MODEL THEORY OF REAL-TREES AND THEIR ISOMETRIES

BY

SYLVIA E. B. CARLISLE

B.A., Carleton College, 2002

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Mathematics
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2009

Urbana, Illinois

Doctoral Committee:

Professor Sławomir Solecki, Chair
Professor C. Ward Henson, Director of Research
Associate Professor Ilya Kapovich
Professor Emeritus Carl G. Jockusch

To my family.

ACKNOWLEDGMENTS

Most importantly, many thanks to my adviser Ward Henson for all his guidance and support. This thesis certainly would not have happened without his help. I would also like to acknowledge the great mathematics teachers I was lucky enough to learn from. There are too many to list here. Finally, thanks to my family and friends for their encouragement. This includes especially my mom Barbara, my dad Tom, my brother Ben and my partner Will.

PREFACE

Continuous logic as developed by Ben Yaacov, Berenstein, Henson and Usvyatsov in [1] is an extension of first-order logic that expands the scope of model-theoretic tools to include structures from analysis and geometry in a natural way. It grew out of several earlier approaches to this challenge, including Henson's positive bounded logic for Banach spaces, Ben Yaacov's metric compact abstract theories and $[0, 1]$-valued logics studied by Chang and Keisler.

An $\mathbb{R}$-tree is a uniquely geodesic metric space such that between any two points there is exactly one arc. In this thesis, we begin by exploring the continuous theory of $\mathbb{R}$-trees, denoted $\mathbb{R}$T. We show this theory has a model companion, denoted rb$\mathbb{R}$T, which has nice model theoretic properties. Since $\mathbb{R}$-trees may be unbounded, we use a many-sorted approach. The idea is to take a pointed metric space (M, d, p) and have signature with a sort for each closed ball of diameter n centered at p. We must also have functions between the sorts that will be interpreted as inclusion maps, preserving the relationships between the balls.

We also study theories of $\mathbb{R}$-trees equipped with isometries, and find model companions to those theories. By an isometry of an $\mathbb{R}$-tree M, we mean a surjective distance preserving function from M to M. Isometries of $\mathbb{R}$-trees fall into two categories. If an isometry f of an $\mathbb{R}$-tree M has a fixed point it is called *elliptic*, otherwise it is *hyperbolic*. The quantity $||f|| := \inf_{x \in M} d(x, f(x))$ is called the *translation distance* of f. If $||f|| = 0$, then f is elliptic. If $||f|| > 0$, then f is hyperbolic and acts as a translation along an axis, which is a copy of $\mathbb{R}$ in M. The points on this axis are moved by exactly distance $||f||$. (See [7, 1.3])

In Chapter 1 we outline the basics of continuous logic and model theory of metric structures necessary for this thesis. Much of this material is adapted from [1]. In Section 1.2 we give a many-sorted, bounded continuous signature L_p which we will use to study $\mathbb{R}$-trees. In Section 1.8 we explain more about this signature and its structures, and set up our many-sorted approach to studying unbounded metric spaces. We also address some issues with definability inherent to this many-sorted approach.

In Chapter 2 we define an L_p-theory $\mathbb{R}$T which axiomatizes the class of $\mathbb{R}$-trees (Theorem 2.2.3) and discuss some useful properties of this theory. We then define the class of *richly branching* $\mathbb{R}$-trees (Definition 2.3.1). We axiomatize this class and denote its theory by rb$\mathbb{R}$T (Theorem 2.3.6). Next, we prove that rb$\mathbb{R}$T is the model companion of $\mathbb{R}$T (Theorem 2.5.4). In Section 2.6, we prove some model theoretic properties of rb$\mathbb{R}$T, including quantifier elimination, completeness and stability.

In Chapter 3 we expand the signature L_p to a signature L_s which is suitable for studying $\mathbb{R}$-trees equipped with isometries. Then, for each $r \in \mathbb{R}^{>0}$ we consider a specific class of $\mathbb{R}$-trees equipped with a hyperbolic isometry f, such that in every member of the class, the translation distance satisfies $||f|| \geq r$ (Definition 3.3.1). We give an L_s-theory H$\mathbb{R}$T$_{r,s}$ which axiomatizes this class (Lemma 3.3.3). Then, we show that a model of H$\mathbb{R}$T$_{r,s}$ is existentially closed if and only if its underlying $\mathbb{R}$-tree is a model of rb$\mathbb{R}$T (Theorem 3.4.5). This lets us find the model companion theory rbH$\mathbb{R}$T$_{r,s}$ (Theorem 3.4.7). In Section 3.5 we prove some model theoretic properties of rbH$\mathbb{R}$T$_{r,s}$ and its completions, including quantifier elimination and stability.

In Chapter 4, we study the class of $\mathbb{R}$-trees equipped with an elliptic isometry. We give an L_s-theory E$\mathbb{R}$T$_s$ which axiomatizes this class (Lemma 4.1.3). Then, we give axioms for a theory rbE$\mathbb{R}$T$_s$ (Definition 4.3.3,) and show it is the model companion of E$\mathbb{R}$T$_s$ (Theorem 4.3.13). In Section 4.4, we prove some model theoretic properties of rbE$\mathbb{R}$T$_s$ and its completions, including quantifier elimination and stability.

TABLE OF CONTENTS

CHAPTER 1

CONTINUOUS LOGIC BACKGROUND

1.1 Metric structures and signatures

This section sets out the definitions of continuous structures and signatures used in this thesis. The material here is adapted and summarized from [1]. An alternative presentation is given in [2].

First, some preliminary definitions. Let (M, d_M) and (N, d_N) be metric spaces. Given a function $f\colon M \to N$, a function $\Delta_f\colon (0,1] \to (0,1]$ is called a *modulus of uniform continuity* for f if, for any $\epsilon \in (0,1]$, whenever $d_M(x,y) < \Delta(\epsilon)$ then $d_N(f(x), f(y)) \leq \epsilon$. Also, if $p \in M$ we will denote the closed ball in (M, d_M) of radius n centered at p by $B_n^M(p)$. If the metric space in which this ball resides is clear, we will omit the superscript M.

Note: In this thesis, if (M_i, d_i) are metric spaces for $i = 1, ..., n$, then unless otherwise specified the metric on $M_1 \times ... \times M_n$ is taken to be the maximum metric

$$d((x_1, ..., x_n), (y_1, ..., y_n)) := \max\{d(x_i, y_i) \mid i = 1, ..., n\}.$$

A many-sorted, bounded *metric structure*

$$\mathcal{M} = \Big(((M^{(s)}, d^{(s)}) \mid s \in S), (F_i \mid i \in I), (c_j \mid j \in J), (R_k \mid k \in K)\Big)$$

consists of:

- a family $((M^{(s)}, d^{(s)}) \mid s \in S)$ of complete bounded metric spaces, with $S \neq \emptyset$, called the *sorts* of the structure;
- a family of functions $(F_i \mid i \in I)$, where for each $i \in I$, there are $s_0, ..., s_n \in S$ such that $F_i\colon M^{(s_1)} \times ... \times M^{(s_n)} \to M^{(s_0)}$ is uniformly continuous;
- a family of constants $(c_j \mid j \in J)$, each from a specific sort;

- a family of predicates $(R_k \mid k \in K)$, where for each $k \in K$, there are $s_1, ..., s_n \in S$ such that $R_k \colon M^{(s_1)} \times ... \times M^{(s_n)} \to \mathbb{R}^{\geq 0}$ is a bounded uniformly continuous function.

In order to work with these metric structures, we need appropriate continuous signatures. A many-sorted, bounded *continuous signature* L consists of:

- a non-empty sort index set S and a family of non-negative real numbers $(C_s \mid s \in S)$;
- a family of function symbols $(F_i \mid i \in I)$, each with an arity $(s_1, ..., s_n; s_0)$ where $s_0, ...s_n \in S$, and a modulus of uniform continuity $\Delta_{F_i} \colon (0, 1] \to (0, 1]$;
- a family of constant symbols $(c_j \mid j \in J)$, each with a given arity $s \in S$;
- a family of predicate symbols $(R_k \mid k \in K)$, each with an arity $(s_1, ..., s_n)$ where $s_1, ..., s_n \in S$, a closed bounded interval I_{R_k} in $\mathbb{R}^{\geq 0}$, and a modulus of uniform continuity Δ_{R_k};
- an infinite number of variables of each sort;
- a metric symbol $d^{(s)}$ for each $s \in S$;

where the sets I, J and K are possibly empty.

Given a continuous signature L. An *L-structure*

$$\mathcal{M} = \Big(((M^{(s)}, d^{(s)}) \mid s \in S), (F_i^{\mathcal{M}} \mid i \in I), (c_j^{\mathcal{M}} \mid j \in J), (R_k^{\mathcal{M}} \mid k \in K) \Big)$$

consists of:

- for each $s \in S$, a metric space $(M^{(s)}, d^{(s)})$ with diameter less than or equal to C_s;
- for each $i \in I$, a function $F_i^{\mathcal{M}} \colon M^{(s_1)} \times M^{(s_2)} \times ... \times M^{(s_n)} \to M^{(s_0)}$ that satisfies the given modulus of uniform continuity Δ_{F_i}, where $(s_1, ..., s_n; s_0)$ is the arity of F_i;
- for each $j \in J$, an element $c_j^{\mathcal{M}}$ of $M^{(s)}$, where s is the arity of c_j;
- for each $k \in K$, a predicate $R_k^{\mathcal{M}} \colon M^{(s_1)} \times ... \times M^{(s_n)} \to I_{R_k}$ that satisfies the given modulus of uniform continuity Δ_{R_k}, where $(s_1, ..., s_n)$ is the arity of R_k.

1.2 Main example: pointed metric spaces

A *pointed metric space* is a metric space with a specified basepoint. As an example of a many-sorted continuous signature and its structures, in this section we define a signature L_p suitable for studying pointed metric spaces. It is this signature that we use in this thesis to study $\mathbb{R}$-trees. The pointed metric spaces are allowed to be unbounded, and the signature L_p provides a framework for studying them using many-sorted bounded continuous logic. Section 1.8 contains more about this signature and a discussion of its use.

The continuous signature L_p consists of:

- the sort index set $S = \mathbb{N}$ and a family of positive real numbers $(C_n = 2n \mid n \in \mathbb{N})$;
- for each $m, n \in \mathbb{N}$ with $m \leq n$, a function symbol $I_{m,n}$ with arity $(m; n)$ and modulus of uniform continuity $\Delta_{m,n}(\epsilon) = \epsilon$;
- constant symbols $(p_n \mid n \in \mathbb{N})$, where each p_n has arity n;
- metric symbols $d^{(n)}$ for $n \in \mathbb{N}$.

Given a pointed metric space (M, d, p), the *corresponding L_p-structure* is

$$\mathcal{M} = \Big((M^{(n)}, d^{(n)}) | n \in \mathbb{N}), p_n^{\mathcal{M}}, I_{m,n}^{\mathcal{M}}\Big)$$

where:

- the n^{th} sort $(M^{(n)}, d^{(n)})$ is the closed ball $B_n^M(p)$ of radius n centered at p in (M, d) with $d^{(n)}$ equal to d restricted to $B_n^M(p)$;
- for each $m, n \in \mathbb{N}$ with $m \leq n$, the function $I_{m,n} \colon M^{(m)} \to M^{(n)}$ is the inclusion map;
- the constants p_n are all interpreted as the point p.

1.3 Continuous logic: syntax and semantics

This section presents the syntax and semantics of continuous logic. The material here is adapted from [1]. For this section, let L be a continuous signature and let

$$\mathcal{M} = \Big(((M^{(s)}, d^{(s)}) \mid s \in S), (F_i \mid i \in I), (c_j \mid j \in J), (R_k \mid k \in K)\Big)$$

be an L-structure. Let M stand for the underlying collection of sorts $((M^{(s)}, d^{(s)}) \mid s \in S)$.

Syntax

In continuous logic, L-*terms* are built up inductively from variables and constants using function symbols in the same way as in first order logic. Each term has an associated sort $s \in S$. The definition is as follows:

- any variable or constant of a given sort s is a term of sort s;
- if F is a function symbol of arity $(s_1, ..., s_n; s_0)$, and $t_1, ..., t_n$ are terms of sort $s_1, ..., s_n$ respectively, then $F(t_1, ..., t_n)$ is a term of sort s_0.

