g_i = f \circ w_i$, then $g_1^{-1}(g_2(\mathbb{R}^n)) = \mathbb{R}^n$, and $g_1 = f \circ w_1 = f \circ w_2 \circ w_2^{-1} \circ w_1 = g_2 \circ (w_2^{-1} w_1)$. We don't take much interest in this building.

Example 4.3. *If $(X, \mathbb{A}, \mathcal{F})$ is a dimension 1 nondiscrete affine $\mathbb{R}$-building, then its underlying metric space is an $\mathbb{R}$-tree.*

It is not clear prima facie that dimension 1 affine $\mathbb{R}$-buildings are $\mathbb{R}$-trees; while any affine $\mathbb{R}$-building is a unique geodesic space, it is not obvious that there is only one *arc* between points in the dimension 1 case. In Proposition 3.3 of [4], Bennett proves the more general case that a dimension 1 affine Λ-building is a Λ-tree.

Example 4.4. *A* **book***, also known as a* **fan of half planes***, is a topological space consisting of a line (called the* spine*) and a collection of one or more half-planes (called the* pages*) each having the spine as its boundary. Any book with at least two pages may be viewed as a nondiscrete affine $\mathbb{R}$-building by letting $\mathbb{A} = (\mathbb{R}^2, W_T)$ where $\bar{W} \cong (\mathbb{Z}/2\mathbb{Z}) \times (\mathbb{Z}/2\mathbb{Z})$ is*

[5] Pun intended.

generated by reflections across the two axes.[6] *We then take as our apartments the union of any two pages and the spine. Similar to Example 4.2, upon picking a representative chart for each apartment, we may close the atlas under pre-composition with elements from* $\bar{W}$.

One may systematically check the requirements $(A1)-(A6)$*; it is perhaps easier to check* $(A1)-(A3)$, (LA), *and* (CO). *The most interesting parts are* (LA) *and* (CO) *as* $(A1)$ *and* $(A3)$ *are true basically by construction, and* $(A2)$ *follows from noting that any two distinct apartments intersect in either a half apartment or the spine of the book. The Large Atlas condition can be checked as follows:*

Given a germ ΔC *of a Weyl chamber* C*, if* C *is based at* c_0*, then if* c_0 *is contained in a page* P *(and not the spine), then* P *contains* ΔC*, and it should be clear that* P *is necessarily unique. If* c_0 *is contained in the spine of the book, then* C *is fully contained in (the closure of) a unique page* P*, hence* P *contains* ΔC*. Thus, if one has two germs* ΔC *and* $\Delta C'$*, then one sees that the apartment composed of the two unique pages containing* ΔC *and* $\Delta C'$ *will contain* ΔC *and* $\Delta C'$*, demonstrating (LA).*

Similarly, if C *and* $-C$ *are opposite Weyl chambers based at the same point* c_0*, then if* c_0 *is contained in the spine, it should be clear that* C *and* $-C$ *are contained in (the closure of) distinct pages, and if* c_0 *is contained in some page* P *(and not the spine), then at least one of* C *and* $-C$ *is fully contained in* P*, and the other must contain a point from the spine and "eventually" be contained in another page. It then follows that the union of these two pages and the spine contains* C *and* $-C$ *and is necessarily the unique apartment containing them.*

We use the following visual to highlight the Exchange Condition (EC), which follows from the above verification of the building requirements:

[6] One need not take T to be trivial despite there being only one "branch."

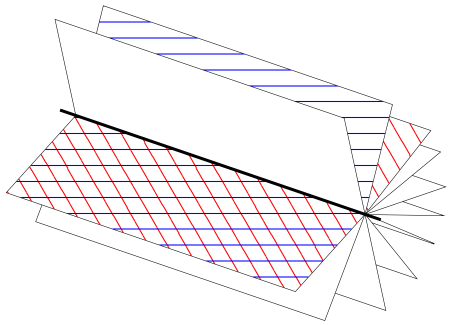

Figure 4.2: Two apartments which intersect in a half-apartment

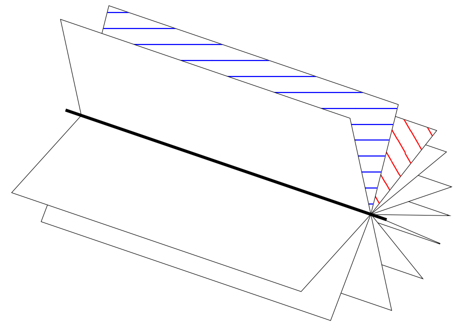

Figure 4.3: The (EC) applied to the two apartments gives a new apartment composed of the "leftover" blue and red half-apartments.

Figure 4.2 highlights two distinct apartments in the book; namely, these are the apartment R colored red and the apartment B colored blue.[7] *Figure 4.3 demonstrates how one may apply the (EC) to find a new apartment: because $R \cap B$ is a half-apartment, we may form a new apartment by taking $(R \oplus B) \cup \partial(R \cap B)$. In the book terminology, the new apartment is the symmetric difference of R and B together with the spine of the book.*

Nonexample 4.1. *As mentioned before, any dimension 1 building is an $\mathbb{R}$-tree. However, it is helpful to see exactly where the axioms fail if one is presented with a dimension 1 space containing a loop.*

Consider the lines $y = 0$, $y = 1 - x$, and $y = 1 + x$ in $\mathbb{R}^2$. One may view each of these lines as isometric embeddings $\sigma_i : \mathbb{R} \hookrightarrow \mathbb{R}^2$ for $i = 1, 2, 3$, respectively. I claim that this cannot *be given a nondiscrete affine $\mathbb{R}$-building structure. More precisely, if one sets $\mathcal{F} := \{\sigma_1, \sigma_2, \sigma_3\}$, $X := \bigcup_{i=1,2,3} \sigma_i(\mathbb{R})$, and $\mathbb{A} := (\mathbb{R}, W_T)$, then there is no $\mathcal{F}' \supseteq \mathcal{F}$ such that $(X, \mathbb{A}, \mathcal{F}')$ is a nondiscrete affine $\mathbb{R}$-building.*

To see this, it shall be easier to work directly with the apartments, rather than the charts. Let $A_i := \sigma_i(\mathbb{R})$ for $i = 1, 2, 3$, and note that each A_i is "cut" into three pieces by the intersections

[7] For those unable to see the color, the lines are oriented in opposing directions.

with the other two apartments. We shall denote these by A_i^j for $j = 1, 2, 3$, as in the following picture:

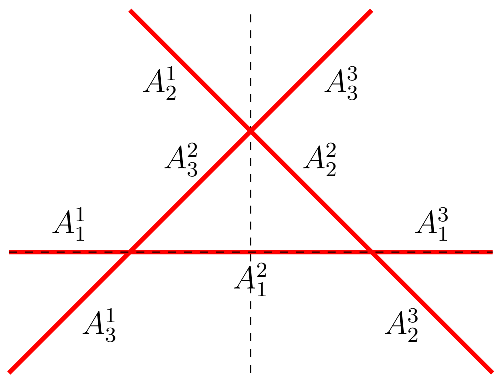

Figure 4.4: A nonbuilding

We note that, by (A3), there must be an apartment B containing $(-2, 0)$ and $(-2, -1)$. As this must be homeomorphic (indeed, isometric) to $\mathbb{R}$, it should be clear that B consists of the points in $A_1^1 \cup A_3^1$. By the Exchange Condition applied to apartments A_3 and B, there is an apartment C consisting of the points in $A_1^1 \cup A_3^2 \cup A_3^3$. Lastly, by applying the Exchange Condition to A_1 and C, there is an apartment D consisting of the points in $A_1^2 \cup A_1^3 \cup A_3^2 \cup A_3^3$, as is highlighted in Figure 4.5.

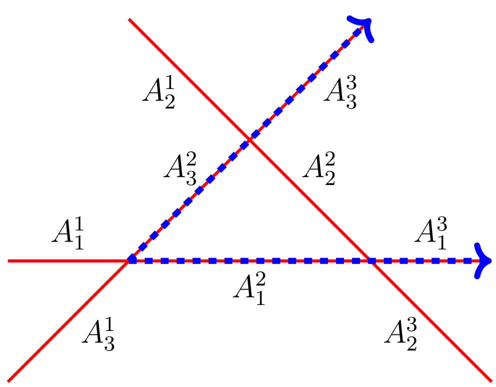

Figure 4.5: The apartment D is the blue/dashed union of rays.

Finally, note that, if D is given by some chart f, then $f^{-1}(\sigma_2(\mathbb{R}))$ is not a closed, convex *subset of $\mathbb{R}$. Indeed, we see that $D \cap A_2$ consists of the points $(0,1)$ and $(1,0)$, which is clearly not convex.*

The above examples and nonexample lend some credence to the statement that "nondiscrete affine $\mathbb{R}$-buildings are higher dimensional analogues of $\mathbb{R}$-trees." In light of this perspective and Nonexample 4.1, one may reasonably expect that, in a building of dimension 2, one may not have three apartments intersecting in hyperplanes forming a hollow triangular prism. This may be shown to be true by utilizing the notion of *wall trees.* First, we outline some terminology that is helpful when discussing dimension 1 buildings.

Definition 4.14. *Let (X,d) be an $\mathbb{R}$-tree and $\mathcal{A}$ a set of lines in X. We say that two rays r_1, r_2 are* **equivalent** *if they eventually agree, in the sense that there is some $y \in r_1 \cap r_2$ such that the subrays r_{i_y} based at y are equal. We then define an* **end** *of the tree to be an equivalence class $[r]$ for a ray r contained in some line $l \in \mathcal{A}$.*

Definition 4.15. *We say that a pair $(X, \mathcal{A})$ is a* **tree with sap**[8] *if X is an $\mathbb{R}$-tree and $\mathcal{A}$ is a set of lines in X such that*

(T1) for any $x, y \in X$, there is some $l \in \mathcal{A}$ such that $x, y \in l$, and

(T2) for any two ends e_1, e_2 in $(X, \mathcal{A})$, there is a line $l \in \mathcal{A}$ such that e_1, e_2 are the ends of the line l.

In this way, affine $\mathbb{R}$-buildings of dimension 1 are trees with sap, and it is not hard to see that the converse is true if one equips a tree with sap with a complete atlas. Bennett proved the following in the case of affine Λ-buildings, but this was known to Tits in the case of affine $\mathbb{R}$-buildings. We mention Bennett for his inclusion of the "tree with sap" terminology.

[8]This terminology was introduced by Ronan in [34], and "sap" is meant as an abbreviation for "**s**ystem of **ap**artments."

Proposition 4.4 (Tits [39], Prop 4)**.** *Suppose that X is a nondiscrete affine $\mathbb{R}$-building of dimension at least 2, and $M \subseteq X$ any wall. Then the parallelism class $\partial_X M$ forms a tree with sap. Moreover, the ends of the tree are in one-to-one correspondence with the half-apartments of ∂X having $\partial_X M$ as their boundary.*

In the above proposition, $\partial X := \{\partial_X S : S \text{ a Weyl simplex in } X\}$. Tits also explains how ∂X is a *spherical* building, and much of the vocabulary between spherical and affine buildings is recycled.

Definition 4.16. *Let X be a nondiscrete affine $\mathbb{R}$-building, and $M \subseteq X$ a wall. Then the $\mathbb{R}$-tree that one obtains from Proposition 4.4 is called a **wall tree with respect to** M.*

Remark 4.2. *Originally, the author found a reference to Proposition 4.4 in [23], where it was stated that the apartments of a wall tree are in one-to-one correspondence with apartments in X. However, this cannot be the case. Consider Example 4.4, and note that, if one takes M to be a wall perpendicular to the spine, then M is contained in a unique apartment, say, composed of pages P_1 and P_2. But if $M' \subseteq X$ is a wall that has nonempty intersection with (the interior of) some page P_3, then it should be clear that P_3 contains a panel of M', and one can check the apartments comprising $P_i \cup P_3$ (for $i = 1, 2$) to see that M' does* not *have finite Hausdorff distance to M, hence these are not parallel. It follows that* all *parallel walls of M must be translates of M in the apartment $P_1 \cup P_2$, and the wall tree with respect to M will be isometric to $\mathbb{R}$.*

The above remark is an unfortunate one, as a correspondence between apartments in a wall tree and apartments in the "ground" building would simplify much analysis of nondiscrete affine $\mathbb{R}$-buildings. Indeed, in an earlier draft, the author thought this correspondence to be correct, as even the beginning of Bennett's proof suggests that, given a chart $f \in \mathcal{F}$ and wall $M \subseteq X$, one may find a pre-wall m such that $f(m) \in \partial_X M$.