Formulas are built inductively as in first order logic, starting with atomic formulas and applying connectives and quantifiers to obtain new formulas. Note that every formula φ will have range contained in a closed bounded interval I_φ contained in $\mathbb{R}^{\geq 0}$, which we must specify. For a formula φ, call I_φ the *range interval of* φ. *Atomic L-formulas* are defined as follows:

- $d^{(s)}(t_1, t_2)$ is an atomic L-formula with range interval $[0, C_s]$, if t_1 and t_2 are terms of sort s;
- $R(t_1, ..., t_n)$ is an atomic L-formula with range interval I_R if R is a predicate symbol with arity $(s_1, ..., s_n)$ and t_i is a term of sort s_i for all $i \in \{1, 2, ...n\}$.

Any continuous function $u\colon [0, \infty)^n \to [0, \infty)$ is a *connective*, and for each sort $s \in S$ and each variable x of sort s we have *quantifiers* $\sup\limits_{s,x}$ and $\inf\limits_{s,y}$. The formal definition of *L-formula* is as follows:

- any atomic L-formula is an L-formula, with range interval as specified above;
- given L-formulas $\varphi_1, ..., \varphi_n$ with range intervals $I_{\varphi_1}, ..., I_{\varphi_n}$ respectively, if $u\colon [0, \infty)^n \to [0, \infty)$ is a connective, then $u(\varphi_1, ..., \varphi_n)$ is an L-formula with range interval $u(I_{\varphi_1}, ..., I_{\varphi_n})$;
- given an L-formula $\varphi(x)$ with free variable x of sort s, $\inf\limits_{s,x} \varphi(x)$ and $\sup\limits_{s,x} \varphi(x)$ are both L-formulas with range interval I_φ.

If the free variables are among $x_1, ..., x_n$, we often denote a term by $t(x_1, ..., x_n)$ to emphasize that fact. Likewise, we may write formulas $\varphi(x_1, ..., x_n)$ if the free variables in a term are among $x_1, ..., x_n$. An *L-sentence* is an L-formula with no free variables. A formula is called *quantifier free* if it is built via the definition of L-formula without using quantifiers.

Semantics

The family $A = (A^{(s)} \mid s \in S)$ is called a *subset* of $((M^s, d^{(s)}) \mid s \in S)$ if $A^{(s)} \subseteq M^{(s)}$ for all $s \in S$. We will often denote this by $A \subseteq M$. Let A be a subset of $((M^{(s)}, d^{(s)}) \mid s \in S)$. We extend the signature L to a signature $L(A)$ by adding a new constant symbol $c(a)$ of the appropriate sort for each $a \in A$. Every L-structure extends canonically to an $L(A)$ structure by taking the interpretation of $c(a)$ to be a for each $a \in A$. We denote this extension by $(\mathcal{M}, a)_{a \in A}$. For the sake of readability, we will write a instead of $c(a)$ when it does not cause confusion.

Let $t(x_1, ..., x_n)$ be an $L(M)$-term of sort s_0, where x_i is of sort s_i for $i = 1, ..., n$. Then $t^{\mathcal{M}}$ denotes the interpretation of t in $\mathcal{M}$. This interpretation $t^{\mathcal{M}}$ is defined as in first order logic (by induction on the definition of L-term) and is a function

$$t^{\mathcal{M}} \colon M^{(s_1)} \times ... \times M^{(s_n)} \to M^{(s_0)}.$$

Next, for each $L(M)$-sentence σ, we define the *value of* σ *in* $\mathcal{M}$ by induction. This value is a real number in I_σ and is denoted $\sigma^{\mathcal{M}}$. Note, the terms in the definition below must have no variables. (So they are terms built up from constants only.)

1.3.1 Definition. 1. $(d^{(s)}(t_1, t_2))^{\mathcal{M}} = d^{(s)}(t_1^{\mathcal{M}}, t_2^{\mathcal{M}})$ for any terms t_1, t_2 of sort s;

2. $(P(t_1, ..., t_n))^{\mathcal{M}} = P^{\mathcal{M}}(t_1^{\mathcal{M}}, ..., t_n^{\mathcal{M}})$ for terms of appropriate sorts;

3. for any $L(M)$-sentences $\sigma_1, ..., \sigma_n$ and any connective $u \colon [0, \infty)^n \to [0, \infty)$,

$$(u(\sigma_1, ..., \sigma_n))^{\mathcal{M}} = u(\sigma_1^{\mathcal{M}}, ..., \sigma_n^{\mathcal{M}});$$

4. for any $L(M)$ formula $\varphi(x)$, where x is of sort s,

$$(\sup_{s,x} \varphi(x))^{\mathcal{M}}$$

is the supremum in I_φ of the set $\{\varphi(a)^{\mathcal{M}} \mid a \in M^{(s)}\}$;

5. for any $L(M)$ formula $\varphi(x)$, where x is of sort s,

$$(\inf_{s,x} \varphi(x))^{\mathcal{M}}$$

is the infimum in I_φ of the set $\{\varphi(a)^{\mathcal{M}} \mid a \in M^{(s)}\}$.

Given an $L(M)$-formula $\varphi(x_1, ..., x_n)$, we use the notation $\varphi^{\mathcal{M}}$ to denote the function from $M^{(s_1)} \times ... \times M^{(s_n)}$ to I_φ defined by

$$\varphi^{\mathcal{M}}(a_1, ..., a_n) = (\varphi(a_1, ..., a_n))^{\mathcal{M}}.$$

Two L-formulas $\varphi(x_1, ..., x_n)$ and $\psi(x_1, ..., x_n)$ are called *logically equivalent* if

$$\varphi^{\mathcal{N}}(x_1, ..., x_n) = \psi^{\mathcal{N}}(x_1, ..., x_n)$$

for any L-structure $\mathcal{N}$. It is a fact that the interpretation $t^{\mathcal{M}}(x_1, ..., x_n)$ of a term $t(x_1, ..., x_n)$, and the interpretation $\varphi^{\mathcal{M}}(x_1, ..., x_n)$ of a formula $\varphi(x_1, ..., x_n)$ are uniformly continuous functions. Moreover (see [1, Theorem 3.5]), for each L-term $t(x_1, ..., x_n)$ and each L-formula $\varphi(x_1, ..., x_n)$ there exist functions Δ_t and Δ_φ from $(0, 1]$ to $(0, 1]$ such that in any L-structure $\mathcal{N}$, the function Δ_t is a modulus of uniform continuity for $t^{\mathcal{N}}(x_1, ..., x_n)$ and Δ_φ is a modulus of uniform continuity for $\varphi^{\mathcal{N}}(x_1, ..., x_n)$.

As we saw above, in continuous logic, the interpretation $\varphi^{\mathcal{M}}$ of a formula $\varphi(x_1, ..., x_n)$ in a structure $\mathcal{M}$ will be a function mapping n-tuples from the underlying metric spaces into a closed bounded interval I_φ. This is in contrast to the situation of first order logic, where the interpretation of a formula is a function that maps an n-tuple from the underlying set of a structure $\mathcal{N}$ into the discrete set {true, false}.

We make the following syntactic definition. An *L-condition* is a statement of the form $\varphi = r$ where φ is a formula and $r \in \mathbb{R}^{\geq 0}$. A *closed L-condition* is an L-condition $\varphi = r$ such that φ is an L-sentence. We often denote an L-condition $\varphi(x_1, ..., x_n) = r$ by $E(x_1, ..., x_n)$ if the free variables of the formula are among $x_1, ..., x_n$. For an L-condition $E(x_1, ..., x_n) = \big(\varphi(x_1, ..., x_n) = r\big)$ and $a_1, ..., a_n \in M^{(s_1)} \times ... \times M^{(s_n)}$, we say that $E(x_1, ..., x_n)$ is *true of* $a_1, ..., a_n$ *in* $\mathcal{M}$ if $\varphi^{\mathcal{M}}(a_1, ..., a_n) = r$. This is denoted $\mathcal{M} \models E[a_1, ..., a_n]$.

1.4 Model theory basics

In this section, we present some basic model theoretic concepts, adapted from [1]. For this section, let L be a continuous signature and let $\mathcal{M}$ and $\mathcal{N}$ be L-structures with sorts $(M^{(s)} \mid s \in S)$ and $(N^{(s)} \mid s \in S)$ respectively. We say $\mathcal{M}$ is a *substructure* of $\mathcal{N}$, denoted $\mathcal{M} \subseteq \mathcal{N}$, if:

- $M^{(s)} \subseteq N^{(s)}$, and $d^{(s)}$ on $M^{(s)}$ is the restriction of $d^{(s)}$ on $N^{(s)}$ for all $s \in S$;
- $F^{\mathcal{N}}$ extends $F^{\mathcal{M}}$ for every function symbol F in L;
- $R^{\mathcal{N}}$ extends $R^{\mathcal{M}}$ for every predicate symbol R in L;
- $c^{\mathcal{N}} = c^{\mathcal{M}}$ for every constant symbol c in L.

The structures $\mathcal{M}$ and $\mathcal{N}$ are *isomorphic* if there is a family of surjective isometries $(g^{(s)}\colon M^{(s)} \to N^{(s)} \mid s \in S)$ with the following properties:

- for each function symbol F of arity $(s_1, ..., s_n; s_0)$ in L, and any $a_1 \in M^{(s_1)}, ..., a_n \in M^{(s_n)}$

$$g^{(s_0)}(F^{\mathcal{M}}(a_1, ..., a_n)) = F^{\mathcal{N}}(g^{(s_1)}(a_1), ..., g^{(s_n)}(a_n));$$

- for each constant symbol c with arity s in L,

$$g^{(s)}(c^{\mathcal{M}}) = c^{\mathcal{N}};$$

- for each predicate symbol R in L of arity $(s_1, ..., s_n)$ in L and any $a_1 \in M^{(s_1)}, ..., a_n \in M^{(s_n)}$,

$$R^{\mathcal{N}}(g^{(s_1)}(a_1), ..., g^{(s_n)}(a_n)) = R^{\mathcal{M}}(a_1, ..., a_n).$$

Such a family of maps is called an *isomorphism* from $\mathcal{M}$ onto $\mathcal{N}$.

An *L-theory* is a set of closed L-conditions. If T is an L-theory, then an L-structure $\mathcal{M}$ is called a *model* of T, denoted $\mathcal{M} \models T$, if $\mathcal{M} \models E$ for every $E \in T$. If $\mathcal{M}$ is an L-structure, we define the L-theory of $\mathcal{M}$ to be the set of closed L-conditions $\mathrm{Th}(\mathcal{M}) := \{E \mid \mathcal{M} \models E\}$. An L-theory T that equals $\mathrm{Th}(\mathcal{M})$ for some L-structure $\mathcal{M}$ is called a *complete* theory. (Note: this is not to be confused with the fact that the underlying metric spaces for our structures are all metrically complete.) Given a class $\mathcal{K}$ of L-structures, we often study the theory of that class, $\mathrm{Th}(\mathcal{K}) := \{E \mid \mathcal{M} \models E \text{ for all } \mathcal{M} \in \mathcal{K}\}$. A class of L-structures $\mathcal{K}$ is *axiomatizable* in L if there exists a set Σ of closed L-conditions such that $\mathcal{K}$ is exactly the class of all models of Σ.

The L-structures $\mathcal{M}$ and $\mathcal{N}$ are called *elementarily equivalent*, denoted $\mathcal{M} \equiv \mathcal{N}$, if for every closed L-condition E, $\mathcal{M} \models E$ if and only if $\mathcal{N} \models E$. In other words, $\sigma^{\mathcal{M}} = \sigma^{\mathcal{N}}$ for every L-sentence σ. A family of maps $(g^{(s)}\colon M^{(s)} \to N^{(s)} \mid s \in S)$ is called an *elementary embedding* if

- each $g^{(s)}$ is an isometric embedding;
- for all formulas $\varphi(x_1, ..., x_n)$ where $x_1, ..., x_n$ are suitable variables, and for any $a_1, ..., a_n$ from $M^{(s_1)}, ..., M^{(s_n)}$ respectively, we have

$$\varphi^{\mathcal{M}}(a_1, ..., a_n) = \varphi^{\mathcal{N}}(g^{(s_1)}(a_1), ..., g^{(s_n)}(a_n)).$$

Note that an isomorphism is always an elementary embedding. If $\mathcal{M} \subseteq \mathcal{N}$ and the inclusion maps on the sorts constitute an elementary embedding, then $\mathcal{M}$ is called an *elementary substructure* of $\mathcal{N}$, and $\mathcal{N}$ is an *elementary extension* of $\mathcal{M}$, denoted $\mathcal{M} \preceq \mathcal{N}$.