The book example is pathological from a certain point of view; namely, it is not a homogeneous space. Further, our choice of wall in the book example was perhaps unnatural. Indeed, we have the following:

Example 4.5. *If M is the spine of the book as in Example 4.4, then the wall tree $\partial_X M$ is a union of rays sharing a common base point; each ray corresponds to a page, and the base point corresponds to the spine of the book. It is then not hard to see that the apartments of the wall tree are in one-to-one correspondence with the apartments of X.*

4.2 Lie Theory

In the KSTT paper, the authors work with a particular kind of Lie group in order to find a finitely presented group with $2^{\mathfrak{c}}$-many non-homeomorphic asymptotic cones under the failure of CH. We introduce some of the Lie theory in order to better understand the context in which they work, and for a more thorough understanding of the Lie theory, we point readers to [19] and [22]. We shall also utilize [28] as a reference for some properties of Riemannian manifolds.

Definition 4.17. *We say that a group G is a* **(real) Lie group** *if G is a (real) smooth manifold such that multiplication and inversion are smooth maps.*

It is equivalent to only require that multiplication be a smooth map, and this in theory would simplify computation when determining if a particular structure is a Lie group, but we shall not come across this very often. We give a few examples of common Lie groups.

Example 4.6. *Our simplest example is just $\mathbb{R}^n$ for some $n \in \mathbb{N}_+$ where the group operation is addition. Clearly this is a real manifold, and of course addition is a smooth operation.*

Example 4.7. *Our more interesting examples are $\mathrm{GL}_n(\mathbb{R})$ or $\mathrm{GL}_n(\mathbb{C})$. These may be viewed as open subsets of $\mathbb{R}^{n^2}$ and $\mathbb{R}^{2n^2}$, respectively, and thus are submanifolds of these spaces [22].*

We see that matrix multiplication is smooth simply because the coordinate-wise computations are each polynomials.

In the previous example, we are choosing to view $\mathrm{GL}_n(\mathbb{C})$ as a real Lie group. Typically, one could view this as a *complex* Lie group, which requires that we have a complex analytic manifold structure such that multiplication is holomorphic, but one may also view any complex Lie group as a real Lie group. We give one more concrete example:

Example 4.8. *The groups* $\mathrm{SL}_n(\mathbb{R})$ *and* $\mathrm{SL}_n(\mathbb{C})$ *for* $n \geqslant 2$ *are real Lie groups.*

The groups $\mathrm{SL}_n(\mathbb{R})$ and $\mathrm{SL}_n(\mathbb{C})$ are not just subgroups of $\mathrm{GL}_n(\mathbb{R})$ and $\mathrm{GL}_n(\mathbb{C})$ as abstract groups, but they are also submanifolds of the general linear groups.

As in any mathematics text, it will be useful to understand when two Lie groups are isomorphic. This is, of course, a bit stronger than a simple group homomorphism, and the additional requirement is expected.

Definition 4.18. *Let* G, H *be Lie groups. We say that* $\varphi : G \to H$ *is a* **Lie group homomorphism** *if it is a group homomorphism that is also a smooth map. We further say that* $\varphi : G \to H$ *is a* **Lie group isomorphism** *if it is a Lie group isomorphism that is also a diffeomorphism.*[9]

Given that any Lie group comes equipped with a manifold structure, we want to discuss how one may view this as a *Riemannian* manifold. For those unfamiliar with Riemannian geometry, the following three terms sound like they ought to be synonyms. However, each of them captures a slightly different notion, and we shall be interested in the "Riemannian distance," which will be a metric in the sense of metric spaces.

Definition 4.19. *A* **Riemannian metric** *on a smooth manifold* M *is a smooth covariant 2-tensor field* g *whose value* g_p *at each* $p \in M$ *is an inner product on the tangent space* T_pM*.*

[9] Recall that a diffeomorphism is a differentiable map between manifolds that also has differentiable inverse.

It turns out that every smooth manifold admits a Riemannian metric [28]; thus, we may view any Lie group as a Riemannian manifold, but it is not prima facie clear that the choice of Riemannian metric has any relation to the group structure.

However, we choose not to delve too deep into manifold theory. In the above definition, one should simply think of the tangent space T_pM as some vector space that is "attached" to each point in the manifold, and so a Riemannian metric gives rise to an inner product on these vector spaces. This gives us a natural way to compute the lengths of vectors in the tangent space. For $X_p \in T_pM$, its length is

$$|X_p|_g := g_p(X_p, X_p)^{1/2}$$

as expected. This allows us to rigorously compute the length of a piecewise smooth curve, which in turn will give us a way to define the distance between two points in the manifold.

Definition 4.20. *Let (M, g) be a Riemannian manifold and $\gamma : [a, b] \to M$ a piecewise smooth curve with nonvanishing velocity (assume $a < b$). The* **length of** γ[10] *is defined by*

$$\mathrm{L}_g(\gamma) := \int_a^b |\gamma'(t)|_g \,\mathrm{d}t.$$

This is, essentially, how one calculates the length of a curve in a calculus course. It is not hard to check that this definition is independent of parameterization. Thus, in a path-connected Riemannian manifold, one may find piecewise smooth curves connecting any two points and measure their lengths. It would be nice if every Riemannian manifold were a geodesic space (meaning that there is *some* path of shortest length), but this is not always the case. Thus, we must make the following small adjustment:

[10]While Lee does not use the following terminology, the author has seen other texts which refer to this as the "Riemannian length." This is helpful for reminding the reader that this notion of length is dependent on the choice of Riemannian metric, but it has the unfortunate effect of creating a triplet of terms which seem, to the new learner, to discuss Riemannian "far-ness."

Definition 4.21. *Let (M, g) be a connected*[11] *Riemannian manifold. We define the* **Riemannian distance** $d_g : M \times M \to \mathbb{R}_{\geqslant 0}$ *by*

$$d_g(p, q) := \inf_{\gamma : p \to q} \mathrm{L}_g(\alpha)$$

as γ ranges over all piecewise smooth curves with nonvanishing velocity from p to q.

As with all groups, a Lie group may act on itself by, say, left multiplication. It is natural to wonder how left multiplication affects the Riemannian metric, and therefore the Riemannian distance. As it turns out, all finite-dimensional Lie groups admit *some* Riemannian metric which is left-invariant; we forgo a formal definition of this as we've avoided a rigorous treatment of tensors and manifolds. However, it will be of interest to note that left-invariant Riemannian metrics give rise to left-invariant Riemannian distances, in the sense that, for every $g, x, y \in G$, we have $d_{\mathrm{g}}(x, y) = d_{\mathrm{g}}(gx, gy)$.

We also record the following fact, which will allow us to view any finite-dimensional Lie group as a metric structure.

Proposition 4.5. *Every Lie group G equipped with a left (or right) invariant metric is complete as a metric space and is homogeneous.*

Later, we shall introduce a class of buildings arising from asymptotic cones of Lie groups. Some of the key pieces that we will need in order to identify the building structure will come from the theory of algebraic groups over real closed fields. We first will discuss real closed fields very briefly, and then explain the connection with algebraic groups.

[11] N.B.: If a space is connected and locally path connected, it is globally path connected. As manifolds are locally homeomorphic to $\mathbb{R}^n$, it is easy to see that connected Riemannian manifolds are path connected.

4.3 Real Closed Fields

We shall give a quick treatment of real closed fields and see an important set of examples. For further details on real closed fields from a classical model theoretic perspective, one may consult [30].

Definition 4.22. *Working entirely in classical model theory, let* $\mathcal{L} = (0, 1, +, \cdot)$ *be the language of rings and* $T := \mathrm{Th}(\mathbb{R})$ *the* $\mathcal{L}$*-theory of the reals viewed as an* $\mathcal{L}$*-structure in the usual way. Then a field* $\mathcal{R}$ *is a* **real closed field** *if, when* $\mathcal{R}$ *is viewed as an* $\mathcal{L}$*-structure in the obvious way,* $\mathcal{R} \models T$.

There are several other equivalent definitions for a field to be real closed. For example, one may also say that $\mathcal{R}$ is real closed if it is not algebraically closed but $\mathcal{R}(\sqrt{-1})$ is. Any real-closed field $\mathcal{R}$ may be equipped with a total order $<$ such that $(\mathcal{R}, <)$ is an ordered field. Further, Theorem 3.3.13 of [30] implies that, if $(F, <)$ is an ordered field, then there is a unique (up to unique isomorphism) real closure $F_{\mathbb{R}}$ such that the canonical ordering on $F_{\mathbb{R}}$ extends $<$ on F.[12] Thus, we shall only speak of "the" real closure of F if F is an ordered field.

We give a few examples of real closed fields:

Example 4.9. *One may begin with the ordered field* $(\mathbb{Q}, <)$ *and extend it to its real closure* $\mathbb{Q}_{\mathbb{R}}$. *We note that if* $\mathbb{Q}_{\mathbb{R}}$ *does not contain all real algebraics, then* $\mathbb{Q}_{\mathbb{R}}(i)$ *will not be algebraically closed. Thus we see that* $\mathbb{Q}_{\mathbb{R}}$ *needs to contain all of the real algebraic numbers. It is also known that adjoining* i *to the real algebraic numbers gives an algebraically closed field, hence* $\mathbb{Q}_{\mathbb{R}}$ *consists precisely of the real algebraic numbers.*

Example 4.10. *In this example, we shall explain how one constructs* **Robinson's asymp-**

[12] We comment that the notation $F_{\mathbb{R}}$ is a bit misleading, as F need not be embeddable in $\mathbb{R}$ as $\mathcal{L}$-structures. Indeed, one may take F to be an ordered field where $|F| > |\mathbb{R}|$.

totic field $^\rho\mathbb{R}_\mathcal{U}$. *First, we let* $^*\mathbb{R}_\mathcal{U}$ *be the* classical *ultrapower*

$$^*\mathbb{R}_\mathcal{U} := \prod_{n\in\mathbb{N}}(\mathbb{R}, 0, 1, +, \cdot, <)/\mathcal{U}.$$

We then let $\rho := [(e^{-n})]_\mathcal{U}$*; that is, it is the equivalence class of the sequence* $(e^{-1}, e^{-2}, \dots)$ *in the classical ultrapower.*[13]

We may then define a subring of $^*\mathbb{R}_\mathcal{U}$ *by* $M_0 := \left\{t \in {}^*\mathbb{R}_\mathcal{U} : |t| < \rho^{-k} \text{ for some } k \in \mathbb{N}_+\right\}$ *where* $|t|$ *denotes the equivalence class obtained by taking the absolute value at each index of some representative for* t.

Then, it is not hard to see that $M_1 := \left\{t \in {}^*\mathbb{R}_\mathcal{U} : |t| < \rho^k \text{ for all } k \in \mathbb{N}_+\right\}$ *is the unique maximal ideal of the ring* M_0*, as* M_1 *contains all of the non-units of* M_0*. Indeed, if* $t = [(t_n)]_\mathcal{U} \in M_0$ *is a unit with respect to* M_0*, then for* $\mathcal{U}$*-almost all coordinates,* $t_n \neq 0$*, hence we may find its inverse by* $[(t'_n)]_\mathcal{U}$ *where* $t'_n = t_n^{-1}$ *for all coordinates in which* $t_n \neq 0$ *and anything else otherwise. It then is easy to see that* $[(t'_n)]_\mathcal{U} \notin M_1$*, because if* k *is such that* $|t| < \rho^{-k}$*, then* $|t_n| < e^{kn}$ *for* $\mathcal{U}$*-almost all* n*, and thus* $|t'_n| > e^{-kn}$ *for* $\mathcal{U}$*-almost all* n.

We thus define $^\rho\mathbb{R}_\mathcal{U}$ *to be the field* M_0/M_1*, which comes equipped with a valuation* $\nu : {}^\rho\mathbb{R}_\mathcal{U} \to \mathbb{R} \cup \{\infty\}$ *defined by* $\nu(t) := \mathrm{st}(\log_\rho(|t|))$*, which induces a norm* $\|t\|_\nu := e^{-\nu(t)}$ *(where, by convention,* $e^{-\infty} = 0$*), and thus a metric* $d_\nu(s,t) := \|s - t\|_\nu$.