Instead of using cardinality to measure the size of structures, in continuous logic the density character of the underlying metric space is used. In our many-sorted setting, we will define the *density character* of an L-structure to be the list $\kappa := (\kappa_s \mid s \in S)$ of the density characters of its sorts. An L-theory T is *κ-categorical* if any two L-structures with the same density character κ are isomorphic.

An L-theory T has *quantifier elimination* if, for any formula $\varphi(x_1, ..., x_m)$, there exist quantifier free formulas $\{\psi_n(x_1, ..., x_m) \mid n \in \mathbb{N}\}$ such that for any $\epsilon > 0$, there exists $N \in \mathbb{N}$ such that for all $n \geq N$ and any $\mathcal{M} \models T$

$$|\varphi^{\mathcal{M}}(x_1, ..., x_m) - \psi_n^{\mathcal{M}}(x_1, ..., x_n)| \leq \epsilon.$$

1.5 Ultraproducts

In this section we define the ultraproduct of a family of L-structures, and give some results concerning of ultraproducts. For this section, let L be a continuous signature.

Let X be a topological space. Let $x_i \in X$ for all $i \in I$ and $x \in X$. Let U be an ultrafilter on I. Then, the *U-ultralimit* of $(x_i)_{i\in I}$ is defined to be

$$\lim_{U,i} x_i = x$$

if for any open set $\mathcal{O}$ containing x, the set $\{i \in I \mid x_i \in \mathcal{O}\}$ is in U.

Next, the *ultraproduct* for bounded metric spaces is defined. (See [1, Chapter 5].) Say $((M_i, d_i) \mid i \in I)$ is a family of uniformly bounded metric spaces indexed by the infinite set I, and let U be an ultrafilter on I. We define $\prod_U M_i = (N, d)$, the ultraproduct over U of $((M_i, d_i) \mid i \in I)$ as follows.

Let $\widehat{N}$ be the cartesian product $\prod_{i\in I} M_i$. Define a pseudometric $\hat{d}$ on $\widehat{N}$ by $\hat{d}((a_i), (b_i)) = \lim_{U,i} d_i(a_i, b_i)$. Let (N, d) be the result of taking the quotient of $(\widehat{N}, \hat{d})$ by the equivalence relation defined by $x \sim_U y$ if and only if $\hat{d}(x, y) = 0$. Let $[a_i]_U$ denote the equivalence class of the sequence $(a_i)_{i\in I}$.

Now that we have a notion of ultraproduct for metric spaces, we define how to take ultraproducts of metric structures. Say $(\mathcal{M}_i \mid i \in I)$ is a family of L-structures indexed by an infinite set I. Let U be a non-principal ultrafilter on I. We define $\mathcal{N} = \prod_U \mathcal{M}_i$, the ultraproduct over U of metric structures $(\mathcal{M}_i \mid i \in I)$ as follows.

- For each sort $s \in S$, let $(N^{(s)}, d^{(s)})$ be the ultraproduct $\prod_U M_i^{(s)}$. (Note that this is an ultraproduct of bounded metric spaces.)
- For each function symbol $F \in L$ and $(a_i^j)_{i \in I}$ for $j = 1, ..., n$ (where the a_i^j are from appropriate sorts) let

$$F^{\mathcal{N}}([a_i^1]_U, ..., [a_i^n]_U) = [F^{\mathcal{M}_i}(a_i^1, ..., a_i^n)]_U.$$

  Note that this is a function from $\prod_U M_i^{(s_1)} \times ... \times \prod_U M_i^{(s_n)} = N^{(s_1)} \times ... \times N^{(s_n)}$ to $\prod_U M_i^{(s_0)} = N^{(s_0)}$, and it satisfies modulus of uniform continuity Δ_F, because each of the factors $F^{\mathcal{M}_i}$ satisfies that modulus (see [1, Chapter 5]).
- For each constant symbol $c \in L$, let $c^{\mathcal{N}} = [c^{\mathcal{M}_i}]_U$.
- For each predicate symbol $R \in L$ and $(a_i^j)_{i \in I}$ for $j = 1, ..., n$ (where the a_i^j are from appropriate sorts) let

$$R^{\mathcal{N}}(([a_i^1]_U, ..., [a_i^n]_U) = \lim_{U,i} R^{\mathcal{M}_i}(a_i^1, ..., a_i^n).$$

  This predicate is a uniformly continuous function from $\prod_U M_i^{(s_1)} \times ... \times \prod_U M_i^{(s_n)} = N^{(s_1)} \times ... \times N^{(s_n)}$ to I_R which satisfies the modulus of uniform continuity Δ_R (see [1, Chapter 5]).

As demonstrated by the following theorem and its corollary, ultraproducts of metric structures are useful in continuous logic for building elementary extensions. An *ultrapower* of a structure $\mathcal{M}$ is an ultraproduct of the family $(\mathcal{M}_i \mid i \in I)$ where $\mathcal{M}_i = \mathcal{M}$ for all $i \in I$.

1.5.1 Theorem (Fundamental theorem of ultraproducts)**.** *([1, Theorem 5.4]) Let $(\mathcal{M}_i \mid i \in I)$ be a family of L-structures. Let U be an ultrafilter on I and let $\mathcal{M} = \prod_U \mathcal{M}_i$ be the U-ultraproduct of $(\mathcal{M}_i \mid i \in I)$. Let $\varphi(x_1, ..., x_n)$ be an L-formula. If $a_j = [a_i^j]_U$ are elements of $\mathcal{M}$ for $j = 1, ..., n$, then*

$$\varphi^{\mathcal{M}}(a_1, ..., a_n) = \lim_{i,U} \varphi^{\mathcal{M}_i}(a_i^1, ..., a_i^n).$$

1.5.2 Corollary. *([1, Corollary 5.5]) Let $\mathcal{N}$ be an L-structure and $\prod_U \mathcal{N}$ be the ultrapower of $\mathcal{N}$. Let j be the embedding of N into $\prod_U N$ which sends x to $[x]_U$. Then j is an elementary embedding of $\mathcal{N}$ into $\prod_U \mathcal{N}$.*

1.6 Types and stability

This section contains definitions and results adapted from Chapters 8 and 14 of [1]. Let L be a continuous signature and $\mathcal{M}$ an L-structure. Let $A \subseteq M$. Let $b_1 \in M^{(s_1)}, ..., b_n \in M^{(s_n)}$. The *type of* $b_1, ..., b_n$ *over* A *in* $\mathcal{M}$, denoted $\mathrm{tp}_{\mathcal{M}}(b_1, ..., b_n/A)$, is the set of $L(A)$-conditions $E(x_1, ..., x_n)$ such that $(\mathcal{M}, a)_{a \in A} \models E[b_1, ..., b_n]$. We say this type has arity $(s_1, ..., s_n)$. Let T_A be a complete $L(A)$-theory. Denote by $S_{s_1,...,s_n}(T_A)$ the collection of types

$$\{\mathrm{tp}_{\mathcal{M}}(b_1, ..., b_n/A) \mid (\mathcal{M}, a)_{a \in A} \models T_A,\ b_1 \in M^{(s_1)}, ..., b_n \in M^{(s_n)}\}.$$

If $A = \emptyset$ we omit it from the notation, and where T is clear we may write $S_{s_1,...,s_n}(A)$ for $S_{s_1,...,s_n}(T_A)$. Given a tuple of sorts $(s_1, ..., s_n)$, we define a metric on $S_{s_1,...,s_n}(A)$ as follows:

$$d(p, q) = \inf\{\max_j d^{\mathcal{M}_A}(b_j, c_j) \mid \mathcal{M}_A \models p[b_1, ..., b_n],\ \mathcal{M}_A \models p[c_1, ..., c_n]\}$$

where $\mathcal{M}_A$ varies over all models of T_A. It is straightforward to see that this is a metric on $S_{s_1,...,s_n}(A)$. We call it the *d-metric* on the type space $S_{s_1,...,s_n}(A)$. The set $S_{s_1,...,s_n}(A)$ with the corresponding d-metric is a complete metric space. (See [1, Proposition 8.8].)

Let T be a complete L-theory and let λ be an infinite cardinal. The theory T is λ-*stable* if for any $\mathcal{M} \models T$, for any $A \subseteq M$ with $|A| \leq \lambda$, for every sort s the space $S_s(T_A)$ has density character $\leq \lambda$ with respect to the d-metric. We say T is *stable* if it is λ-stable for some λ. For more about types and stability in continuous logic see [1] and [2].

1.7 More model theory

This section presents more definitions and results of model theory for metric structures, adapted from [1]. For this section let L be a continuous signature, T an L-theory and $\mathcal{M}$ an L-structure. A set $\Sigma(x_1, ..., x_n)$ of L-conditions (with free variables among $x_1, ..., x_n$)is called *satisfiable in* $\mathcal{M}$ if there exist $a_1, ..., a_n$ in $\mathcal{M}$ such that $\mathcal{M} \models E[a_1, ..., a_n]$ for every $E(x_1, ..., x_n) \in \Sigma$. Let κ

be a cardinal. A model $\mathcal{M}$ of T is called *κ-saturated* if for any set of parameters $A \subseteq M$ with cardinality $< \kappa$ and any set $\Sigma(x_1, ..., x_n)$ of $L(A)$-conditions, if every finite subset of $\Sigma(x_1, ..., x_n)$ is satisfiable in $(\mathcal{M}, a)_{a \in A}$, then the entire set $\Sigma(x_1, ..., x_n)$ is satisfiable in $(\mathcal{M}, a)_{a \in A}$.

1.7.1 Proposition. *([1, Proposition 7.6]) For any countably incomplete ultrafilter U on I, the U-ultraproduct of a family of L-structures $(\mathcal{M}_i \mid i \in I)$ is ω_1-saturated.*

Note that any non-principal ultrafilter on $\mathbb{N}$ is countably incomplete.

1.7.2 Proposition. *([1, Proposition 7.10]) For any cardinal κ, any L-structure $\mathcal{M}$ has a κ-saturated elementary extension.*

Saturated structures have many useful properties. For example, in an ω-saturated structure all quantifiers are realized exactly. The proposition below captures this idea. It is stated for structures with just one sort, but an analogous statement holds for many-sorted structures.

1.7.3 Proposition. *([1, Proposition 7.7]) Let $\mathcal{M}$ be an L-structure and suppose $E(x_1, ..., x_m)$ is the L-condition*

$$(Q^1_{y_1} ... Q^n_{y_n} \varphi(x_1, ..., x_m, y_1, ..., y_n)) = 0$$

where each Q^i is either inf *or* sup *and φ is quantifier free. Let $\mathcal{E}(x_1, ..., x_m)$ be the mathematical statement*

$$\widetilde{Q}^1_{y_1} ... \widetilde{Q}^n_{y_n} (\varphi(x_1, ..., x_m, y_1, ..., y_n) = 0)$$

where each $\widetilde{Q}^i$ is $\exists y_i$ if $Q^i_{y_i}$ is $\inf_{y_i}$ and is $\forall y_i$ if $Q^i_{y_i}$ is $\sup_{y_i}$. If $\mathcal{M}$ is ω-saturated, then for any elements $a_1, ..., a_m$ of M, we have $\mathcal{M} \models E[a_1, ..., a_m]$ if and only if $\mathcal{E}(a_1, ..., a_n)$ is true in M.

If Λ is a linearly ordered set, a *Λ-chain of L-structures* is a family of L-structures $(\mathcal{M}_\lambda \mid \lambda \in \Lambda)$ such that $\mathcal{M}_\alpha \subseteq \mathcal{M}_\beta$ for all $\alpha < \beta < \lambda$. The L-structure $\mathcal{M}$ with each sort $M^{(s)}$ equal to the completion of the union $\bigcup_{\lambda \in \Lambda} M^{(s)}_\lambda$ is called the *union of the chain* and denoted $\bigcup_{\lambda \in \Lambda} \mathcal{M}_\lambda$. A chain of structures $(\mathcal{M}_\lambda \mid \lambda \in \Lambda)$ is an *elementary chain* if $\mathcal{M}_\alpha \preceq \mathcal{M}_\beta$ for all $\alpha < \beta$. A class $\mathcal{K}$ of L-structures is called *inductive* if it is closed under unions of chains. We say T is a *$\forall\exists$-theory* if T is an L-theory axiomatized by closed L-conditions of the form

$$\sup_{s_1, x_1} ... \sup_{s_k, x_k} \inf_{t_1, y_1} ... \inf_{t_l, y_l} \psi(x_1, ..., x_k, y_1, ..., y_l)$$

where ψ is quantifier free.