The above construction is typical, but uses classical ultraproducts. It is also not immediately clear that $^\rho\mathbb{R}_\mathcal{U}$ is real closed; certainly $^*\mathbb{R}_\mathcal{U}$ is real closed by the classical version of Łoś' Theorem, but it is not obvious that this property is preserved after taking a subring and quotient-ing by its maximal ideal.

[13] It general, one may take ρ to be any infinitesimal from $^*\mathbb{R}_\mathcal{U}$.

Luther shows that the *continuous* ultraproduct

$$\prod_{n\in\mathbb{N}_+}(\mathbb{R}, \sqrt[n]{d}, 0, 1, +, \cdot)/\mathcal{U} \cong ({}^\rho\mathbb{R}_\mathcal{U}, d_\nu, 0, 1, +, \cdot)$$

where $\cong$ means $\mathcal{L}$-isomorphism in the continuous language $\mathcal{L} = (0, 1, +, \cdot)$. In particular, these structures are *isometric*. It is important to note that this does *not* immediately tell us that ${}^\rho\mathbb{R}_\mathcal{U}$ shares the same classical first order theory as $\mathbb{R}$. However, using Theorem 3.3.5 of [30], one can axiomatize within continuous logic *most* of the necessary pieces to imply that ${}^\rho\mathbb{R}_\mathcal{U}$ is real closed; namely, we may axiomatize that, for any $a \in {}^\rho\mathbb{R}_\mathcal{U}$ either a or $-a$ is a square, and any polynomial of odd degree has a root in ${}^\rho\mathbb{R}_\mathcal{U}$. It is worth writing out this second part:

$$\varphi_{r,n} := \sup_{\bar{a}}\Big|_{\bar{0}}^{\bar{r}} \inf_{x}\Big|_{s}^{t} \left(a_{2n+1}x^{2n+1} + \cdots + a_1 x + a_0 \dot{=} 0\right)$$

where s is taken to be larger than $\sum |a_i|$. It is worth pointing out that because ${}^\rho\mathbb{R}_\mathcal{U}$ is $\aleph_1$-saturated, the inf quantifier indeed operates as an existential quantifier. Thus, because for each $r \in \mathbb{R}_+$ and $n \in \mathbb{N}$ we have $(\mathbb{R}, \sqrt[n]{d}, 0, 1, +, \cdot) \models \varphi_{r,n}$, then ${}^\rho\mathbb{R}_\mathcal{U} \models \varphi_{r,n}$ as well. This, combined with recognizing that -1 is not the sum of squares in ${}^\rho\mathbb{R}_\mathcal{U}$ (i.e. ${}^\rho\mathbb{R}_\mathcal{U}$ is *formally real*) demonstrates that ${}^\rho\mathbb{R}_\mathcal{U}$ is real closed.

Later, we shall make very brief mention of ${}^\rho\mathcal{O}_\mathcal{U} \subseteq {}^\rho\mathbb{R}_\mathcal{U}$, the *maximal o-convex subring* of Robinson's asymptotic field. If $(X, <)$ is a totally ordered set, we say that $O \subseteq X$ is *order-convex* or *o-convex* if, for any $x, y \in O$ such that $x < y$, if $z \in X$ is such that $x < z < y$, then $z \in O$. As ${}^\rho\mathbb{R}_\mathcal{U}$ is real closed, we may equip it with an ordering $<$ such that $({}^\rho\mathbb{R}_\mathcal{U}, <)$ is an ordered field. It should be clear that, as ${}^\rho\mathcal{O}_\mathcal{U}$ is a proper sub*ring* of Robinson's asymptotic field, all of its elements must be infinitesimal, and it is not hard to see that ${}^\rho\mathcal{O}_\mathcal{U}$ should consist exactly of these elements.

Algebraic Groups

We shall now discuss algebraic groups. The details will be important for understanding how one obtains a root system and hence a Weyl group. We will see how one may obtain a nondiscrete affine $\mathbb{R}$-building from algebraic groups, and the key result from [18] shall be that, in some sense, much of the building structure is maintained as we switch from working over one real closed field to another. The upshot is that we shall see how one may obtain a Weyl group from an algebraic group, and this will be invariant as we work over Robinson's asymptotic field regardless of the choice of ultrafilter.

An algebraic group is, intuitively, an algebraic variety over some field k that is also a group in a precise way. We shall mostly utilize [24] for concepts and results in this section, although readers are encouraged to consult [31] for a more modern, scheme-theoretic approach to the theory of algebraic groups.

Definition 4.23. *An **algebraic group** is an algebraic variety G over an algebraically closed K together with an identity element $e \in G$ and morphisms of varieties for multiplication (denoted μ) and inversion (denoted ι) that make G into a group.*

Example 4.11. *If we fix our field K, then we obtain two important examples of algebraic groups. We denote by $\mathbf{G}_a$ the affine line $\mathbb{A}^1$ together with $\mu(x,y) := x+y$, $e=0$, and $\iota(x) = -x$, and we call this the **additive group**. Similarly, we denote by $\mathbf{G}_m$ the affine open subset $K^* \subseteq \mathbb{A}^1$ with $\mu(x,y) = xy$, $\iota(x) = x^{-1}$, and $e = 1$, and call this the **multiplicative group**.*

Perhaps our most important example is the following:

Example 4.12. *The algebraic group GL_n is given by all $n \times n$ invertible matrices with entries from K, and the group operations are as expected. It is called the **general linear group**.*

The importance of this example is that it allows us to identify several other algebraic groups by finding closed subgroups of the general linear group.[14] In fact, we can see that $\mathbf{G}_m$ may itself be identified with GL_1. The **diagonal subgroup** D_n is a closed subgroup of GL_n simply by requiring the nondiagonal entries be zero, and further, D_n is naturally isomorphic with the n-fold direct product of $\mathbf{G}_m$. This example is also of importance, and receives its own terminology:

Definition 4.24. *We say that an algebraic group is a* ***torus*** *if it is isomorphic to* D_n *for some* $n \in \mathbb{N}_+$*.*

We comment that the maximal tori shall give rise to a Weyl group, in the same sense as we saw when first defining affine $\mathbb{R}$-buildings. Really, our presentation is somewhat reversed in terms of the history, as one develops roots systems and Weyl groups from Lie theory and the theory of algebraic groups and then uses these notions to identify building structure. We have taken this sort of "backwards" approach to remind the reader that, for this version of continuous model theory, we care much more about the metric structure arising from the building metric than we care about, say, the Zariski topology.

In fact, our presentation of algebraic groups is also incomplete (though certainly not backwards!) Humphreys takes the approach of developing the theory of algebraic groups as group objects in the category of algebraic varieties over K, just as we have (briefly) done. But one then may reasonably ask about the "k-rational points" where k is an arbitrary subfield of K. This leads to the notion of a k-group, which is really an algebraic group where G and its morphisms for multiplication and inversion are "defined over k" in a precise sense (cf. §34 of [24]). The notion of being "defined over k" leads to natural defintions such as **k-tori** (which are tori defined over k), and we further say that a **k-split torus** T is a k-torus that is k-isomorphic to d-many copies of $\mathbf{G}_m$ (where $d = \dim T$).

If G is a reductive k-group (which includes all semisimple groups), then it turns out that

[14]N.B.: Here, "closed" refers to being closed in the Zariski topology.

all maximal k-split tori are conjugate under $G(k)$, the k-rational points of G. This means that we may define the k-rank of G to be the common dimension of these maximal k-split tori. Then, if S is a k-split torus (which we will take to be maximal in our setting), then one obtains a k-Weyl group ${}_k\bar{W}$ by ${}_k\bar{W} := N_G(S)/C_G(S)$.

In a moment, we shall construct a building from an algebraic k-group where $k = {}^\rho\mathbb{R}_\mathcal{U}$ and $K = k(\sqrt{-1})$. This may make it seem as though our situation has become horribly complex, as this seems to imply that (under the failure of CH) we have $2^{2^{\aleph_0}}$-many fields k over which to consider k-groups, but we will point out how one obtains isomorphic Weyl groups regardless of ultrafilter $\mathcal{U}$. We state the key result from [18]:

Theorem 4.1. *Let k be a real closed field and $K := k(\sqrt{-1})$ its algebraic closure. Suppose $\mathbf{G}$ is a connected, semisimple algebraic group defined over k contained in $\mathrm{GL}_n(K)$. Then there is an algebraic group $\mathbf{G}_1$ defined over $\mathbb{Q}_\mathbb{R}$ such that $\mathbf{G}$ and $\mathbf{G}_1$ are isomorphic over k.*

Further, if $\mathbf{G}_1 \subseteq \mathrm{GL}_n(K)$ and $\mathbf{G}_2 \subseteq \mathrm{GL}_m(K)$ are connected algebraic groups defined over $\mathbb{Q}_\mathbb{R}$ which are isomorphic over k, then they are isomorphic over $\mathbb{Q}_\mathbb{R}$.

4.4 KSTT Buildings

In the KSTT paper, the authors work with G, which is any connected semisimple Lie group with at least one absolutely simple[15] factor S such that $\mathbb{R}-\mathrm{rank}(S) \geqslant 2$. This implies that G is noncompact. Then, if Γ is a uniform lattice in G (meaning that Γ is a discrete subgroup of G such that G/Γ is compact; these always exist in our context by the Borel-Harish-Chandra Theorem), then it is known that Γ is finitely presented and quasi-isometric to G. This is the necessary setup for the authors to construct an explicit example of a finitely presented group with $2^\mathfrak{c}$-many non-homeomorphic asymptotic cones under the failure of CH.

[15]The notion of a Lie group being "absolutely simple" will not be necessary to understanding our main theorem.

So, in their context, they look specifically at $\text{Cone}_{\mathcal{U}}(\Gamma, e, (\overrightarrow{r_n}))$ where $(\overrightarrow{r_n})$ is a particular increasing sequence of reals with unbounded limit (and hence ultralimit). Due to a plethora of quasi-isometries, they equivocate this with several other asymptotic cones, as a quasi-isometry between two spaces will induce a bi-Lipschitz homeomorphism between their asymptotic cones (with respect to the same ultrafilter) by Proposition 3.9. The difficulty as a model theorist is that our notion of isomorphism is stronger than simply a bi-Lipschitz homeomorphism.

So, to be more precise, we're working with a number of different groups:

- $\mathbf{G}$ is a semisimple algebraic group defined over $\mathbb{Q}_{\mathbb{R}}$, the real closure of $\mathbb{Q}$ in $\mathbb{R}$.
- $\mathbf{G}(\mathbb{R})$ is the group of $\mathbb{R}$-points of $\mathbf{G}$ and is a real Lie group. $\mathbf{G}(\mathbb{R})^\circ$ is its (Hausdorff) connected component[16], and $[\mathbf{G}(\mathbb{R}) : \mathbf{G}(\mathbb{R})^\circ] < \infty$.
- If G is any connected semisimple Lie group, then there exists a semisimple algebraic group $\mathbf{G}$ defined over $\mathbb{Q}_{\mathbb{R}}$ such that $G/Z(G)$ and $\mathbf{G}(\mathbb{R})^\circ$ are isomorphic as Lie groups (see 1.14.6 of [11]). This allows us to define $\mathbb{R}$-rank(G) and G being absolutely simple according to how these are defined for the corresponding algebraic groups. The fact that this is well-defined follows from Theorem 4.1.

Using this notation, it is known that $G/Z(G)$ and G are quasi-isometric. Likewise, if Γ is a uniform lattice in G, then Γ and G are quasi-isometric. Further, in Lie theory, a symmetric space is the quotient G/K of G by a maximal compact subgroup K, and it is known that G, and thus Γ are quasi-isometric to the symmetric space G/K. We thus collect the following quasi-isometries $\sim_{QI}$, all of which induce bi-Lipschitz homeomorphisms between their respective asymptotic cones (of a fixed ultrafilter $\mathcal{U}$ and scaling sequence (r_n)):

$$\mathbf{G}(\mathbb{R})^\circ \cong G/Z(G) \sim_{QI} G \sim_{QI} \Gamma \sim_{QI} G/K.$$

[16]This is the terminology used to denote the connected component containing the identity

Fortunately for us as continuous model theorists, because we are working with connected semisimple Lie groups of noncompact type, if any two such Lie groups are isomorphic, then equipped with their left-invariant Riemannian metrics, they shall be isometric as Riemannian manifolds (possibly after scaling).