1.7.4 Proposition. *The class of models* Mod(T) *of a* $\forall\exists$*-theory* T *is an inductive class.*

Proof. Let T be an $\forall\exists$-theory. Let $\sup\limits_{s_1,x_1} ... \sup\limits_{s_k,x_k} \inf\limits_{t_1,y_1} ... \inf\limits_{t_l,y_l} \psi(x_1,...,x_k,y_1,...,y_l) = 0$ be an axiom of T, where ψ is quantifier free. Let Λ be a linearly ordered set and let $(\mathcal{M}_\lambda \mid \lambda \in \Lambda)$ be a chain of L-structures from Mod(T). Let $\mathcal{M} = \bigcup_{\lambda\in\Lambda} \mathcal{M}_\lambda$ be the union of the chain. For any $a = a_1,...,a_k \in \mathcal{M}$ from suitable sorts, there exists $\lambda \in \Lambda$ so that $a_1,...,a_k \in \mathcal{M}_\lambda$. The $L(a)$-condition $\inf\limits_{t_1,y_1} ... \inf\limits_{t_l,y_l} \psi(a_1,...,a_k,y_1,...,y_l) = 0$ is true in $\mathcal{M}_\lambda$, since $\mathcal{M}_\lambda \models T$. Then in $\mathcal{M}_\lambda$ there exist witnesses to the fact that $\inf_{t_1,y_1} ... \inf_{t_l,y_l} \psi(a_1,...,a_k,y_1,...,y_l) = 0$ in $\mathcal{M}_\lambda$. That is for all $\epsilon > 0$ there exist $b_1^\epsilon,...,b_l^\epsilon \in \mathcal{M}_\lambda$ such that $\psi^{\mathcal{M}}(a_1,...,a_k,b_1^\epsilon,...,b_l^\epsilon) \leq \epsilon$. We know $\mathcal{M}_\lambda \subseteq \mathcal{M}$, and therefore these witnesses are also in $\mathcal{M}$. Thus, $\inf_{t_1,y_1} ... \inf_{t_l,y_l} \psi(a_1,...,a_k,y_1,...,y_l) = 0$ is true in $\mathcal{M}$. The $a_1,...,a_k \in \mathcal{M}$ were arbitrary, so $\sup\limits_{s_1,x_1} ... \sup\limits_{s_k,x_k} \inf\limits_{t_1,y_1} ... \inf\limits_{t_l,y_l} \psi(x_1,...,x_k,y_1,...,y_l) = 0$ is true in $\mathcal{M}$. Therefore, $\mathcal{M} \models T$. □

1.7.5 Proposition. *([1, Proposition 7.2]) If* $(\mathcal{M}_\lambda \mid \lambda \in \Lambda)$ *is an elementary chain and* $\lambda \in \Lambda$*, then* $\mathcal{M}_\lambda \preceq \bigcup_{\lambda\in\Lambda} \mathcal{M}_\lambda$.

Next, we discuss existential closure, model completeness, and the definition of a model companion.

1.7.6 Definition. An inf*-formula* of L is a formula of the form

$$\inf_{s_1,y_1} ... \inf_{s_n,y_n} \varphi(x_1,...,x_k,y_1,...,y_n)$$

where $\varphi(x_1,...,x_k,y_1,...,y_n)$ is quantifier-free.

1.7.7 Definition. Let T be an L-theory for a given continuous signature L. Let $\mathcal{M} \models T$. We say $\mathcal{M}$ is an *existentially closed (e.c.)* model of T if, for any inf-formula

$$\inf_{s_1,y_1} ... \inf_{s_n,y_n} \varphi(x_1,...,x_k,y_1,...,y_n)$$

any $a_1,...,a_k \in M$ and any $\mathcal{N} \models T$ that is an extension of $\mathcal{M}$ we have:

$$\inf_{s_1,y_1} ... \inf_{s_n,y_n} \varphi(a_1,...,a_k,y_1,...,y_n)^{\mathcal{N}} = \inf_{s_1,y_1} ... \inf_{s_n,y_n} \varphi(a_1,...,a_k,y_1,...,y_n)^{\mathcal{M}}.$$

An L-theory T is *model complete* if any embedding between models of T is an elementary embedding.

The following lemma and proposition are stated for a signature L with a single sort, but are easily extended to the many-sorted case.

1.7.8 Lemma. *If $\mathcal{M} \subseteq \mathcal{N}$ are models of T so that for any inf-formula*

$$\inf_{s_1,y_1} \dots \inf_{s_n,y_n} \varphi(x_1,...,x_k,y_1,...,y_n)$$

and any $a_1,...,a_k \in M$

$$\inf_{s_1,y_1} \dots \inf_{s_n,y_n} \varphi(a_1,...,a_k,y_1,...,y_n)^{\mathcal{N}} = \inf_{s_1,y_1} \dots \inf_{s_n,y_n} \varphi(a_1,...,a_k,y_1,...,y_n)^{\mathcal{M}}$$

(that is, $\mathcal{M}$ is existentially closed in $\mathcal{N}$), then there exists an elementary extension $\mathcal{M}'$ of $\mathcal{M}$ and an embedding $f\colon \mathcal{N} \to \mathcal{M}'$ so that $f(a) = a$ for all $a \in \mathcal{M}$.

Proof. Let $\mathcal{M} \subseteq \mathcal{N}$ such that $\mathcal{M}$ and $\mathcal{N}$ are models of T. Assume, for any inf-formula

$$\inf_{s_1,y_1} \dots \inf_{s_n,y_n} \varphi(x_1,...,x_k,y_1,...,y_n)$$

and any $a_1,...,a_k \in M$

$$\inf_{s_1,y_1} \dots \inf_{s_n,y_n} \varphi(a_1,...,a_k,y_1,...,y_n)^{\mathcal{N}} = \inf_{s_1,y_1} \dots \inf_{s_n,y_n} \varphi(a_1,...,a_k,y_1,...,y_n)^{\mathcal{M}}.$$

Let κ be a cardinal larger than the cardinality of the underlying set N of $\mathcal{N}$. Let $\mathcal{M}'$ be a κ-saturated elementary extension of $\mathcal{M}$. Then, $(\mathcal{M}', a)_{a\in M} \equiv (\mathcal{M}, a)_{a\in M}$, and $(\mathcal{M}', a)_{a\in M}$ is κ-saturated (since $\kappa > \text{card}(N) \geq \text{card}(M)$). Moreover, any closed condition

$$\inf_{s_1,y_1} \dots \inf_{s_n,y_n} \varphi(a_1,...,a_k,y_1,...,y_n) = 0$$

(where $a_1,...,a_k$ are in appropriate sorts of $\mathcal{M}$ and φ is quantifier-free) that is true in $\mathcal{N}$ must also be true in $\mathcal{M}$, and hence in $\mathcal{M}'$. This allows us to build an embedding of $\mathcal{N}$ in $\mathcal{M}'$ over $\mathcal{M}$. $\square$

1.7.9 Proposition. *The L-theory T is model complete if and only if every model of T is an existentially closed model of T.*

Proof. The left to right direction is immediate from the definitions. Now, assume that every model $\mathcal{M}$ of T is an existentially closed model of T. Let $\mathcal{M}_1 \subseteq \mathcal{N}_1$ be models of T. We will show that $\mathcal{M}_1$ is an elementary substructure of $\mathcal{N}_1$, from which it follows that T is model complete. By Lemma 1.7.8, there exists an elementary extension $\mathcal{M}_2$ of $\mathcal{M}_1$ and an embedding $f\colon \mathcal{N}_1 \to \mathcal{M}_2$ so that $f(a) = a$ for all a in the underlying metric space M_1 of $\mathcal{M}_1$. By renaming the elements of the image of $\mathcal{N}_1$, we may assume $\mathcal{M}_1 \subseteq \mathcal{N}_1 \subseteq \mathcal{M}_2$. Since f is the identity on $\mathcal{M}_1$, this renaming does not change the fact that $\mathcal{M}_1 \preceq \mathcal{M}_2$.

Next, again by Lemma 1.7.8, there exists an elementary extension $\mathcal{N}_2$ of $\mathcal{N}_1$ and an embedding $f\colon \mathcal{M}_2 \to \mathcal{N}_2$ so that $f(a) = a$ for all a in the underlying metric space N_1 of $\mathcal{N}_1$. By renaming as above we may assume $\mathcal{N}_1 \subseteq \mathcal{M}_2 \subseteq \mathcal{N}_2$. Since f is the identity on $\mathcal{N}_1$, this renaming does not change the fact that $\mathcal{N}_1 \preceq \mathcal{N}_2$.

We proceed in this fashion to get an increasing chain of models $\mathcal{M}_1 \subseteq \mathcal{N}_1 \subseteq \mathcal{M}_2 \subseteq \mathcal{N}_2 \subseteq ...$ so that $\mathcal{M}_i \preceq \mathcal{M}_{i+1}$ and $\mathcal{N}_i \preceq \mathcal{N}_{i+1}$ for all $i \in \mathbb{N}$. Let $\mathcal{W}$ be the union of the elementary chain $(\mathcal{M}_n \mid n \geq 1)$, which is also the union of the elementary chain $(\mathcal{N}_n \mid n \geq 1)$. By Proposition 1.7.5 we have $\mathcal{M}_1 \preceq \mathcal{W}$ and $\mathcal{N}_1 \preceq \mathcal{W}$. Since $\mathcal{M}_1 \subseteq \mathcal{N}_1$, this implies $\mathcal{M}_1 \preceq \mathcal{N}_1$, as claimed. □

1.7.10 Definition. A *model companion* of T is an L-theory S such that:

- every model of T extends to a model of S;
- every model of S extends to a model of T;
- S is model complete.

1.7.11 Proposition. *If the L-theories S and S' are both model companions of T, then S is equivalent to S', that is,* $\mathrm{Mod}(S) = \mathrm{Mod}(S')$.

Proof. Assume S and S' are both model companions of T. It suffices to show $\mathrm{Mod}(S) \subseteq \mathrm{Mod}(S')$, since then we could switch S and S' in that argument to get $\mathrm{Mod}(S') \subseteq \mathrm{Mod}(S)$. Let $\mathcal{M} \models S$. Then, there exists an extension $\mathcal{W}$ of $\mathcal{M}$ so that $\mathcal{W} \models T$, and there exists an extension $\mathcal{N}$ of $\mathcal{W}$ so that $\mathcal{N} \models S'$. Therefore, there exists an extension $\mathcal{N}$ of $\mathcal{M}$ that is a model of S'. Similarly, for any $\mathcal{N} \models S'$ there exists an extension $\mathcal{M}$ of $\mathcal{N}$ so that $\mathcal{M} \models S$.

We use these facts to build an increasing chain $\mathcal{M}_1 \subseteq \mathcal{N}_1 \subseteq M_2 \subseteq \mathcal{N}_2 \subseteq ...$ so that each $\mathcal{M}_i$ is a model of S and each $\mathcal{N}_i$ is a model of S'. Since both S and S' are model complete, we know $\mathcal{M}_i \preceq \mathcal{M}_{i+1}$ for all $i \in \mathbb{N}^{>0}$ and $\mathcal{N}_i \preceq \mathcal{N}_{i+1}$ for all $i \in \mathbb{N}^{>0}$. Let $\mathcal{W}$ be the union of the elementary chain $(\mathcal{M}_n \mid n \geq 1)$, which is also the union of the elementary chain $(\mathcal{N}_n \mid n \geq 1)$.. By Proposition 1.7.5 we have $\mathcal{M}_1 \preceq \mathcal{W}$ and $\mathcal{N}_1 \preceq \mathcal{W}$. Since $\mathcal{M}_1 \subseteq \mathcal{N}_1$, this implies $\mathcal{M}_1 \equiv \mathcal{N}_1$. Therefore any model of S is elementarily equivalent to a model of S', so $\mathrm{Mod}(S) \subseteq \mathrm{Mod}(S')$. □

1.7.12 Definition. We say the L-theory T has *amalgamation over substructures* if: for any substructures $\mathcal{M}_0$, $\mathcal{M}_1$ and $\mathcal{M}_2$ of models of T and embeddings $f_1\colon \mathcal{M}_0 \to \mathcal{M}_1$, $f_2\colon \mathcal{M}_0 \to \mathcal{M}_2$, there exists a model $\mathcal{N}$ of T and embeddings $g_i\colon \mathcal{M}_i \to \mathcal{N}$ such that $g_1 \circ f_1 = g_2 \circ f_2$.