Thus, it suffices to work with $\mathbf{G}(\mathbb{R})^\circ$ where $\mathbf{G}$ is some semisimple algebraic group defined over $\mathbb{Q}_\mathbb{R}$. In particular, the KSTT paper works with the semisimple algebraic group SL_m as a clean toy example, which will have an absolutely simple factor S with $\mathbb{R}$-rank$(S) \geqslant 2$ exactly when $m \geqslant 3$. As we mentioned, its $\mathbb{R}$-rank is given by $m - 1$.

We are now prepared to introduce the notion of a KSTT building, which will be the main object of focus for us.

Let $\mathbf{G}$ be a semisimple algebraic group over $\mathbb{Q}_\mathbb{R}$ whose $\mathbb{R}$-rank is at least 2. In [26], they outline a description of how one can construct a building. In particular, they define the set $X := \mathbf{G}({}^\rho\mathbb{R}_\mathcal{U})/\mathbf{G}({}^\rho\mathcal{O}_\mathcal{U})$[17]. Then, we denote by $\mathbf{S}$ a maximal $\mathbb{Q}_\mathbb{R}$-split torus, and the model space is given by $\mathbf{S}({}^\rho\mathbb{R}_\mathcal{U})/\mathbf{S}({}^\rho\mathcal{O}_\mathcal{U})$ whose underlying finite dimensional vector space is viewed as $\mathbb{R}^n$. Again, we may recover the spherical part of the Weyl group by $\bar{W} = N_\mathbf{G}(\mathbf{S})/C_\mathbf{G}(\mathbf{S})$. Finally, the atlas is given by $\mathcal{F} := \{\phi_g : g \in \mathbf{G}({}^\rho\mathbb{R}_\mathcal{U})\}$ where ϕ_g maps $v\mathbf{S}({}^\rho\mathcal{O}_\mathcal{U}) \mapsto gv\mathbf{G}({}^\rho\mathcal{O}_\mathcal{U})$.

They mention that one obtains the same "structure, rank, Tits diagram and building type" when working over an arbitrary real closed field $\mathcal{R}$ as when working over $\mathbb{Q}_\mathbb{R}$. In particular, the type of root system and Weyl group W_T is independent of our choice of real closed field. This is particularly important as we note that branching of apartments may only occur along walls (by (A2)), and as we shall see, each wall has many branches, which makes these fairly complicated buildings.

Definition 4.25. *We say that* $(X, \mathbb{A}, \mathcal{F})$ *is a* **KSTT building** *if it is constructed as above, and we denote it as* $\Delta_{\mathrm{aff}}(\mathbf{G}, {}^\rho\mathbb{R}_\mathcal{U}, {}^\rho\mathcal{O}_\mathcal{U})$.

[17]N.B. This is the set of cosets, as $\mathbf{G}({}^\rho\mathcal{O}_\mathcal{U})$ need not be normal in $\mathbf{G}({}^\rho\mathbb{R}_\mathcal{U})$.

Theorem 4.2. *A KSTT building is a thick nondiscrete affine* $\mathbb{R}$*-building such that the point space is complete (as a metric space) and the atlas is complete.*

Proof. This is a special case of a theorem proved by Kramer and Tent [27]. The thickness was noted by Kleiner and Leeb. [25] ▫

Theorem 4.3. *Suppose that* X *is some symmetric space* $X := G/K$ *where* G *is a semisimple Lie group of rank* $\geqslant 1$*, and* $\mathbf{G}$ *over* $\mathbb{Q}_{\mathbb{R}}$ *has been chosen such that* $\mathbf{G}(\mathbb{R})^{\circ} = G$*. Then* $\mathrm{Cone}_{\mathcal{U}}(X, e, (\overrightarrow{r_n}))$ *is* isometric *to the point space of the affine* $\mathbb{R}$*-building* $\Delta_{\mathrm{aff}}(\mathbf{G}, {}^{\rho}\mathbb{R}_{\mathcal{U}}, {}^{\rho}\mathcal{O}_{\mathcal{U}})$*.*

Proof. This was proved by Kramer and Tent [27], and independently by Thornton [37]. ▫

Now, as model theorists, we shall want to work with the symmetric spaces of semisimple Lie groups as $\mathcal{L}$-structures in our version of continuous model theory. We shall work in $\mathcal{L} = (\star)$, the language with only one constant symbol. When equipped with its bi-invariant Riemannian metric, we note that the Riemannian manifold of a symmetric space is complete as a metric space [22], and we make the convention that $\star$ is always interpreted as the (equivalence class of the) identity element of the Lie group. This allows us to view the symmetric space X of any semisimple Lie group G as an $\mathcal{L}$-structure.

Main Result

We now develop some deeper building theory that will apply to asymptotic cones of symmetric spaces from semisimple Lie groups. As we shall see, some of these results establish that certain trees associated to nondiscrete affine $\mathbb{R}$-buildings satisfy the assumptions of Lemma 3.1.

Definition 4.26. *Let* $(X, \mathbb{A}, \mathcal{F})$ *be a nondiscrete affine* $\mathbb{R}$*-building. We say that* X *is* **thick** *if any wall is the boundary of at least* 3 *half-apartments with disjoint interiors.*

It is known that KSTT buildings are thick (c.f. [25]). It turns out that thickness in a building which is an asymptotic cone will imply "infinite thickness," in that every wall will be the boundary of infinitely many half-apartments. This will follow from a simple argument involving types and the fact that ultraproducts are $\aleph_1$-saturated.

In what follows, we shall work with buildings that are thick, $\aleph_1$-saturated,[18] atlas complete, and equipped with the full affine Weyl group W. Thus, as KSTT buildings satisfy each of these properties, everything that we state shall immediately be true of KSTT buildings. For brevity, we shall call these "nice" buildings:

Definition 4.27. *A nondiscrete affine $\mathbb{R}$-building is said to be* ***nice*** *if it is thick, $\aleph_1$-saturated, atlas complete, and equipped with the full affine Weyl group.*

Lemma 4.1 (Many Apartments)**.** *Suppose that X is a nice building. Given any chart $f : \mathbb{A} \hookrightarrow X$ and half-apartment $H^+_{\alpha,t} \subseteq \mathbb{A}$, there are infinitely many apartments $(A_i)_{i\in\omega}$ such that $A_i \cap f(\mathbb{A}) = f(H^+_{\alpha,t}) = A_i \cap A_j$ when $i \neq j$.*

To prove this, we adapt the proof of Lemma 7.10 from [6]. We first comment on the sketch of their proof, as this illuminates the strategy we employ. In their paper, they specifically work in *bounded* continuous model theory, and so consider pointed metric spaces where all points are at distance $< r$ from the base point p. In showing that any ω-saturated richly branching $\mathbb{R}$-tree must have at least countably many branches at any point a such that $d(p, a) < r$, they argue by contradiction. They pick representatives from each of the finitely many branches, and then along one of these branches, they find a new branch that is ϵ-close to the point a; this ϵ-closeness allows them to construct a type, and any realization of this type *must* belong to a new branch. In our setting, our "branches" look different, most notably because they are higher dimensional! However, we shall see that $\aleph_1$-saturation will help us get around this obstacle.

[18] N.B.: More precisely, because buildings are complete, one may view any building as an $\mathcal{L}$-structure for $\mathcal{L}$ the empty language upon "forgetting" the building structure, and thus saturation makes sense.

Remark 4.3. *We first note that, if we let* $Q_+^k := \{(q_1, \ldots, q_k) \in \mathbb{Q}^k : q_1 \geqslant 0\}$, *and if* $f(H_{\alpha,t})$ *is any hyperplane in a nondiscrete affine* $\mathbb{R}$*-building, then after choosing a basepoint* $x_0 \in H_{\alpha,t}$ *and orthonormal basis* $\mathcal{B} = \{b_1, \ldots, b_k\}$ *of* $\mathbb{R}^k$ *such that* $x_0 + b_1 \in H_{\alpha,t}^+$ *and* $b_1 \perp H_{\alpha,t}$, *there is a "natural" isometric embedding of* Q_+^k *into* $f(H_{\alpha,t}^+)$. *This is induced by mapping* $\vec{0} \mapsto x_0$ *and* $\vec{e}_i \mapsto x_0 + b_i$ *and composing this with the chart* f. *One can then naturally extend this to an isometric embedding of* $\mathbb{Q}^k$ *into any apartment (noting that a hyperplane bounding two half-apartments must be specified).*

Proof. We argue by contradiction, supposing that there are only finitely many such apartments, which we denote as $A_1, \ldots, A_n$ with charts $f_1, \ldots, f_n$, respectively. Again, we may assume that the f_i are such that $f_i(H_{\alpha,t}^+) = f(H_{\alpha,t}^+)$.

Let $Q_+^k := \{(q_1, \ldots, q_k) \in \mathbb{Q}^k : q_1 \geqslant 0\}$ as in the remark. We shall consider sets of parameters and free variables that are indexed by Q_+^k in this "natural" way.

Now fix some basepoint $x_0 \in f(H_{\alpha,t})$, let $\left(a_{\bar{q}}^{(j)}\right)_{\bar{q} \in Q_+^k}$ be the set of points in $f_j(H_{-\alpha,t}^+)$ as described above, and similarly let $\left(a_{\bar{q}}^{(0)}\right)_{\bar{q} \in Q_+^k}$ be the set of points in $f(H_{\alpha,t}^+)$. These will be the parameters for a type p which we will construct below. We comment that there are clearly only countably many parameters.

Our type p will have free variables $(x_{\bar{q}})_{\bar{q} \in \mathbb{Q}^k}$. The formulae in p are as follows:

$$\text{(For each } \bar{q}, \bar{r} \in \mathbb{Q}^k\text{)} \qquad \left|d(x_{\bar{q}}, x_{\bar{r}}) - d_{\mathbb{Q}^k}(\bar{q}, \bar{r})\right| \dot{\leqslant} 1/m \quad \text{(Almost iso. to dense} \subseteq \text{of a flat)}$$

$$\text{(For each } \bar{q} \in Q_+^k\text{)} \qquad d\left(x_{\bar{q}}, a_{\bar{q}}^{(0)}\right) \doteq 0 \qquad \text{(Coincides with } f(H_{\alpha,t}^+)\text{)}$$

$$\text{(For each } \bar{q} \in \mathbb{Q}^k \smallsetminus Q_+^k, \bar{r} \in Q_+^k\text{)} \qquad \left|d\left(x_{\bar{q}}, a_{\bar{r}}^{(j)}\right) - d_{\mathbb{Q}^k}(\bar{q}, \bar{r})\right| \dot{\leqslant} 1/m \quad \text{(}\epsilon\text{-close to new branch)}$$

where m ranges over $\mathbb{N}$ and $j = 1, \ldots, n$. We also add $d(x_{\bar{q}}, \star) \leqslant \lambda_{\bar{q}}$ for sufficiently large $\lambda_{\bar{q}}$.

We clarify some of the above notation. First, we recall that to a continuous model theorist, $\varphi\dot{\leqslant}\psi$ is shorthand for the statement $\max(\varphi - \psi, 0)$, and $\varphi\dot{=}\psi$ is shorthand for $|\varphi - \psi|$. Further, note that in the first and third types of formulae, we are *not* actually including the metric $d_{\mathbb{Q}^k}$ as a formal symbol, but rather one would hypothetically compute this value and write it into the formula.

I claim that this set of formulae is finitely satisfiable. Fix a finite subset of p, and let M be the largest instance of m. Then pick some $0 < \epsilon < 1/M$. We know that there is branching at every special hyperplane, and further that every hyperplane fixed by a reflection from the affine Weyl group is special. Thus, we may find an apartment B given by a chart g such that $g(\mathbb{A}) \cap f_1(\mathbb{A}) = f_1(H^+_{\alpha,t+\epsilon})$. I claim that $g(\mathbb{A})$ satisfies this finite subset of p in the following way: assign the $x_{\bar{q}}$'s with $\bar{q} \in Q^k_+$ to the points $\left(a^{(0)}_{\bar{q}}\right)_{\bar{q}\in Q^k_+}$, and assign the $x_{\bar{r}}$'s with $\bar{r} \in \mathbb{Q}^k \setminus Q^k_+$ to the points $g(\mathbb{A}) \setminus f_1(\mathbb{A})$.