1.7.13 Proposition. *Let S and T be L-theories such that S is the model companion of T. Assume T has amalgamation over substructures. Then S has quantifier elimination.*

Proof. Let L, S and T be as above. A theory S has quantifier elimination if and only if the following holds: if $\mathcal{M}, \mathcal{N} \models S$, then every embedding of a substructure of $\mathcal{M}$ into $\mathcal{N}$ can be extended to an embedding of $\mathcal{M}$ into an elementary extension of $\mathcal{N}$. (See [1] Proposition 13.6.) It is straightforward to show this property holds for S. Let $\mathcal{M}$ and $\mathcal{N}$ be models of S. Let $\mathcal{A} \subseteq \mathcal{M}$ and let $f\colon \mathcal{A} \to \mathcal{N}$ be an embedding. Let $\mathcal{W}$ be an amalgam of $\mathcal{M}$ and $\mathcal{N}$ over $\mathcal{A}$ as models of T. Without loss of generality, we may assume $\mathcal{N} \subseteq \mathcal{W}$. Let $\mathcal{W}' \models S$ be an extension of $\mathcal{W}$. Since S is model complete we know $\mathcal{W}'$ is an elementary extension of $\mathcal{N}$. And, clearly $\mathcal{M}$ is embedded in $\mathcal{W}'$ by a map that extends f. □

Definability is one of the central notions in model theory. The following definition is adapted from [1, Chapter 9]. Let I_P be a bounded interval in $\mathbb{R}^{\geq 0}$. A predicate $P\colon M^{(s_1)} \times ... \times M^{(s_n)} \to I_P$ is *definable in* $\mathcal{M}$ *over* A if an only if there is a sequence $(\varphi_k \mid k \geq 1)$ of $L(A)$-formulas such that the predicates $\varphi_k^{\mathcal{M}}$ converge to P uniformly on $M^{(s_1)} \times ... \times M^{(s_n)}$. A closed set $D \subseteq M^{(s_1)} \times ... \times M^{(s_n)}$ is *definable in* $\mathcal{M}$ *over* A if and only if the predicate $\mathrm{dist}(x, D)$ is definable in $\mathcal{M}$ over A. A function $f\colon M^{(s_1)} \times ... \times M^{(s_n)} \to M^{(s_0)}$ is *definable in* $\mathcal{M}$ *over* A if and only the predicate $d^{(s_0)}(f(x_1, ..., x_n), y)$ is definable in $\mathcal{M}$ over A. We will use the term 0-*definable* when the set of parameters A is empty.

1.7.14 Definition. Let T be an L-theory and $\varphi(x_1, ..., x_n)$ an L-formula. By *the zerosets of the formula* $\varphi(x_1, ..., x_n)$, we mean the sets

$$\{(x_1, ..., x_n) \in M^{(s_1)} \times ... \times M^{(s_n)} \mid \varphi^{\mathcal{M}}(x_1, ..., x_n) = 0\}$$

as $\mathcal{M}$ varies over the models of T.

The zerosets of $\varphi(x_1, ..., x_n)$ are said to be *uniformly definable in models of* T if and only if there exists a definable predicate $\psi(x_1, ..., x_n)$ such that for all models $\mathcal{M}$ of T and all $a \in M$ of appropriate sort

$$\psi^{\mathcal{M}}(a_1, ..., a_n) = \mathrm{dist}((a_1, ..., a_n), \{(b_1, ..., b_n) \in M^n \mid \varphi^{\mathcal{M}}(b_1, ..., b_n) = 0\}).$$

That is, for each model $\mathcal{M}$ of T the distance to the zeroset of $\varphi(x_1, ..., x_n)^{\mathcal{M}}$ is given by the interpretation of the *same* definable predicate ψ. In the rest of this thesis, if we say a certain family of sets is uniformly definable in all models of a theory T, we mean that there exists an L-formula $\varphi(x_1, ..., x_n)$ so that the given family of sets is exactly the family of zerosets of $\varphi(x_1, ..., x_n)$, and that they are uniformly definable as defined above.

The utility of definable sets is demonstrated by the following theorem.

1.7.15 Theorem. *([1, Theorem 9.17]) For a closed set $D \subseteq M^{(s)}$, the following are equivalent:*

1. *D is definable in $\mathcal{M}$ over A.*
2. *For any predicate*

$$P\colon (M^{(s_1)} \times ... \times M^{(s_n)}) \times M^{(s)} \to I_P$$

that is definable in $\mathcal{M}$ over A, the predicate $Q\colon M^{(s_1)} \times ... \times M^{(s_n)} \to I_Q$ defined by

$$Q(x) = \inf\{P(x_1, ..., x_n, y) \mid y \in D\}$$

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

It is straightforward to extend this theorem to closed sets $D \subseteq M^{(s_1)} \times ... \times M^{(s_n)}$.

As in first order logic, in continuous logic there is a notion of an *extension by definition* whereby we extend a signature and its structures by adding a symbol for a definable predicate or function without really changing the expressive power of that signature. See [1, Chapter 9] for more about definability and extension by definition.

1.8 L_p-structures

This section presents the method used in this thesis for studying certain kinds of unbounded metric spaces using many-sorted bounded continuous logic. There is an alternative approach in [3], which is equivalent to what we do here for the structures considered in this thesis.

Recall from Section 1.2 that the continuous signature L_p consists of:

- a non-empty sort index set $S = \mathbb{N}$ and the family of positive real numbers $(2n \mid n \in \mathbb{N})$;
- for each $m, n \in \mathbb{N}$ with $m \leq n$, a function symbol $I_{m,n}$ with arity $(m; n)$ and modulus of uniform continuity $\Delta_{m,n}(\epsilon) = \epsilon$;

- constant symbols $(p_n \mid n \in \mathbb{N})$, where p_n has arity n;
- metric symbols $d^{(n)}$ for each $n \in \mathbb{N}$.

In addition, recall that given a pointed metric space (M, d, p), we define the *corresponding L_p-structure*

$$\mathcal{M} = \Big(((M^{(n)}, d^{(n)}) | n \in \mathbb{N}), p_n^{\mathcal{M}}, I_{m,n}^{\mathcal{M}} \Big)$$

so that:

- the n^{th} sort $(M^{(n)}, d^{(n)})$ is the closed ball of radius n centered at p in (M, d) with the metric restricted to that ball;
- for each $m, n \in \mathbb{N}$ with $m \leq n$, the function $I_{m,n} \colon M^{(m)} \to M^{(n)}$ is the inclusion map;
- the constants p_n are all interpreted as the point p.

Thus, given a pointed metric space, we know how to construct the (unique) corresponding L_p-structure. Next, we define an L_p-theory T_p so that for every $\mathcal{M} \models T_p$, there is a unique pointed metric space (M, d, p) such that $\mathcal{M}$ is isomorphic to the L_p structure corresponding to (M, d, p). First, we need some more definitions and lemmas.

A metric space M is a *length space* if, for any points $x, y \in M$, the distance between x and y is equal to the infimum of the lengths of the rectifiable paths between x and y.

1.8.1 Definition. ([5, Chapter 3, Section 1]) We say a metric space M has the *approximate midpoint property* if: for any $x, y \in M$, for any $\epsilon > 0$, there exists $z \in M$ such that

$$|d(x,z) - \frac{d(x,y)}{2}| \leq \epsilon \text{ and } |d(y,z) - \frac{d(x,y)}{2}| \leq \epsilon.$$

1.8.2 Fact. ([5, Chapter 3, Section 1]) A complete metric space M is a length space if and only if it has the approximate midpoint property.

1.8.3 Definition. ([5, Chapter I, 1.3]) A metric space (X, d) is *geodesic* if for any $x, y \in X$ there exists an isometric embedding $\gamma : [0, d(x,y)] \to X$ with $\gamma(0) = x$ and $\gamma(d(x,y)) = y$. A *geodesic segment* is the image of such a path.

1.8.4 Fact. ([5, Chapter I, 1.4]) A complete metric space M is a geodesic space if and only if for any two points $x, y \in M$ there exists a midpoint z between x and y. That is, there exists z such that:

$$d(x,z) = \frac{d(x,y)}{2}, \text{ and } d(y,z) = \frac{d(x,y)}{2}.$$

1.8.5 Lemma. *If a complete metric space M is a length space, then for any $r \in \mathbb{R}$ and any $a \in M$, the open ball $\mathcal{O}_r(a)$ with center a and radius r is dense in the closed ball $B_r(a)$.*

Proof. Let M be a complete length space and let $r \in \mathbb{R}$ and $a \in M$. Towards contradiction, assume that $\mathcal{O}_r(a)$ is not dense in $B_r(a)$. Then there exists $y \in B_r(a) \setminus \mathcal{O}_r(a)$ and $\epsilon > 0$ such that for all $x \in \mathcal{O}_r(a)$, we know $d(x,y) > \epsilon$. Because M is a length space and $d(a,y) = r$, we may find a path γ between a and y with length $l(\gamma) < r + \frac{\epsilon}{2}$. (Let γ refer to both the path and its image.) By the connectedness of γ, we may find $x \in \gamma$ such that $\frac{\epsilon}{2} \leq d(x,y) < \epsilon$. Then,

$$d(a,x) + d(x,y) \leq l(\gamma) < r + \frac{\epsilon}{2}$$

and because $\frac{\epsilon}{2} \leq d(x,y)$, we conclude $d(a,x) < r$. So, $x \in \mathcal{O}_r(a)$ and $d(x,y) < \epsilon$. This is a contradiction. □

The theory T_p defined below gives axioms for a class of L_p-structures for which we may identify a unique underlying pointed metric space.

1.8.6 Definition. Let T_p be the following L_p-theory:

1. for all $m, n \in \mathbb{N}$ with $m \leq n$ the axiom

$$d^{(n)}(I_{m,n}(p_m), p_n) = 0;$$

2. for all $m, n \in \mathbb{N}$ with $m \leq n$ the axiom

$$\sup_{m,x} \sup_{m,y} |d^{(n)}(I_{m,n}(x), I_{m,n}(x)) - d^{(m)}(x,y)| = 0;$$

3. for all $n \in \mathbb{N}$ the axiom

$$\sup_{n,x} d^{(n)}(x, I_{n,n}(x)) = 0;$$

4. for all $m, n \in \mathbb{N}$ with $m < n$ the axiom

$$\sup_{n,y} \min\{m \dot{-} d^{(n)}(y, p_n), \quad \inf_{m,x} d^{(n)}(I_{m,n}(x), y)\} = 0;$$

5. for all $j, m, n \in \mathbb{N}$ with $j \leq m \leq n$, the axiom

$$\sup_{j,x} d^{(n)}(I_{m,n} \circ I_{j,m}(x), I_{j,n}(x)) = 0;$$

6. for all $n \in \mathbb{N}$ the axiom

$$\sup_{n,x} \sup_{n,y} \inf_{n,z} \max\{|d(x,z) - \frac{d(x,y)}{2}|, \quad |d(y,z) - \frac{d(x,y)}{2}|\} = 0.$$

1.8.7 Lemma. *If (M, d, p) is a pointed metric space and $\mathcal{M}$ is the corresponding L_p-structure, then $\mathcal{M}$ is a model of the axioms in (1)-(5) from the definition of T_p.*

Proof. This is clear from the definition of the L_p-structure corresponding to (M, d, p). □

1.8.8 Lemma. *Let (M, d, p) be a pointed metric space such that each ball closed ball $B_n(p)$ is a length space. Let $\mathcal{M}$ be the corresponding L_p-structure. Then, $\mathcal{M} \models T_p$.*

Proof. This follows from Lemma 1.8.7 and the fact that if each $B_n(p) = M^{(n)}$ is a length space, then the axioms in (6) from the definition of T_p are true in $\mathcal{M}$. □

1.8.9 Lemma. *Let $\mathcal{M} = \Big(((M^{(n)}, d^{(n)}) \mid n \in \mathbb{N}), p_n^{\mathcal{M}}, I_{m,n}^{\mathcal{M}}\Big)$ be a model of T_p. Then, $\big((M^{(n)}, d^{(n)}, p_n^{\mathcal{M}}), I_{m,n}^{\mathcal{M}}\big)$ is a directed system of pointed metric spaces indexed by $\mathbb{N}$. If (W, d_W, q) is the direct limit of this directed system, then the corresponding L_p-structure*

$$\mathcal{W} = \Big(((W^{(n)}, d_W^{(n)}) \mid n \in \mathbb{N}), p_n^{\mathcal{W}}, I_{m,n}^{\mathcal{W}}\Big)$$

is isomorphic to $\mathcal{M}$.

Proof. Assume $\mathcal{M}$ is as above. By the axioms in (1) and (2) in Definition 1.8.6, we know that each function $I_{m,n}^{\mathcal{M}} \colon M^{(m)} \to M^{(n)}$ is an isometric embedding which takes $p_m^{\mathcal{M}}$ to $p_n^{\mathcal{M}}$. The axioms in (3) imply that $I_{n,n}$ is the identity on $(M^{(n)}, d^{(n)})$. The axioms in (4) imply that for each $m \leq n$, for every $y \in M^{(n)}$ if $d(y, p_n^{\mathcal{M}}) < m$, then $\inf_{m,x} d^{(n)}(I_{m,n}^{\mathcal{M}}(x), y) = 0$. Therefore, the closure of the image of $I_{m,n}$ contains the open ball centered at $p_n^{\mathcal{M}}$ of radius m in $M^{(n)}$. By the axioms in (6), each sort $(M^{(n)}, d^{(n)})$ is a length space, and therefore the open ball centered at $p_n^{\mathcal{M}}$ of radius m in $M^{(n)}$ is dense in the closed ball centered at $p_n^{\mathcal{M}}$ of radius m in $M^{(n)}$. By uniform continuity of $I_{m,n}^{\mathcal{M}}$, and the completeness of each sort we conclude that $I_{m,n}^{\mathcal{M}}$ maps $M^{(m)}$ isometrically onto the closed ball of radius m and center p contained in $M^{(n)}$. The axioms in

(5) imply that $I^{\mathcal{M}}_{j,m} \circ I^{\mathcal{M}}_{m,n} = I^{\mathcal{M}}_{j,n}$ for $j \leq m \leq n$. Therefore, $\big((M^{(n)}, d^{(n)}, p^{\mathcal{M}}_n), I^{\mathcal{M}}_{m,n}\big)$ is a directed system.

Suppose $\mathcal{W}$ is as in the statement of the lemma. Then, (W, d_W, q) is the direct limit of $\big((M^{(n)}, d^{(n)}, p^{\mathcal{M}}_n), I^{\mathcal{M}}_{m,n}\big)$, and we have isometric embeddings $g_n : M^{(n)} \to W$ for $n \in \mathbb{N}$ such that $g_n(p^{\mathcal{M}}_n) = q$, and for all $m \leq n$ we know $g_n = g_m \circ I^{\mathcal{M}}_{m,n}$. Note that g_n maps onto $W^{(n)}$. The family of functions $(g_n \mid n \in \mathbb{N})$ is an isomorphism between the L_p-structures $\mathcal{M}$ and $\mathcal{W}$. □

1.8.10 Corollary. *Assume (M, d, p) is a pointed metric space such that the corresponding L_p-structure $\mathcal{M}$ is a model of T_p. Then $\big((M^{(n)}, d^{(n)}, p^{\mathcal{M}}_n), I^{\mathcal{M}}_{m,n}\big)$ is a directed set of pointed metric spaces indexed by $\mathbb{N}$. If (W, d_W, q) is the direct limit of this directed set, then (M, d, p) and (W, d_W, q) are isomorphic as pointed metric spaces. That is, there exists an isometry from (M, d) onto (W, d_W) which sends p to q.*

Proof. Take the union of the family of functions $(g_n \mid n \in \mathbb{N})$ from the proof above. □

So, for a model $\mathcal{M}$ of T_p, there is a unique underlying pointed metric space (up to isomorphism), equal to the direct limit of the sorts of $\mathcal{M}$. For any $\mathcal{M} \models T_p$ we will call this space the *underlying metric space of* $\mathcal{M}$. From now on, all of the pointed metric spaces considered in this thesis will be such that their corresponding L_p-structures are models of T_p. In this many-sorted setting, the proliferation of sorts and the sort by sort way we defined notions such as embedding, subset and isomorphism often gets in the way of the clarity of statements and arguments. Here are some conventions intended to help avoid some of the notational and expositional complications of the many-sorted setting.

Let (M, d, p) be a pointed metric space and $\mathcal{M} \models T_p$ its corresponding L_p-structure.

- For legibility, the inclusion map symbols $I_{m,n}$ will be left out of formulas.
- Because the metrics $d^{(n)}$ on the sorts are all just restrictions of the metric on M, we will leave off the superscripts.
- We will often use p instead of p_n, and often use p instead of $p^{\mathcal{M}}$.
- We will often refer to the underlying metric space of $\mathcal{M}$ by M instead of referring to the sorts.
- Let $(A^{(n)} \mid n \in \mathbb{N}) \subseteq (M^{(n)} \mid n \in \mathbb{N})$ be such that $A^{(m)} = A^{(n)} \cap M^{(m)}$ for all $m < n$. Then, there exists a unique subset $A \subseteq M$ such that $A^{(n)} = A \cap M^{(n)}$ for each $n \in \mathbb{N}$.

In the other direction, clearly any subset $A \subseteq M$ gives a unique family $(A \cap M^{(n)} \mid n \in \mathbb{N}) \subseteq (M^{(n)} \mid n \in \mathbb{N})$. This implies that the underlying metric space of any substructure of $\mathcal{M} \models T_p$ is uniquely determined up to isomorphism. For models of T_p, we will often simply refer to a subset A of the underlying space instead of referring to the images of A in each sort.

- Say $\mathcal{M} \models T_p$ and $\mathcal{N} \models T_p$ are two L_p-structures with underlying metric spaces (M, d_M, p_M) and (N, d_N, p_N) respectively. Let $f\colon M \to N$ be an isometric embedding such that $f(p_M) = p_N$. Then the collection of maps

$$(f^{(n)}\colon M^{(n)} \to N^{(n)} \mid n \in \mathbb{N})$$

where $f^{(n)}$ is the restriction of f to $M^{(n)}$, is an embedding of L_p-structures. If $f\colon M \to N$ is onto, then $(f^{(n)}\colon M^{(n)} \to N^{(n)} \mid n \in \mathbb{N})$ is an isomorphism of L_p-structures. Moreover, every embedding or isomorphism between L_p structures is given by an embedding on the underlying metric spaces in this manner. When it is not confusing to do so, we will use f as an abbreviation for $(f^{(n)}\colon M^{(n)} \to N^{(n)} \mid n \in \mathbb{N})$.

- If the underlying metric space M of $\mathcal{M}$ has density character κ, then we say $\mathcal{M}$ has density character κ. The inclusion maps between the sorts guarantee that the density characters of the sorts are increasing. The density of the underlying space M is equal to the supremum of the densities of the sorts.

- The following convention will be useful in Chapters 3 and 4 when we want to add functions to our L_p-structures. Let $k, s \in \mathbb{N}$. Let L_s be an extension of L_p by a family of function symbols $\{f_n \mid n \in \mathbb{N}\}$ where f_n has arity $(n; n+s)$. Let T' be an extension of the L_p-theory T_p by a family of axioms

$$\{\sup_{m,x} d(f_n(I_{m,n}(x)), I_{m+s,n+s}(f_m(x))) = 0 \mid m, n \in \mathbb{N}; \quad m \leq n\}.$$

Let $\mathcal{M} \models T'$ be an L_s-structure. Then, when we take the direct limit of the sorts to find the underlying metric space for $\mathcal{M}$, we may also take the direct limit (union) of each of the functions $\{f_n^{\mathcal{M}} \mid n \in \mathbb{N}\}$ and get a function $f^{\mathcal{M}}$ on the underlying metric space such that $f_n^{\mathcal{M}}$ is the restriction of f to $B_n(p) = M^{(n)}$. We will call this the *underlying function of* $\mathcal{M}$. As with the metric, in L_s-formulas we may use the symbol f instead of f_n.

Next, the *ultraproduct* for pointed (possibly unbounded) metric spaces is defined. Say $((M_i, d_i, p_i) \mid i \in I)$ is a family of pointed metric spaces, and let U be an ultrafilter on I. We define $\prod_U (M_i, d_i, p_i) = (N, d, p)$, the U-ultraproduct of $((M_i, d_i, p_i) \mid i \in I)$ as follows. Let

$$\widehat{N} = \{(x_i)_{i \in I} \mid x_i \in M_i, \text{ and there exists } c \in \mathbb{R},\ \forall i \in I\ d_i(x_i, p_i) \leq c\}.$$

Define a pseudometric $\hat{d}$ on $\widehat{N}$ by $\hat{d}((a_i), (b_i)) = \lim_{U,i} d_i(a_i, b_i)$. Let (N, d, p) be the result of taking the quotient of $(\widehat{N}, \hat{d})$ where p is the image of (p_i) in this quotient. In the case that the pointed metric spaces in question are uniformly bounded, this ultraproduct is isometric to the ultraproduct of bounded metric spaces as defined in Section 1.5, and the basepoint and the condition that sequences in $\widehat{N}$ be bounded are superfluous.

1.8.11 Lemma. *Say $((M_i, d_i, p_i) \mid i \in I)$ is a family of pointed metric spaces such that the corresponding L_p-structures $\mathcal{M}_i$ are all models of T_p. Let U be an ultrafilter. Then the L_p-structure $\mathcal{N}$ corresponding to the ultraproduct $\prod_U (M_i, d_i, p_i)$ is isomorphic to the L_p-structure $\prod_U \mathcal{M}_i$.*

Proof. Let $((M_i, d_i, p_i) \mid i \in I)$ and $\mathcal{M}_i$ be as above. By the Fundamental Theorem of Ultraproducts (Theorem 1.5.1) we know that $\prod_U \mathcal{M}_i$ is a model of T_p. To make the notation nicer, let $\prod_U \mathcal{M}_i = \mathcal{M}$. By the definition of the ultraproduct of many-sorted metric structures (see Section 1.5),

$$\mathcal{M} = \prod_U \mathcal{M}_i = \Big(((\prod_U M_i^{(n)}, \delta^{(n)}) \mid n \in \mathbb{N}), p_n^{\mathcal{M}}, I_{m,n}^{\mathcal{M}} \Big)$$

where the metric $\delta^{(n)} = \lim_{i,U} d_i^{(n)}$, the basepoint $p^{\mathcal{M}} = [p^{\mathcal{M}_i}]_U$ and $I_{m,n}^{\mathcal{M}}([a_i]_U) = [I_{m,n}^{\mathcal{M}_i}(a_i)]_U$. Let (N, d, p) be the ultraproduct $\prod_U (M_i, d_i, p_i)$ of pointed metric spaces, and denote elements of N by $[a_i]_U^N$. It suffices to show that for each $n \in \mathbb{N}$, there is an isometry from $\prod_U M_i^{(n)}$ onto $B_n^N(p)$ which sends $p_n^{\mathcal{M}}$ to p. If such isometries exist, then they form an isomorphism between the L_p-structure $\prod_U \mathcal{M}_i = \mathcal{M}$ and the L_p-structure $\mathcal{N}$ corresponding to the pointed metric space $\prod_U (M_i, d_i, p_i) = (N, d, p)$. Note that for all $[a_i]_U \in \prod_U M_i^{(n)}$, the representative $(a_i)_{i \in I}$ must be such that $a_i \in M_i$, and $d_i(a_i, p_i) \leq n$ for all $i \in I$. Thus,

$$(a_i)_{i \in I} \in \widehat{N} = \{(x_i)_{i \in I} \mid x_i \in M_i, \text{ and there exists } c \in \mathbb{R}, d_i(x_i, p_i) \leq c\ \forall i \in I\}.$$

which means that $[a_i]_U^N$ is an element of the pointed ultraproduct $(N, d, p) = \prod_U (M_i, d_i, p_i)$. Also, since $d_i(a_i, p_i) \leq n$ for all $i \in I$, clearly $d([a_i]_U^N, [b_i]_U^N) = \lim_{U,i} d_i(a_i, b_i) \leq n$. Thus,

$[a_i]_U^N \in B_n^N(p)$. Therefore, we may define $f\colon \prod_U M_i^{(n)} \to B_n^N(p)$ by $f([a_i]_U) = [a_i]_U^N$.