It should be clear that each formula of the second kind is satisfied. Likewise, we note that each formula of the first kind is satisfied, as $g(\mathbb{A})$ contains each of the interpretations of the variables $x_{\bar{q}}$, and we note that the preimages of these interpretations maps to the set $\left\{x_0 + \bar{q} : \bar{q} \in Q^k_+\right\} \cup \{x_0 + (-q_1 - \epsilon, \ldots, q_k) : q_2, \ldots, q_k \in \mathbb{Q}, q_1 \in \mathbb{Q}_{>0}\}$. That is, the interpretations of the free variables are *almost* any isometric embedding of $\mathbb{Q}^k$, except for a "strip" of width ϵ at the branching hyperplane. By using the triangle inequality, it should be clear that, for each $\bar{q}, \bar{r} \in \mathbb{Q}^k$, $\left|d\left(x_{\bar{q}}, x_{\bar{r}}\right) - d_{\mathbb{Q}^k}(\bar{q}, \bar{r})\right| \leqslant \epsilon < 1/m$.

By a similar argument, one can show the satisfaction of formulae of the third kind by finding, for each $j \neq 1$, an apartment made of half apartments $f_j(H^+_{-\alpha,t})$ and $g(H^+_{-\alpha,t})$. This is possible by applying the Exchange Condition to f and f_j. In the case that $j = 1$, one can simply use the apartment constructed from applying the Exchange Condition to f_1 and g.

It follows that the formulae in p are finitely satisfiable, hence p is satisfiable and thus a type. Because p uses only countably many free variables and parameters, and because each formula only mentions finitely many free variables, then p must be satisfied in X by the $\aleph_1$-saturation of ultraproducts.

Suppose that $C := (c_{\bar{q}})_{\bar{q}\in\mathbb{Q}^k}$ satisfies p. Then it should be clear that the completion of C in X is isometric to $\mathbb{R}^k$, hence is a maximal flat and therefore an apartment which we denote as $\bar{C}$. The formulae of the second kind imply that $\bar{C} \cap f(\mathbb{A}) = f(H^+_{\alpha,t})$, and the formulae of the third kind imply that $\bar{C} \cap f_i(\mathbb{A}) = f(H^+_{\alpha,t})$. Thus, $\bar{C}$ is a new apartment satisfying the conditions of the lemma and is not among the $A_1, \ldots, A_n$, contradicting our assumption. ▫

We refer to the next lemma as the "Realized Intersections Lemma." This is analogous to Theorem 2.3 in [2], but adapted to our context of nondiscrete affine $\mathbb{R}$-buildings.

Lemma 4.2 (Realized Intersections)**.** *Suppose that X is a nice building. Let $A \subseteq X$ be an apartment given by some chart f and $S \subseteq A$ be any W_T-convex subset of A. Then there is an apartment $B \subseteq X$ such that $A \cap B = S$.*

Proof. Let $f^{-1}(S) = \bigcap_{i=1}^{k} H^+_{\alpha_i,t_i}$ where $\{\alpha_1, \ldots, \alpha_k\} \subseteq \Phi$. We assume WLOG that the intersection is nontrivial, otherwise we have $A \cap B = S = A$, and we can choose $B = A$. For the simplicity of notation, we write H^+_i for $H^+_{\alpha_i,t_i}$. We prove this by induction on k; that is, as any W_T-convex set is simply the intersection of finitely many half-apartments, we choose to prove by induction on the number of half-apartments involved in the intersection.

The case $k = 1$ is handled by Lemma 4.1. Namely, if $f^{-1}(S) = H^+_{\alpha,t}$, then there are infinitely many A_i such that $A_i \cap f(\mathbb{A}) = f(H^+_{\alpha,t}) = S$, and so we may pick one of them.

Suppose that the statement is true for some arbitrary k. We show that it is true for $k + 1$. That is, suppose that A is an apartment given by chart f and $S \subseteq A$ is some W-convex set such that $f^{-1}(S) = \bigcap_{i=1}^{k+1} H^+_i$. We assume by induction that there is some B_1 given by g_1

such that

$$A \cap B_1 = f\left(\bigcap_{1 \leqslant j \leqslant k} H_j^+\right) = g_1\left(\bigcap_{1 \leqslant j \leqslant k} H_j^+\right).$$

Now, by Lemma 4.1, there is some apartment B_2 given by g_2 such that $B_2 \cap B_1 = g_1(H_{k+1}^+)$, and because there are infinitely many apartments meeting B_1 at $g_1(H_{k+1}^+)$, we also may assume that $g_2(-H_{k+1}^+) \cap f(\mathbb{A}) = \varnothing$ (where $-H_{k+1}^+$ is shorthand for $H_{-\alpha_{k+1},t_{k+1}}^+ \setminus H_{\alpha_{k+1},t_{k+1}}$, the "opposite" pre-half apartment without its boundary wall). Without loss of generality, we assume that g_1 and g_2 agree on H_{k+1}^+. I claim that

$$A \cap B_2 = g_1\left(\bigcap_{1 \leqslant j \leqslant k+1} H_j^+\right) = f\left(\bigcap_{1 \leqslant j \leqslant k+1} H_j^+\right).$$

Note that the second equality holds because g_1 and f agree on $\bigcap_{1 \leqslant j \leqslant k} H_j^+$. We explain the first equality.

Note that

$$\begin{aligned} A \cap B_2 = f(\mathbb{A}) \cap g_2(\mathbb{A}) &= f(\mathbb{A}) \cap (g_2(H_{k+1}^+) \cup g_2(-H_{k+1}^+)) \\ &= \Big(f(\mathbb{A}) \cap g_2(H_{k+1}^+)\Big) \cup \Big(f(\mathbb{A}) \cap g_2(-H_{k+1}^+)\Big). \end{aligned}$$

Then, our expression becomes

$$= \Big(f(\mathbb{A}) \cap g_1(H_{k+1}^+)\Big) \cup \Big(f(\mathbb{A}) \cap g_2(-H_{k+1}^+)\Big) = g_1\left(\bigcap_{1 \leqslant j \leqslant k+1} H_j^+\right) \cup \varnothing.$$

□

Remark 4.4. *We very briefly remark on the notion of parallel walls. As was mentioned before, parallel pre-walls may be determined by the same root. When working in $\mathbb{R}^n$, this is clear, and much of Bennett's work was to extend these notions to the case where our model space is a particular kind of $\mathbb{Q}$-module. In what follows, we shall be discussing parallelism classes of walls, and we refer to [4] for concrete references to certain claims.*

We note that, by Corollary 3.11 of [4], both walls and pre-walls are parallel if and only if they are translates of each other, after we have fixed a common apartment in the case of walls. Technically, if M_1, M_2 are (pre-)walls, there are several different translations which suffice; however, in the case of $\mathbb{R}$-buildings, by basic linear algebra on the model space, we may always choose a translation by a vector v in the direction of the normal vector. Such a translation is unique and $\|v\| = d(M_1, M_2)$, and so in the case of (pre-)walls, we choose this translation by convention.

Lemma 4.3. *Suppose that $(X, \mathbb{A}, \mathcal{F})$ is a nondiscrete affine $\mathbb{R}$-building. Let $M \subseteq X$ be a wall and $x \in X$ a point contained in some wall parallel to M. Then there is a unique wall M_x parallel to M containing x.*

Proof. To show uniqueness, we suppose that M'_x is another wall parallel to M containing x. Then because $M_x \sim M'_x$, by Corollary 3.11 of [4], there is some apartment $g(\mathbb{A})$ containing M_x, M'_x, and their corresponding pre-walls are translates of each other in the model space. In particular, they are determined by the same root, and because $M_x \cap M'_x \neq \varnothing$, we see that their pre-walls must coincide, and $M_x = M'_x$. □

Remark 4.5. *To appreciate our main result, we note the following Theorem 2 from [39], which tells us how one may determine the isomorphism type of a nondiscrete affine $\mathbb{R}$-building using the parallel classes of walls and Weyl simplices. It is then very natural to suspect that elementary equivalence of these parallel classes should give us elementary equivalence of the full buildings, but a precise proof eludes us.*

We briefly comment that, associated to an $\mathbb{R}$-tree, one may define a "projective valuation" ω defined on quadruples of distinct points. Tits defines ω by first letting $\kappa(a, b, c)$ be the unique point amongst $[a, b], [b, c]$, and $[a, c]$ in the $\mathbb{R}$-tree (c.f. Fact 3.2), and then $\omega(a, b; c, d)$ is the signed distance between $\kappa(a, b, c)$ and $\kappa(a, b, d)$, where the sign is $+$ if $\kappa(a, b, c)$ precedes $\kappa(a, b, d)$ on $[a, b]$ (where a precedes b by assumption).

Thus, given some wall $M \subseteq X$, we obtain a natural projective valuation ω_M. A similar result is true for panels in X.

Theorem 4.4 (Tits [39], Theorem 2)**.** *If the spherical building ∂X is thick, then the building $(X, \mathbb{A}, \mathcal{F})$ is determined up to unique isomorphism by the set of valuations ω_M associated to walls of ∂X.*

Of course, in our setting, we cannot expect the same set of valuations between buildings arising as asymptotic cones of symmetric spaces for the simple reason that, under the failure of CH, our buildings will not be isomorphic! However, we are able to show that, if one fixes a wall in two nice buildings and focuses simply on the points contained in some parallel wall, then these "sub-buildings" will be e.e.

Definition 4.28. *Suppose that X is a nondiscrete affine $\mathbb{R}$-building, and fix a wall $M \subseteq X$. Then we define the M-**parallel points** of X to be the set*

$$X_M := \{x \in M' : M' \sim M\}.$$

We are now ready to prove the following:

Theorem 4.5. *Suppose that X and Y are nice buildings with isomorphic model spaces $\mathbb{A}$ and $\mathbb{A}'$, respectively, and let $M \subseteq X$ and $M' \subseteq Y$ be walls. Then $X_M \equiv Y_{M'}$.*

We note that this is distinct from saying that the respective wall trees are elementarily equivalent, as wall trees are specifically $\mathbb{R}$-trees. Intuitively, we are saying that even if we consider a "tree" whose leaves are hyperplanes, then we still maintain elementary equivalence.

Proof. As in Lemma 3.1, we shall not take into account the $r > 0$ bounding the field of play. Fix an EF game of length n with $\epsilon > 0$ as in Corollary 2.1. It suffices to assume that the game consists of all atomic formulae of the form $d(x_i, x_j)$ and $d(\star, x_i)$ for $i, j \in \{1, \ldots, n\}$.

We shall explain Player II's strategy by utilizing her perfect strategy between the wall trees. Namely, because our buildings are nice, we know that each wall is the boundary of infinitely many half apartments, and this corresponds to each point having infinite valency in the wall tree. The fact that each branch contains a ray follows from the fact that any two points are contained in an apartment.

We begin by setting $M_0 := M$ and $M_0' := M$. If $\varphi : \mathbb{A} \to \mathbb{A}'$ is an isomorphism, we may then extend this to $\tilde{\varphi} := \psi \circ \varphi$ where $\psi : \mathbb{A}' \to \mathbb{A}'$ is an isometry sending $\varphi(M_0)$ to M_0'. Player II shall simultaneously play an EF game between the wall trees $\partial_X M_0$ and $\partial_Y M_0'$.

On Player I's k^{th} move, he selects, say, $a_k \in X$. By Lemma 4.3, there is a unique wall M_k containing a_k parallel to M_0, hence M_0 and M_k are of bounded distance $\rho_{0,k}$ from each other. They are contained in a common apartment $A_{0,k}$, and thus Player II may view M_k as the image of M_0 under a translation $\tau_{0,k} : A_{0,k} \to A_{0,k}$ which translates M_0 orthogonally (from the perspective of the model space).

Player II responds as follows:

- She first notes Player I's move as a move in the EF game between $\partial_X M_0$ and $\partial_Y M_0'$.
- She then selects a wall $M_k' \in \partial_Y M_0'$ according to her perfect strategy in the EF game between the wall trees. We let $\tau_{0,k}' : A_{0,k}' \to A_{0,k}'$ denote the translation sending M_k' to M_0' orthogonally.
- Then, she notes that for each $1 \leqslant i \leqslant k$, $\tau_{0,i}^{-1}(a_i) \in M_0$. That is, she may translate each of the plays from X so that they live in the "base wall" M_0.
- Then, using the model space isomorphism $\tilde{\varphi}$, she plays $b_k := \tau_{0,k}' \circ \tilde{\varphi} \circ \tau_{0,k}^{-1}(a_k)$.

We claim that this is a perfect strategy. Fix plays a_i, a_j, b_i, b_j. Then we note that $d(M_i, M_j) = d(M_i', M_j')$ by the fact that Player II utilized a perfect strategy in the EF game between $\partial_X M_0$

and $\partial_Y M_0'$. Further, consider the map $\tau_{i,j} : M_i \to M_j$ defined by

$$\tau_{i,j} := \tau_{0,j} \circ \tau_{0,i}^{-1}\big|_{M_i}$$

and similarly for the other building. That is, we first translate the wall containing the i^{th} play back to the base wall and then translate this to the wall containing the j^{th} play. I claim that $d_X(\tau_{i,j}(a_i), a_j) = d_Y(\tau_{i,j}'(b_i), b_j)$. Indeed, this follows from noting that translations are isometries, hence

$$d_X(\tau_{i,j}(a_i), a_j) = d_X(\tau_{0,j}^{-1} \circ \tau_{i,j}(a_i), \tau_{0,j}^{-1}(a_j)) = d_X(\tau_{0,i}^{-1}(a_i), \tau_{0,j}^{-1}(a_j))$$

$$= d_Y(\tau_{0,j}' \circ \tilde{\varphi} \circ \tau_{0,i}^{-1}(a_i), \tau_{0,j}' \circ \tilde{\varphi} \circ \tau_{0,j}^{-1}(a_j)) = d_Y(\tau_{0,j}' \circ \tilde{\varphi} \circ \tau_{0,i}^{-1}(a_i), b_j)$$

$$= d_Y(\tau_{0,j}' \circ ((\tau_{0,i}')^{-1} \circ \tau_{0,i}') \circ \tilde{\varphi} \circ \tau_{0,i}^{-1}(a_i), b_j) = d_Y(\tau_{i,j}'(b_i), b_j).$$

Then, we note that $\tau_{i,j}$, though defined only on a wall, induces a translation $\bar{\tau}_{i,j}$ on any apartment $A_{i,j}$ containing the walls M_i and M_j, and similarly for the other building. In fact, $\tau_{i,j}$ and $\tau_{i,j}'$ necessarily translate the same distance due to Player II's perfect strategy in the EF game between the wall trees.

Thus, we note that $d_X(\bar{\tau}_{i,j} \circ \tau_{i,j}(a_i), a_j) = d_Y(\bar{\tau}_{i,j}' \circ \tau_{i,j}'(b_i), b_j)$. This essentially is due to the Pythagorean Theorem, as one may compute the distance between two points $x, y \in \mathbb{R}^n$ by computing the distance between two orthogonal hyperplanes H_x and H_y containing x and y respectively, then computing the distance between x and the image of y under the unique orthogonal translation sending H_y to H_x.

Finally, we see that $d_X(\bar{\tau}_{i,j} \circ \tau_{i,j}(a_i), a_j) = d_Y(\bar{\tau}_{i,j}' \circ \tau_{i,j}'(b_i), b_j)$ is exactly $d_X(a_i, a_j) = d_Y(b_i, b_j)$ as desired. □

It is worthwhile to note that, because the walls were arbitrary, this also shows that $X_M \equiv X_{M'}$

for any two walls $M, M' \subseteq X$.

Remark 4.6. *As mentioned earlier, we had suspected that the EF game between the wall trees was enough to deduce the elementary equivalence of two nice buildings with isomorphic model spaces. The caveat is that if one fixes a wall, its set of parallel walls do not necessarily cover the whole building, at least in general. It does not seem likely that this is the case for nice buildings.*

Another earlier attempt at proving the elementary equivalence of nice buildings was to use the (LA) condition to play EF games. In that strategy, Player II may select, at turn k, some apartment $A_k \subseteq \mathcal{A}$ (or $B_k \subseteq \mathcal{B}$) which contains a_k (resp. b_k) and $\Delta_\star C$ for some fixed fundamental Weyl chamber in the respective building. The retraction map guaranteed by (A5) actually has a concrete definition (c.f. Proposition 8.17 of [23]), and it turns out that the retraction onto any A_k (resp. B_k) will be an isomorphism when restricted to any A_j (resp. B_j). This, combined with (A2) and the Realized Intersections Lemma, allows Player II to find some new apartment B_k (resp. A_k) and, using explicit isomorphisms, determine where to place her k^{th} play. It is easy to see that such a strategy shall preserve distances to $\star$ in the respective buildings, but a difficulty arises in trying to show that the distances between *points are preserved if, say, the intersection of their corresponding apartments does not contain a half-apartment, as then the Exchange Condition does not apply.*

Final Remarks

We conclude this dissertation with a few remarks on potential directions for further investigation.

It was not necessary for our main theorem, but as we mentioned in the introduction to this chapter, one may associate to any nondiscrete affine $\mathbb{R}$-building a *spherical building at infinity.* We give a *very* brief overview of spherical buildings, and the curious reader may consult [34] or [1] for specifics of spherical buildings.

Spherical Buildings

Given our model space, we'll first introduce *spherical buildings*, which turn out to be simpler than the nondiscrete affine $\mathbb{R}$-buildings. In fact, a key reason for introducing these first is all nondiscrete affine $\mathbb{R}$-buildings have associated to them a natural spherical building, which is usually termed the "*building at infinity*." It would behoove us to have an idea of what a spherical building is prior to declaring that it exists within some larger building.

There's a lot of theory on spherical buildings, and for a full treatment, one usually begins with chamber systems, Coxeter groups, and then Coxeter complexes. We will skip some of this introductory material as it will not be necessary in this work. We first comment, though, that a finite Coxeter complex of rank n has a geometric realization which is a certain triangulation of the $(n-1)$-sphere. Indeed, in [34], he explains that a Coxeter group W acts faithfully on a particular vector space V of dimension n, and the fixed points of the reflections from W will be hyperplanes, which we may think of as triangulating a sphere S^{n-1} centered at the origin.

In this case, W is a finite group and is a reflection group. Indeed, this is exactly the *spherical Weyl group* we mentioned before.

Definition 4.29. *We say that Δ is a* **building of spherical type** *if it is a simplicial complex that can be expressed as a union of subcomplexes Σ (called* **apartments***) such that:*

(S1) Each Σ is a Coxeter complex.

(S2) For any two simplices $A, B \in \Delta$, there is an apartment Σ containing both A and B.

(S3) If Σ and Σ' are two apartments containing simplices A, B, then there is an isomorphism

$\varphi : \Sigma \to \Sigma'$ fixing A and B pointwise.

With this in mind, we pose a few questions. The converse of the first question is known to be true due to work by Luther in demonstrating the definability of the building at infinity within continuous model theory.

Question 1. *If X and Y are affine Λ-buildings with ∂X and ∂Y their respective buildings at infinity, then does $\partial X \equiv \partial Y$ imply $X \equiv Y$?*

Question 2. *It seems unlikely that one could formulate the building requirements as bona fide $\mathcal{L}$-sentences within this continuous logic. Is there a version of continuous logic (perhaps a second-order continuous logic?) in which one can cleanly axiomatize affine $\mathbb{R}$- or affine Λ-buildings?*

Question 3. *In [26], they mentioned that it is not known whether there exists a subgroup of a complex Lie group with more than one asymptotic cone up to homeomorphism. This question still seems to be unresolved, and it seems natural to conjecture that there is always a unique such asymptotic cone due to the uncountable categoricity of algebraically closed fields.*

Question 4. *Luther was able to show the definability of* a *field using the spherical building at infinity and some projective geometry, but it is not clear that this field is Robinson's asymptotic field, or that the metric from Robinson's asymptotic field is a T-formula (sometimes called a "definable predicate"). Is it true that ${}^{\rho}\mathbb{R}_{\mathcal{U}} \equiv {}^{\rho}\mathbb{R}_{\mathcal{V}}$ for nonprincipal ultrafilters $\mathcal{U}$ and $\mathcal{V}$? One may be able to demonstrate this by showing that nice buildings are e.e. and finishing the demonstration that Robinson's asymptotic field is definable within the class of KSTT buildings.*

Question 5. *Is there an alternative way to show the e.e. of KSTT buildings? Namely, it seems odd that, under CH, there should be a unique asymptotic cone up to homeomorphism (which need not imply that the cones are e.e.), and yet the analysis practically requires one* *to identify the asymptotic cone. Could one find a more general approach demonstrating the elementary equivalence of structures that follow this dichotomy?*

Bibliography

[1] P. Abramenko and K. Brown. *Buildings: Theory and Applications*. Springer, 2010.

[2] P. Abramenko and H. Van Maldeghem. Intersections of apartments. *J. Comb. Theory Ser. A.*, 117(4):440–453, May 2010.

[3] I. Ben Yaacov, A. Berenstein, C. W. Henson, and A. Usvyatsov. Model theory for metric structures. *Model Theory with Applications to Algebra and Analysis*, page 315–427, 2008.

[4] C. D. Bennett. *Affine Λ-buildings*. University of Chicago, Department of Mathematics, 1990.

[5] C. D. Bennett, P. N. Schwer, and K. Struyve. On axiomatic definitions of non-discrete affine buildings. *Advances in Geometry*, 14(3):381–412, Jul 2014.

[6] S. Carlisle and C. W. Henson. Model theory of ℝ-trees. *J. Log. Anal.*, 12, Dec. 2020.

[7] I. Chiswell. *Introduction to Λ-Trees*. World Scientific Publishing, Singapore, Singapore, Dec. 2001.

[8] M. Clay and D. Margalit. *Office Hours with a Geometric Group Theorist*. Princeton University Press, 2018.

[9] C. Druţu. Quasi-isometry invariants and asymptotic cones. *Int. J. Algebra Comput.*, 12(01n02):99–135, Feb. 2002.

[10] A. Dyubina and I. Polterovich. Explicit constructions of universal ℝ-trees and asymptotic geometry of hyperbolic spaces. *Bull. Lond. Math. Soc.*, 33(6):727–734, Nov. 2001.

[11] P. B. Eberlein. *Geometry of nonpositively curved manifolds*. Chicago Lectures in Mathematics. University of Chicago Press, Chicago, IL, Apr. 1997.

[12] S. N. Evans. *Probability and real trees*. Lecture notes in mathematics. Springer, Berlin, Germany, 2008 edition, Oct. 2007.

[13] I. Farah and S. Shelah. A dichotomy for the number of ultrapowers. *Journal of Mathematical Logic*, 10(01n02):45–81, 2010.

[14] I. Goldbring. *Ultrafilters throughout mathematics*. Graduate Studies in Mathematics. American Mathematical Society, Providence, RI, Aug. 2022.

[15] I. Goldbring and B. Hart. On the theories of McDuff's II_1 factors. *Int. Mathem. Res. Not.*, Aug. 2016.

[16] M. Gromov. Groups of polynomial growth and expanding maps (with an appendix by jacques tits). *Publications Mathématiques de l'IHÉS*, 53:53–78, 1981.

[17] M. Gromov. *Asymptotic Invariants of Infinite Groups.* Institut des Hautes Etudes Scientifiques [IHES], 1992.

[18] F. Grosshans. Semi-simple algebraic groups defined over a real closed field. *Amer. J. Math.*, 94(2):473, Apr. 1972.

[19] B. C. Hall. *Lie groups, lie algebras, and representations.* Graduate texts in mathematics. Springer International Publishing, Cham, Switzerland, 2 edition, May 2015.

[20] B. Hart. Math 712 lecture 6 slides, October 2012.

[21] B. Hart. Ehrenfeucht-fraïssé games in continuous logic. preprint on webpage at `https://math.mcmaster.ca/~bradd/publications.html`, 2018.

[22] S. Helgason. *Differential Geometry, Lie Groups, and Symmetric Spaces.* Academic Press, 1978.

[23] P. Hitzelberger. *Generalized affine buildings: Automorphisms, affine Suzuki-Ree Buildings and Convexity.* PhD thesis, Westfälische Wilhelms-Universität Münster, 2009.

[24] J. E. Humphreys. *Linear Algebraic Groups.* Graduate texts in mathematics. Springer, New York, NY, June 2012.

[25] B. Kleiner and B. Leeb. Rigidity of quasi-isometries for symmetric spaces and Euclidean buildings. *Publications Mathématiques de l'IHÉS*, 86:115–197, 1997.

[26] L. Kramer, S. Shelah, K. Tent, and S. Thomas. Asymptotic cones of finitely presented groups. *Advances in Mathematics*, 193(1):142–173, 2005.

[27] L. Kramer and K. Tent. Asymptotic cones and ultrapowers of lie groups. *The Bulletin of Symbolic Logic*, 10(2):175–185, 2004.