For all $[a_i]_U, [b_i]_U \in \prod_U M_i^{(n)}$,

$$\delta^{(n)}([a_i]_U, [b_i]_U) = \lim_{U,i} d_i(a_i, b_i)$$

by the definition of the ultraproduct of bounded metric structures. We also know

$$\lim_{U,i} d_i(a_i, b_i) = d([a_i]_U^N, [b_i]_U^N)$$

by the definition of the ultraproduct for pointed metric spaces. Therefore

$$\delta^{(n)}([a_i]_U, [b_i]_U) = d([a_i]_U^N, [b_i]_U^N) = d(f([a_i]_U), f([b_i]_U))$$

which proves f is well defined, since if $\delta^{(n)}(f([a_i]_U^N), f([b_i]_U^N)) = 0$, then $d([a_i]_U^N, [b_i]_U^N) = 0$, and also shows f is an isometric embedding.

It remains to show that f is onto. It suffices to show that for any $[b_i]_U^N \in B_n^N(p)$, there is a representative $(a_i)_{i\in I}$ of $[b_i]_U^N$ so that for every $i \in I$, we know $a_i \in M_i^{(n)}$. In that case, we know $[a_i]_U \in \prod_U M_i^{(n)}$ and $f([a_i]_U) = [a_i]_U^N = [b_i]_U^N$.

Let $[b_i]_U^N \in B_n^N(p)$. If $\{i \mid d_i(b_i, p_i) \le n\} \in U$, then we are done. So, assume $\{i \mid d_i(b_i, p_i) > n\} \in U$. It follows from this assumption that $d(p, [b_i]_U^N) = n$, and we may assume without loss of generality that $2n > d_i(b_i, p_i) > n$ for all $i \in I$.

For each $i \in I$, let γ_i be a rectifiable path from p_i to b_i contained in $M_i^{(2n)}$ such that

$$d_i^{(2n)}(p_i, b_i) + \frac{1}{2^i} \ge length(\gamma_i) \ge d_i^{(2n)}(p_i, b_i).$$

These paths exist because for each $i \in I$, the L_p-structure $\mathcal{M}_i$ is a model of T_p, and therefore each $M_i^{(2n)}$ is a length space. (Note that $\mathcal{M}_i \models T_p$ also implies $d_i^{(2n)}$ is a restriction of the metric d_i on $\mathcal{M}_i$ to the closed $2n$-ball $M_i^{(2n)} \subseteq M_i$.) The image of each γ_i is connected, so for each $i \in I$ we may find a point a_i on γ_i such that $d_i^{(2n)}(a_i, p_i) = n$. Therefore $a_i \in M_i^{(n)}$.

Let γ_i^1 denote γ_i restricted to $[0, \gamma_i^{-1}(a_i)]$, so γ_i^1 is a rectifiable path from p_i to a_i. Let γ_i^2 denote γ_i restricted to $[\gamma_i^{-1}(a_i), \gamma_i^{-1}(b_i)]$, so γ_i^2 is a rectifiable path from a_i to b_i. Then,

$$d_i^{(2n)}(a_i, b_i) \le length(\gamma_i^2) = length(\gamma_i) - length(\gamma_i^1) \le d_i^{(2n)}(p_i, b_i) + \frac{1}{2^i} - d_i^{(2n)}(p_i, a_i).$$

Therefore,

$$\lim_{U,i} d_i^{(2n)}(a_i, b_i) \leq \lim_{U,i} d_i^{(2n)}(p_i, b_i) + \lim_{U,i} \frac{1}{2^i} - \lim_{U,i} d_i^{(2n)}(p_i, a_i) = n + 0 - n = 0.$$

So,

$$d([a_i]_U^N, [b_i]_U^N) = \lim_{U,i} d_i(a_i, b_i) = \lim_{U,i} d_i^{(2n)}(a_i, b_i) = 0$$

and for each $i \in I$ the distance $d_i(a_i, p_i) = d_i^{(2n)}(a_i, p_i) \leq n$, which implies that $a_i \in M_i^{(n)}$.

□

We now discuss a notion of definability for subsets of and functions on the underlying metric space M of a model $\mathcal{M}$ of T_p. A closed set D that is a subset of $M^{(n_1)} \times ... \times M^{(n_k)}$ can be viewed as a subset of $(M^{(m)})^k$ for any $m \geq \max\{n_1, ..., n_k\}$, and D also corresponds to a unique closed subset of the underlying metric space M^k.

1.8.12 Definition. Let $\mathcal{M} \models T_p$ with underlying metric space (M, d, p). Let $A \subseteq M$ be a subset and let $D \subseteq M^k$ be a closed subset. We say D is *definable in* $\mathcal{M}$ *over* A if the intersection $D \cap (M^{(n)})^k$ is definable in $\mathcal{M}$ over A for each $n \in \mathbb{N}$.

If for each $n \in \mathbb{N}$ the sets $D \cap (M^{(n)})^k$ are uniformly definable in models $\mathcal{M}$ of T_p, then we say D is *uniformly definable* in models of T_p.

1.8.13 Definition. Let $\mathcal{M} \models T_p$ with underlying metric space (M, d, p). Let $A \subseteq M$ be a subset and $f \colon M^k \to M$ a uniformly continuous function. We say f is *definable in* $\mathcal{M}$ *over* A if the restriction $f|(M^{(n)})^k$ is definable in $\mathcal{M}$ over A for each $n \in \mathbb{N}$.

If the functions $f|(M^{(n)})^k$ are uniformly definable in models of T_p, then we say f is *uniformly definable* in models of T_p.

In formulating the signature L_p we made an arbitrary choice to use closed balls of integer radius as the sorts of our many-sorted structures. It is natural to ask if this choice has any effect on the model theoretic properties of the structures we study (especially, of $\mathbb{R}$-trees). In fact it has no such effect, as the following discussion (in conjunction with Theorem 1.7.15 above) indicates.

We show in what follows that closed balls centered at the basepoint are uniformly definable in models of T_p via quantifier free formulas. We will need the following useful connective.

1.8.14 Definition. ([1, Definition 6.1]) Define the binary function $\dot{-} \colon [0, \infty) \times [0, \infty) \to [0, \infty)$ by $x \dot{-} y = \max\{x - y, 0\}$.

1.8.15 Lemma. *Let $r \in \mathbb{R}^{>0}$ and $n \in \mathbb{N}$ with $r \leq n$. Let $\varphi_n(x)$ be the quantifier-free formula $d(x,p) \dot{-} r$ where x has sort n. Suppose $\mathcal{M} \models T_p$ and (M, d, p) is the underlying metric space of $\mathcal{M}$. If M is a geodesic space, then the closed ball $B_r(p) \subseteq M^{(n)}$ is 0-definable. Indeed, for $x \in M^{(n)}$ we have* $\operatorname{dist}(x, B_r(p)) = \varphi_n(x)^{\mathcal{M}}$.

Proof. Assume the situation in the hypotheses of the lemma. We show that $\varphi_n^{\mathcal{M}}(x)$ is equal to the distance function $\operatorname{dist}(x, B_r(p) \cap M^{(n)})$. For $x \in M^{(n)}$ we know $\varphi_n^{\mathcal{M}}(x) = 0$ if and only if $d(x,p) \leq r$, i.e. if and only if $x \in B_r(p)$. Now, let $x \notin B_r(p)$. Then, $\varphi_n^{\mathcal{M}}(x) = d(x,p) - r$. Let γ be a geodesic segment from p to x with $\gamma(0) = p$. Then,

$$\operatorname{dist}(x, B_r(p) \cap M^{(n)}) \leq d(x, \gamma(r)) = d(x,p) - d(p, \gamma(r)) = d(x,p) - r = \varphi_n^{\mathcal{M}}(x).$$

Now towards contradiction, assume $d(x,p) - r > \operatorname{dist}(x, B_r(p) \cap M^{(n)})$. Then, there exists a point in $c \in B_r(p)$ with $d(c,x) < d(x,p) - r$. So,

$$d(x,p) \leq d(p,c) + d(c,x) \leq r + d(c,x) < r + d(x,p) - r = d(x,p)$$

and thus $d(x,p) < d(x,p)$, which is a contradiction. Therefore, $\operatorname{dist}(x, B_r(p) \cap M^{(n)}) = \varphi_n(x)^{\mathcal{M}}$. $\square$

1.8.16 Lemma. *Assume $\mathcal{M} \models T_p$. Let $\mathcal{N}$ be an ω-saturated elementary extension of $\mathcal{M}$. Then the underlying metric space N of $\mathcal{N}$ is a geodesic space.*

Proof. Assume $\mathcal{M}$ and $\mathcal{N}$ are as specified above. For each $n \in \mathbb{N}$ let θ_n be the L_p-sentence

$$\sup_{n,x} \sup_{n,y} \inf_{n,z} \max\{|d(x,z) - \frac{d(x,y)}{2}|, |d(y,z) - \frac{d(x,y)}{2}|\}.$$

For each n we know that $\theta_n^{\mathcal{M}} = 0$ because $\mathcal{M} \models T_p$. Because $\mathcal{N}$ is an elementary extension of $\mathcal{M}$, we know $\theta_n^{\mathcal{N}} = 0$ is true. Since $\mathcal{N}$ is ω-saturated, Proposition 1.7.3 yields that in each of the θ_n the $\inf_{n,z}$ quantifier is realized exactly in $N^{(n)}$. Thus, between any x and y in $N^{(n)}$ there is a midpoint. Therefore between any x and y in N there is a midpoint. By Fact 1.8.4, N is geodesic. $\square$

1.8.17 Theorem. *For any $r \in \mathbb{R}^{\geq 0}$ and $n \in \mathbb{N}$ the family of sets $M^{(n)} \cap B_r^M(p)$ is uniformly definable in models of T_p.*

Proof. Let $\mathcal{M} \models T_p$ and take an ω-saturated elementary extension $\mathcal{N}$. Then the underlying metric space N of $\mathcal{N}$ is a geodesic space by Lemma 1.8.16. By Lemma 1.8.15, the intersection $N^{(n)} \cap B_r^N(p)$ is a definable subset in $\mathcal{N}$ over $\{p\}$. Moreover, the distance to $N^{(n)} \cap B_r^N(p)$ in N is given by interpreting the quantifier-free formula $\varphi_n(x) = d(x,p) \dot{-} r$ in $\mathcal{N}$, where the variable x is from the sort $N^{(n)}$. Then, since $\varphi_n(x)$ was quantifier free, the distance in M to $B_r^M(p)$ is given by the interpretation of the formula $\varphi_n(x) = d(x,p) \dot{-} r$ in $\mathcal{M}$ (where the variable x is from the sort $M^{(n)} \subseteq N^{(n)}$.) Since the distance to the ball is given by the same formula in every model of T_p, the balls are uniformly definable in models of T_p. $\square$

CHAPTER 2

MODEL THEORY OF $\mathbb{R}$-TREES

2.1 Introduction to $\mathbb{R}$-trees

This section contains a number of basic definitions and facts about $\mathbb{R}$-trees. Much of this material can be found in [5], [6] and [13].

First, a standard bit of notation. Given a metric space X, a subset Y of X and $\delta > 0$, we denote the set $\{x \in X \mid \mathrm{dist}(x, Y) \leq \delta\}$ by Y^δ.

An $\mathbb{R}$-*tree* is a metric space X such that between any two points in X there is a unique arc, and that arc is a geodesic segment. (An *arc* is the image of a topological embedding $f\colon [a, b] \to X$ of $[a, b] \subset \mathbb{R}$ where $f(a) = x$ and $f(b) = y$.) In an $\mathbb{R}$-tree, $[a, b]$ denotes the unique geodesic segment between a and b. Setwise, $[a, b] = [b, a]$, but when we write the segment as $[a, b]$, we mean that the corresponding isometric embedding $\gamma\colon [0, d(a, b)] \to X$ is such that the initial point $\gamma(0) = a$ and the terminal point $\gamma(d(a, b)) = b$. When it does not cause confusion, we may use $[a, b]$ to refer to the isometric embedding as well as its image. It is a fact (see [6, Lemma 2.4.14]) that the completion of an $\mathbb{R}$-tree is an $\mathbb{R}$-tree. In this thesis most of the $\mathbb{R}$-trees we consider are complete, since metric structures are required to be based on complete metric spaces.