[28] J. M. Lee. *Introduction to Riemannian manifolds.* Graduate texts in mathematics. Springer International Publishing, Cham, Switzerland, 2 edition, Jan. 2019.

[29] M. Luther. *Unbounded continuous logic and metric geometry.* PhD thesis, McMaster University, 2015.

[30] D. Marker. *Model Theory : An Introduction.* Graduate Texts in Mathematics. Springer, New York, NY, 2002 edition, Aug. 2002.

[31] J. S. Milne. *Cambridge studies in advanced mathematics: Algebraic groups: The theory of group schemes of finite type over a field series number 170.* Cambridge University Press, Cambridge, England, Oct. 2017.

[32] A. Y. Ol'shanskii, D. V. Osin, and M. V. Sapir. Lacunary hyperbolic groups. *Geom. Topol.*, 13(4):2051–2140, Apr. 2009.

[33] A. Y. Ol'shanskii and M. V. Sapir. A finitely presented group with two non-homeomorphic asymptotic cones. *Int. J. Algebra Comput.*, 17(02):421–426, Mar. 2007.

[34] M. Ronan. *Lectures on Buildings.* Academic Press, 1989.

[35] T. Sider. *Logic for Philosophy.* Oxford University Press, 2010.

[36] S. Thomas and B. Velickovic. Asymptotic cones of finitely generated groups. *Bull. Lond. Math. Soc.*, 32(2):203–208, Mar. 2000.

[37] B. Thornton. *Asymptotic cones of symmetric spaces.* PhD thesis, University of Utah, 2002.

[38] J. Tits. *Buildings of Spherical Type and Finite BN-pairs.* Springer-Verlag, 1974.

[39] J. Tits. Immeubles de type affine. In L. A. Rosati, editor, *Buildings and the Geometry of Diagrams*, pages 159–190, Berlin, Heidelberg, 1986. Springer Berlin Heidelberg.

[40] L. van den Dries and A. J. Wilkie. Gromov's theorem on groups of polynomial growth and elementary logic. *J. Algebra*, 89(2):349–374, Aug. 1984.

Appendix A

Notational Conventions

We list here some of our notational conventions, many of which were chosen in order to simplify what are already lengthy statements and formulae. These conventions are standard throughout the dissertation unless otherwise specified.

1. x, y shall denote single variables unless otherwise specified. Similarly for x_i, y_i, though these shall be single variables within a tuple or sequence of variables.

2. Similarly, a, b shall be used to denote elements of a structure; c shall be reserved for formal constant symbols within our logic.

3. Finite tuples will be represented by a bar, as in $\bar{x} = (x_1, \ldots, x_n)$.

4. A sequence indexed by I shall be written as $(x_i)_{i \in I}$, or, if I is understood, we shall simply write (x_i) and understand that the parentheses indicate that this is a (potentially) infinite sequence of "whatever" rather than specifically a finite tuple.

5. The index of an element within a sequence should indicate its underlying indexing set in the sense that (x_i) shall be indexed by I, (x_j) shall be indexed by J, and (x_n) shall be indexed by $\mathbb{N}$. Thus, given conventions 1., 5., and 6., we *could* write $(n) = (0, 1, 2, \ldots)$.

6. Given the previous convention, $\varphi((x_i);(a_j))$ denotes an $\mathcal{L}$-formula with free variables $(x_i)_{i\in I}$ and parameters $(a_j)_{j\in J}$, though we will not always specify the free variables in an $\mathcal{L}$-formula.

7. When working with many-sorted structures, each variable must be associated to a particular sort symbol, but this shall be understood; in particular, we shall *not* typically write something to the effect of $x_{i,S}$ to denote a variable x_i associated to the sort S.

8. Calligraphic $\mathcal{M}$ and $\mathcal{N}$ denote $\mathcal{L}$-structures and not simply a universe of an $\mathcal{L}$-structure. The universe for a sort S from an $\mathcal{L}$-structure $\mathcal{M}$ shall typically be denoted M_S. However...

9. ...if $\varphi((x_i))$ is an $\mathcal{L}$-formula whose x_i belong to (potentially) different sorts, if we shall want to find some sequence $(a_i)_{i\in I}$ from an $\mathcal{L}$-structure $\mathcal{M}$ such that each a_i comes from the same sort as x_i, we shall typically only write that $(a_i)_{i\in I} \in \mathcal{M}$, or indeed even $(a_i) \in \mathcal{M}$, as $\mathcal{M}$ comprises all universes of sorts, and it shall be understood that each a_i belongs to the universe of the correct sort, as otherwise it would be nonsense to use in the $\mathcal{L}$-formula as given.

10. If $(X, d, \star)$ is a pointed metric space and $x \in X$, we shall write $\|x\|$ to denote $d(x, \star)$.

11. If $(X, d, \star)$ is a pointed metric space, we shall write $B_r(X)$ (resp. $\bar{B}_r(X)$) to denote the open (resp. closed) ball of radius r centered at $\star$. Given the previous convention, this means that

$$B_r(X) := \{x \in X : \|x\| < r\}.$$

12. If $\bar{r}$ is a finite tuple[1] of positive reals of length n and $X := \prod(X_k, d_k, \star_k)$ is a product n-many of pointed metric spaces, then we make the convention

$$B_{\bar{r}}(X) := \left\{(x_1, \ldots, x_n) \in \prod X_k : x_k \in B_{r_k}(X_k) \text{ for each } k = 1, \ldots, n\right\}.$$

[1]Our convention says it must be a finite tuple, but we add this reminder for emphasis.

Even further, if $\bar{x}$ is a finite tuple of variables, and if the associated sort for x_k is S_k, and if $\mathcal{M}$ is an $\mathcal{L}$-structure, then similarly to above, we shall write $\bar{x} \in B_{\bar{r}}(\mathcal{M})$ to mean that $x_i \in B_{r_i}(M_{S_i})$ where M_{S_i} again denotes the universe for the sort S_i in $\mathcal{M}$. This convention also extends to infinite tuples (x_i), but this may not appear in the dissertation.

Appendix B

Towards an EF Strategy from Elementary Equivalence

We present here the challenges that we have faced in trying to prove that, for $\mathcal{L}$-structures $\mathcal{M}$ and $\mathcal{N}$, $\mathcal{M} \equiv \mathcal{N}$ implies that Player II has a winning strategy in all general EF games, and we will present an easy result should this hurdle be overcome.

B.1 Adapting Hart's Lemma

The key piece to demonstrating $\mathcal{M} \equiv \mathcal{N}$ implies that Player II has a winning strategy in all EF games is in showing that ϵ'-equivalence of structures should somehow lead to a winning strategy in any length 1 EF game. Essentially, we are trying to use $\mathcal{L}$-sentences to tell us that we may play on $\mathcal{L}$-formulae, in some sense.

In [21], the way that this is demonstrated is by the following lemma. We note that this is stated for bounded continuous logic, hence the use of the quantifier $\inf_x$ is valid. As we shall see, part of the difficulty of adapting this lemma lies in the quantifiers.

First, we let $\mathcal{L}_{c_1,\ldots,c_n}$ denote the expansion of the language $\mathcal{L}$ by a new constant symbols $c_1, \ldots, c_n$, which do not come with controllers in bounded continuous logic.

Lemma B.1. *Suppose that $F = \{\varphi_1(c), \ldots, \varphi_k(c)\}$ is a finite set of $\mathcal{L}_c$ sentences, and fix $\epsilon > 0$. Then there is a finite set $\tilde{F}$ of $\mathcal{L}$-sentences such that, for any $\mathcal{L}$-structures $\mathcal{M}$ and $\mathcal{N}$, if $\mathcal{M} \equiv_\epsilon^{\tilde{F}} \mathcal{N}$, then for any $a \in \mathcal{M}$, there is $b \in \mathcal{N}$ such that $(\mathcal{M}, a) \equiv_{4\epsilon}^{F} (\mathcal{N}, b)$.*

First, we shall see why this enough to provide Player II with a winning strategy for any general EF game between $\mathcal{M}$ and $\mathcal{N}$, provided that $\mathcal{M} \equiv \mathcal{N}$. Hart gives the argument as follows.

First, suppose we are playing the length n game consisting of $\mathcal{L}$-formulae $\{\varphi_1(\bar{x}), \ldots, \varphi_k(\bar{x})\}$ and $\epsilon > 0$. Then we define a sequence $F_n, \ldots, F_0$ of $\mathcal{L}_{c_1,\ldots,c_i}$-sentences as follows:

- First, $F_n := \{\varphi_1(c_1, \ldots, c_n), \ldots, \varphi_k(c_1, \ldots, c_n)\}$.
- Then, if we have defined F_{i+1} as a finite set of $\mathcal{L}_{c_1,\ldots,c_{i+1}}$-sentences, we let $F_i := \tilde{F}_{i+1}$ from the lemma be our set of $\mathcal{L}_{c_1,\ldots,c_i}$-sentences.
- This results in F_0 a finite set of $\mathcal{L}$-sentences.

Now, Player II may develop a strategy such that, on turn i of her game, she has that

$$(\mathcal{M}, a_1, \ldots, a_i) \equiv_{\epsilon/4^{n-i}}^{F_i} (\mathcal{N}, b_1, \ldots, b_i).$$

If we may guarantee this, then we will have shown that Player II has a winning strategy. This is proved by induction. The base case follows from the fact that $\mathcal{M} \equiv \mathcal{N}$. The inductive case follows from the lemma.

There is nothing in the previous proof which relies on inherent differences between bounded continuous logic versus our setting (aside from ensuring that all plays respect the controllers for the respective constants). Thus, the main hurdle to overcome lies in the proof of Lemma B.1.

We now present Hart's proof of Lemma B.1. His strategy is to define $\theta_S(x)$ where S, intuitively, assigns to each $\varphi_i(x)$ some "guess" of the value of its interpretation, which will be correct up to ϵ. $\theta_S^{\mathcal{M}}(a)$ then is true when S is the correct guess. Thus, our set $\tilde{F}$ is simply the formulae $\inf_x \theta_S(x)$ as S ranges across all (necessarily finitely many) possible guesses. Because there is an "ϵ-correct" guess for any $a \in \mathcal{M}$, we will have $\mathcal{M} \models \inf_x \theta_S(x)$, thus $\mathcal{N} \models \inf_x \theta_S(x) \dot{\leqslant} \epsilon$, and we can find a $b \in \mathcal{N}$ that makes $\theta_S^{\mathcal{N}}(b)$ almost true.

Again, it is worth noting that the following proof is presented in the context of a different logic.

Proof. Pick an ϵ-dense set $r_1^i, \ldots, r_l^i$ in the range of $\varphi_i(x)$. Now, for any arbitrary function $S : \{1, \ldots, k\} \to \{1, \ldots, l\}$, define

$$\theta_S(x) := \max_i(|\varphi_i(x) - r_{S(i)}^i| \dot{\leqslant} \epsilon).$$

We let $\tilde{F}$ be the set of $\mathcal{L}$-sentences $\inf_x \theta_S(x)$ as S ranges across all such functions. Suppose that $\mathcal{M} \equiv_{\epsilon}^{\tilde{F}} \mathcal{N}$, and fix $a \in \mathcal{M}$. Now choose S such that

$$|\varphi_i^{\mathcal{M}}(a) - r_{S(i)}^i| \leqslant \epsilon$$

for each $i = 1, \ldots, k$. We note that such an S must exist because we have chosen ϵ-dense subsets of the images of the φ_i.

By the ϵ-equivalence on $\tilde{F}$, we see that $(\inf_x \theta_S(x))^{\mathcal{N}} \leqslant \epsilon$, hence we may find $b \in \mathcal{N}$ such that $\theta_S^{\mathcal{N}}(b) \leqslant 2\epsilon$. Then we see that $|\varphi_i^{\mathcal{N}}(b) - r_{S(i)}^i| \leqslant 3\epsilon$ for each $i = 1, \ldots, k$, and by Triangle Inequality, we see that $|\varphi^{\mathcal{M}}(a) - \varphi^{\mathcal{N}}(b)| \leqslant 4\epsilon$. □

There are several changes which must be made to this lemma in order for it to work correctly in our setting. Firstly, we note that we may not take arbitrary $a \in \mathcal{M}$ and $b \in \mathcal{N}$, but that we must ensure that a, b respect the controllers for the constant symbols. Truthfully, this is not such a big issue, as the controllers will simply correspond to the field of play for an EF game in our setting anyway. That is to say, requiring that $\|a\| < \lambda_c$ where λ_c is the controller for a new constant symbol is only guaranteeing that an application of our adapted lemma will result in a valid play in the EF game.

Thus, one may think to rewrite the lemma as follows:

Idea 1. *Suppose that $F := \{\varphi_0(c), \ldots, \varphi_k(c)\}$ is a finite set of $\mathcal{L}_c$-sentences where c is some constant symbol with controller λ_c, and fix $\epsilon > 0$. Then there is a finite set $\tilde{F}$ of $\mathcal{L}$-sentences such that for any $\mathcal{L}$-structures $\mathcal{M}$ and $\mathcal{N}$, if $\mathcal{M} \equiv_\epsilon^{\tilde{F}} \mathcal{N}$, then for any $a \in \mathcal{M}$ such that $\|a\| < \lambda_c$, there is $b \in \mathcal{N}$ with $\|b\| < \lambda_c$ such that, as $\mathcal{L}_c$ structures, $(M, a) \equiv_{4\epsilon}^F (N, b)$.*

Perhaps one would need to change 4 to some other number. However, we shall now point out the issues that arise in trying to prove this idea.

We note that each $\varphi_i(x)$ is controlled, and the controller λ_i for $\varphi_i(x)$ will give us a bound on the image of $\varphi_i(x)$ when we restrict our inputs to the ball of radius λ_c. In other words, we may still ask for an ϵ-dense set of $[0, \lambda_i(\lambda_c)]$ and define S and $\theta_S(x)$ as normal.

The issue arises when we define $\tilde{F}$, as our quantifiers change. The first thought is to say that $\tilde{F}$ consists of all $\mathcal{L}$-sentences $\inf_x \big|_{\lambda_c}^{\lambda_c+1} \theta_S(x)$ as S ranges over all possible functions. However, this only tells us that

$$\mathcal{N} \models \left(\inf_x \big|_{\lambda_c}^{\lambda_c+1} \theta_S(x)\right) \dot{\leqslant} \epsilon,$$

and this does *not* guarantee some $b \in \mathcal{N}$ with $\|b\| < \lambda_c$ so that $\theta_S(b) \dot{\leqslant} 2\epsilon$. This is because $\inf\limits_{x \in B_{\lambda_c}(\mathcal{N})} \theta_S(x)$ may, in fact, be arbitrarily large! This is because our quantifiers only take the average inf and sup, and so it could be that, up to the boundary of the field of play, there is no $b \in \mathcal{N}$ which is even close to winning for Player II.

However, we *do* still learn that there is some $t \in (\lambda_c, \lambda_c + 1)$ such that $\inf_{x \in B_t(\mathcal{N})} \theta_S(x) \dot{\leqslant} \epsilon$, and so perhaps we should change the bounds on our quantifier so that λ_c is on top!

We note that the play a must satisfy $\|a\| < \lambda_c$, hence there will be some $\gamma < \lambda_c$ so that $\|a\| < \gamma$, and thus we would see that $\mathcal{M} \models \inf_x \Big|_{\gamma}^{\lambda_c} \theta_S(x)$ as well. However, we are unable to pick finitely many γ_i so that *any* choice of a will satisfy $\|a\| < \gamma_i$ for some i. This leads to the next idea.

Idea 2. *Perhaps there is a logical connective (i.e., a controlled function) which acts as countable disjunction (or approximately does).*

If this were the case, we may instead consider the $\mathcal{L}$-sentences

$$\inf_x \Big|_{\lambda_c - 1/k}^{\lambda_c} \theta_S(x)$$

for sufficiently large k and take their countable disjunction (call this Ψ_S). Then, for some S, we would see that $\mathcal{M} \models \Psi_S$, hence $\mathcal{N} \models \Psi_S \dot{\leqslant} \epsilon$, which should demonstrate that there is *some* k where $\mathcal{N} \models \inf_x \Big|_{\lambda_c - 1/k}^{\lambda_c} \theta_S(x) \dot{\leqslant} \epsilon$ (or perhaps 2ϵ in the case of approximate disjunction), hence there is some $t \in (\lambda_c - 1/k, \lambda_c)$ such that $\inf_{x \in B_t(\mathcal{N})} \theta_S(x) \dot{\leqslant} \epsilon$, and we may proceed as usual.

However, it seems countable disjunction should not exist for this logic. Indeed, what would it mean to say that $f : \mathbb{R}_{\geqslant 0}^{\mathbb{N}} \to \mathbb{R}_{\geqslant 0}$ acts as countable disjunction? It should say that if (r_n) is a sequence of nonnegative reals, then $f((r_n)) = 0$ if and only if at least one of the $r_n = 0$. We record the inexistence of such a controlled function here.

Fact B.1. *There is no controlled function $f : \mathbb{R}_{\geqslant 0}^{\mathbb{N}} \to \mathbb{R}_{\geqslant 0}$ such that $f((r_n)) = 0$ if and only if at least one of the $r_n = 0$.*

Proof. Suppose such an f existed with controllers (λ, N, δ), and fix a sequence (r_n) of positive reals. By assumption, we see that $f((r_n)) > 0$, and we may define $\epsilon := f((r_n))/2$ and a sequence $(\rho_n) := (r_n + 1)$. We thus get a finite set $\Gamma := N((\rho_n), \epsilon)$, and we pick some $k \in \mathbb{N} \setminus \Gamma$.

Now, consider the sequence (r'_n) which is the same as (r_n) except that $r_k = 0$. By the assumption on f, it follows that $f((r'_n)) = 0$. However, because for each $n \in \Gamma$ we have that $\|r_n\|, \|r'_n\| < \rho_n$ and $d(r_n, r'_n) = 0 < \delta((\rho_n), \epsilon)$, we find that

$$|f((r_n)) - f((r'_n))| \leqslant \epsilon,$$

but the LHS is $f((r_n))$ and the RHS is $f((r_n))/2$, a contradiction. □

This is an issue, but perhaps asking for countable disjunction was too strong? We note that, as a sequence in terms of k, we have that

$$\left(\inf_x \Big|_{\lambda_c - 1/k}^{\lambda_c} \theta_S(x)\right)^{\mathcal{M}}$$

is decreasing and is eventually 0. Perhaps there is some other connective which behaves like taking a limit of decreasing sequences?

In Example 2.4.5 of [29], Luther mentions a notion of *forced limits* which have appeared in other works. This also does not work in our context, as one must know that the original limit converges quickly enough in order to conclude that it agrees with the forced limit.

It thus seems a bit hopeless to adapt this lemma to our setting, at least by changing the bounds on our quantifiers and trying to connect them in a nice way.

Our last attempt is to change the number of $\mathcal{L}$-sentences appearing in F and $\tilde{F}$. In our logic, we are no longer confined by finitary connectives, hence we may potentially work with countable or uncountable F and $\tilde{F}$.

Idea 3. *Suppose that* $F := \{\varphi_0(c), \dots, \varphi_k(c), \dots\}$ *is a countable set of* $\mathcal{L}_c$*-sentences where* c *is some constant symbol with controller* λ_c*, and fix* $\epsilon > 0$*. Then there is a countable set* $\tilde{F}$ *of* $\mathcal{L}$*-sentences such that for any* $\mathcal{L}$*-structures* $\mathcal{M}$ *and* $\mathcal{N}$*, if* $\mathcal{M} \equiv_{\epsilon}^{\tilde{F}} \mathcal{N}$*, then for any* $a \in \mathcal{M}$ *such that* $\|a\| < \lambda_c$*, there is* $b \in \mathcal{N}$ *with* $\|b\| < \lambda_c$ *such that, as* $\mathcal{L}_c$ *structures,* $(M, a) \equiv_{4\epsilon}^{F} (N, b)$*.*

The issue here is retaining the countability of $\tilde{F}$. In this case, in defining $\tilde{F}$, we may not allow S to range over all functions $S : \mathbb{N} \to \{1, \dots, l\}$, as this would result in a $\tilde{F}$ of cardinality $2^{\aleph_0}$.

If we instead consider taking all functions $S : \{1, \dots, N\} \to \{1, \dots, l\}$ as we vary $N \geqslant l$, we may carry out most of the proof by setting $\tilde{F}$ to consist of the formulae

$$\inf_x \Big|_{\lambda_c - 1/k}^{\lambda_c} \theta_{S,N}(x)$$

where $\theta_{S,N}(x)$ is essentially the same as $\theta_S(x)$, but where we have specified that the domain of S is $\{0, \dots, N\}$. The issue with this approach is that we may only find $b \in \mathcal{N}$ which works for finitely many formulae at once.

The very last approach we consider is dropping the countability in the rewrite of the lemma. However, this would require *arbitrary* conjunction in order to define $\theta_S(x)$, and it is not clear that arbitrary conjunction is expressible as a controlled function.

B.2 A Conditional Result

Regardless, should an approach be found to prove that $\mathcal{M} \equiv \mathcal{N}$ implies that Player II has a winning strategy for all EF games, we have the following result, which could easily be adapted to other continuous logics.

Theorem B.1. *Let $\mathcal{L}$ be a single sorted language with no relation symbols, and suppose that (G, d) and (H, d') are groups endowed with their word metrics such that they are also $\mathcal{L}$-structures with base points e_G and e_H, respectively. If $(G, d) \equiv (H, d')$, then for any appropriate $\mathcal{U}$ and (r_n), we have*

$$\mathrm{Cone}_{\mathcal{U}}(G, e_G, (r_n)) \equiv \mathrm{Cone}_{\mathcal{U}}(H, e_H, (r_n)).$$

Proof. First, to simplify notation, we let $\mathcal{M} := (G, d, e_G)$ and $\mathcal{N} := (H, d', e_H)$. We first shall show that if $(G, d, e_G) \equiv (H, d', e_H)$, then for any positive ρ, $(G, \frac{d}{\rho}, e_G) \equiv (H, \frac{d'}{\rho}, e_H)$, and we denote their corresponding $\mathcal{L}$-structures by $\mathcal{M}_\rho$ and $\mathcal{N}_\rho$. Fix a finite set Γ of atomic $\mathcal{L}$-formulae with finitely many free variables $x_1, \ldots, x_n$, positive reals r_n, and $\epsilon > 0$. Then, we need to show that Player II has a winning strategy $\mathcal{G}(\mathcal{M}_\rho, \mathcal{N}_\rho, \Gamma, \bar{r}, \epsilon)$.

We note that Player II has a winning strategy in the game $\mathcal{G}(\mathcal{M}, \mathcal{N}, \Gamma, \rho\bar{r}, \rho\epsilon)$ where we have defined $\rho\bar{r} := \rho r_1, \ldots, \rho r_n$. We also note that any strategy for the game $\mathcal{G}(\mathcal{M}, \mathcal{N}, \Gamma, \rho\bar{r}, \rho\epsilon)$ is a valid strategy for the game $\mathcal{G}(\mathcal{M}_\rho, \mathcal{N}_\rho, \Gamma, \bar{r}, \epsilon)$ in the sense that $\|a_i\|_{\mathcal{M}} < \rho r_i$ if and only if $\|a_i\|_{\mathcal{M}_\rho} < r_i$, and similarly for any b_i and the structures $\mathcal{N}$ and $\mathcal{N}_\rho$. Then, because $\mathcal{L}$ has no relation symbols, we see that each of our $\mathcal{L}$-formulae in Γ are of the form $d(s, t)$ for $\mathcal{L}$-terms s, t, and thus we see that

$$\left| d^{\mathcal{M}}(\bar{a}) - d^{\mathcal{N}}(\bar{b}) \right| < \rho\epsilon \iff \left| \frac{d^{\mathcal{M}}(\bar{a})}{\rho} - \frac{d^{\mathcal{N}}(\bar{b})}{\rho} \right| = \left| d^{\mathcal{M}_\rho}(\bar{a}) - d^{\mathcal{N}_\rho}(\bar{b}) \right| < \epsilon.$$

Finally, by the Fundamental Theoremm of Ultraproducts, we see that for any $\mathcal{L}$-sentence φ,

$$\varphi^{\mathcal{M}} = \lim_{\mathcal{U}} \varphi^{\left(G, \frac{d}{r_n}, e_G\right)} = \lim_{\mathcal{U}} \varphi^{\left(H, \frac{d'}{r_n}, e_H\right)} = \varphi^{\mathcal{N}}.$$

□