Let M be an $\mathbb{R}$-tree and $a \in M$. Call the connected components of $M \setminus \{a\}$ *branches* at a. Let the *degree* of a point $a \in M$ be the cardinal number of branches at a. If there are three or more branches at $a \in M$, then we call a a *branch point*. The *height* of a branch β at a is $\sup\{d(a, x) | x \in \beta\}$ if that supremum exists, and is ∞ otherwise. A *subtree* of M is any subspace of M that is itself an $\mathbb{R}$-tree. Any connected subspace of M is a subtree. Moreover, subtrees are convex. That is, given any two points a, b in a subtree N, the geodesic segment $[a, b] \subseteq M$ is contained in N. Thus, the intersection of two subtrees is a subtree. A *ray* in an $\mathbb{R}$-tree is an isometric copy of $\mathbb{R}^{\geq 0}$. If $a \in M$ a *ray at a* is a ray so that the image of 0 under the isometric embedding of $\mathbb{R}^{\geq 0}$ into M is a.

2.1.1 Lemma. *(See [6, Lemma 1.9]) If M is an $\mathbb{R}$-tree and $\mathcal{T}_1$, $\mathcal{T}_2$ are disjoint, closed, non-empty subtrees of M, then there exists a unique shortest geodesic segment with its initial point in $\mathcal{T}_1$ and its terminal point in $\mathcal{T}_2$. Moreover, for all $b \in \mathcal{T}_1$ and $c \in \mathcal{T}_2$, the geodesic segment from b to c must contain this segment.*

2.1.2 Lemma. *(See [6, Chapter 2]) If M is an $\mathbb{R}$-tree and $a, b, c \in M$, then*

1. $d(a,b) + d(b,c) = d(a,c) + 2\operatorname{dist}(b, [a,c])$.

2. $b \in [a,c]$ *if and only if* $d(a,c) = d(a,b) + d(b,c)$.

3. $b \in [a,c]$ *if and only if a and c are in different branches at b.*

2.1.3 Lemma. *Let $\mathcal{T}$ be a closed subtree of M and let $a, b \in M$. Let $e_a \in \mathcal{T}$ be the closest point to a, and let $e_b \in \mathcal{T}$ be the closest point to b. If $e_a \neq e_b$, then*

$$d(a,b) = d(a,e_a) + d(e_a,e_b) + d(b,e_b).$$

Proof. Assume $\mathcal{T}$ is a closed subtree of M and $a, b \in M$. Let $e_a \in \mathcal{T}$ be the closest point to a, and let $e_b \in \mathcal{T}$ be the closest point to b. Assume $e_a \neq e_b$. First we show $e_a \in [a,b]$. Otherwise, a and b are on the same branch at e_a, which implies e_a must be the closest point to both a and b in $\mathcal{T}$, a contradiction. So, $e_a \in [a,b]$, which implies $d(a,b) = d(a,e_a) + d(b,e_a)$ by Lemma 2.1.2. Because e_b is the closest point to b in $\mathcal{T}$ and $e_a \in \mathcal{T}$, by Lemma 2.1.1 we know that $e_b \in [b, e_a]$. Therefore, $d(a,b) = d(a,e_a) + d(b,e_a) = d(a,e_a) + d(e_a,e_b) + d(b,e_b)$.

□

2.1.4 Definition (*Gromov product)***.** For a metric space M and $x, y, w \in M$, define

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

2.1.5 Definition. Let $\delta > 0$. A metric space M is *δ-hyperbolic* if, for all $x, y, z, w \in M$

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

If M is geodesic, then M is δ-hyperbolic if and only if, given $a, b, c \in M$ and any geodesic segments $[a,b], [b,c]$, and $[c,a]$, the segment $[a,b]$ is contained in $([b,c] \cup [c,a])^\delta$. A metric space is *0-hyperbolic* if it is δ-hyperbolic for all $\delta > 0$.

2.1.6 Lemma. *([13, Proposition 6.13]) Any $\mathbb{R}$-tree is 0-hyperbolic. Moreover, any 0-hyperbolic metric space embeds isometrically in an $\mathbb{R}$-tree.*

Next, we define what it means for an $\mathbb{R}$-tree to be finitely generated, and prove some properties of finitely generated $\mathbb{R}$-trees.

2.1.7 Definition. If $A \subseteq M$ is a subset of the $\mathbb{R}$-tree M, let E_A denote the smallest subtree containing A. We call this the $\mathbb{R}$-*tree generated by* A. Note that

$$E_A = \bigcup \{[a_1, a_2] \mid a_1, a_2 \in A\}.$$

The closure $\overline{E_A}$ of E_A is the smallest closed subtree containing A.

2.1.8 Definition. An $\mathbb{R}$-tree M is *finitely generated* if there exists a finite subset $A \subseteq M$ such that $M = E_A$.

Note that if A is finite, $E_A = \overline{E_A}$. Therefore, finitely generated $\mathbb{R}$-trees are complete.

2.1.9 Definition. Let M be an $\mathbb{R}$-tree. If $c \in M$ is such that there do *not* exist $a, b \in M \setminus \{c\}$ with $c \in [a, b]$, then c is called an *endpoint* of M. Equivalently, an endpoint is a point with degree one.

2.1.10 Lemma. *If an $\mathbb{R}$-tree M is finitely generated and C is the set of endpoints of M, then*

1. *if B is a generating set, then $C \subseteq B$;*

2. *the set C generates M.*

Thus, C is the unique least set of generators for M.

Proof. Let M be a finitely generated $\mathbb{R}$-tree. Then M is complete and the diameter D of M is finite. Let B be a generating set for M.

Proof of (1): Assume there is an endpoint $c \in M$ not contained in B. Then, there must exist $a, b \in B$ such that $c \in [a, b]$. But, this is a contradiction because c is an endpoint.

Proof of (2): Let $a \in M$. Let S_a be the set of all segments $[b, c] \subseteq M$ such that $a \in [b, c]$ and order S_a by inclusion. This is a partial ordering on S_a. Let $\{[b_i, c_i] \mid i \in \alpha\}$ be a chain in this partial ordering of segments containing a, where α is some cardinal. Let I be the closure of $\bigcup_{i \in \alpha}[b_i, c_i]$. Then, I is a geodesic segment in M. Clearly $a \in I$, and the length of I is at most D. Therefore $I \in S_a$, and I is an upper bound for the chain, since $[b_i, c_i] \subseteq I$ for all $i \in \mathbb{N}$.

The chain was arbitrary, so any chain has an upper bound. Therefore, by Zorn's Lemma there exists a maximal element of S_a. Let $[b_a, c_a]$ denote such a maximal element. The elements b_a and c_a must be endpoints of M. (Say, for instance, that b_a is not an endpoint. Then there exist $e, f \in M$ such that $b_a \in [e, f]$, and either $[e, c_a]$ or $[f, c_a]$ will contain $[b_a, c_a]$. This would mean $[b_a, c_a]$ was not maximal in S_a.) Therefore, for each $a \in M$, there exist endpoints b_a and c_a so that $a \in [b_a, c_a]$. So, M is generated by the set of its endpoints, and this generating set is as small as possible by (1). □

2.2 The theory of $\mathbb{R}$-trees

In this section we give axioms in the signature L_p for the theory of pointed $\mathbb{R}$-trees and investigate some properties of this theory.

2.2.1 Definition. Let $\mathcal{K}$ be the class of models $\mathcal{M}$ of T_p whose underlying metric space (M, d, p) is a pointed $\mathbb{R}$-tree.

2.2.2 Definition. Let $\mathbb{R}\mathrm{T}$ be the following collection of axioms:

1. the axioms of T_p from Definition 1.8.6; recall specifically the axioms from item (6) of Definition 1.8.6: for each $n \in \mathbb{N}$, the axiom

$$\sup_{n,x} \sup_{n,y} \inf_{n,z} \max\{|d(x,z) - \frac{d(x,y)}{2}|, |d(y,z) - \frac{d(x,y)}{2}|\} = 0;$$

2. for each $n \in \mathbb{N}$ the axiom

$$\sup_{n,x} \sup_{n,y} \sup_{n,z} \sup_{n,w} \big(\min\{(x \cdot z)_w, (y \cdot z)_w\} \dot{-} (x \cdot y)_w \big) = 0.$$

2.2.3 Theorem. *The class $\mathcal{K}$ is exactly the class of models of the L_p-theory $\mathbb{R}\mathrm{T}$. That is, $\mathbb{R}\mathrm{T}$ axiomatizes $\mathcal{K}$.*

Proof. Let $\mathcal{M} \models \mathbb{R}\mathrm{T}$. Then $\mathcal{M} \models T_p$ and therefore has a unique underlying metric space (M, d, p) by Lemma 1.8.9. If we let $x, y, z, w \in M$ and let $m = \max\{d(x,p), d(y,p), d(z,p), d(w,p)\}$, then the axiom

$$\sup_{m,x} \sup_{m,y} \sup_{m,z} \sup_{m,w} \big(\min\{(x \cdot z)_w, (y \cdot z)_w\} \dot{-} (x \cdot y)_w \big) = 0$$

from (2) implies that the metric condition for 0-hyperbolicity given in 2.1.5 holds for x, y, z, w. The points x, y, z, w were arbitrary, so we may conclude that (M, d, p) is 0-hyperbolic. Then by

Lemma 2.1.6 we know that (M, d, p) embeds isometrically in an $\mathbb{R}$-tree. The axioms in item (6) of Definition 1.8.6 guarantee that for any $x, y \in M$, for any $\epsilon > 0$, there exists a point $z \in M$ such that

$$\left|d(x,z) - \frac{d(x,y)}{2}\right| \leq \epsilon, \text{ and } \left|d(y,z) - \frac{d(x,y)}{2}\right| \leq \epsilon.$$

Given a pair of points x, y in M, the existence of these approximate midpoints and the completeness of M may be used to construct a path between x and y. So M is path-connected. A path-connected subspace of an $\mathbb{R}$-tree is an $\mathbb{R}$-tree. Therefore (M, d, p) is a pointed $\mathbb{R}$-tree. Conversely, let (M, d, p) be a pointed $\mathbb{R}$-tree and let $\mathcal{M}$ be the corresponding L_p-structure. Then $\mathcal{M}$ clearly satisfies the axioms in (1)-(5) of the definition of T_p. The fact that M is a uniquely geodesic space with convex balls implies that $\mathcal{M}$ satisfies the axioms in (6). So, $\mathcal{M} \models T_p$, and the fact that M is 0-hyperbolic implies that $\mathcal{M}$ satisfies axioms in (2). □

For a model $\mathcal{M}$ of $\mathbb{R}$T, we will refer to the underlying metric space (M, d, p) of $\mathcal{M}$ as the *underlying* $\mathbb{R}$*-tree of* $\mathcal{M}$.

Now, we present a short discussion of some definability issues in models of $\mathbb{R}$T. Let M be an $\mathbb{R}$-tree and for $r \in [0,1]$ define the function $\nu_r \colon M \times M \to M$ by: $\nu_r(x_1, x_2) =$ the point in $[x_1, x_2]$ with distance $rd(x_1, x_2)$ from x_1 and distance $(1-r)d(x_1, x_2)$ from x_2. Define $\nu_r^{(n)}$ to be the restriction of ν_r to $B_n(p)$.

2.2.4 Lemma. *Let $r \in [0,1]$. The function $\nu_r \colon M \times M \to M$ is a uniformly 0-definable function in models of* $\mathbb{R}$T.

Proof. Let $\mathcal{M} \models \mathbb{R}$T, and let (M, d, p) be the underlying $\mathbb{R}$-tree of $\mathcal{M}$. Let $r \in [0,1]$. For $n \in \mathbb{N}$ let ψ_n be the formula

$$\max\{d(x_1, y) \dot{-} rd(x_1, x_2), d(x_2, y) \dot{-} (1-r)d(x_1, x_2)\}$$

where the variables are from the sort $M^{(n)}$. In $\mathcal{M}$, the distance $d(\nu_r^{(n)}(x_1, x_2), y)$ is equal to $\psi_n^{\mathcal{M}}(x_1, x_2, y)$. So for any $n \in \mathbb{N}$ the function $\nu_r^{(n)}$ is 0-definable via this quantifier-free formula in any model of $\mathbb{R}$T. □

Define $\mu = \nu_{1/2}$. We will use the midpoint function μ extensively in Chapter 3.

Next, we discuss of amalgamation over substructures for the L_p-theory $\mathbb{R}$T. We begin by proving any substructure of a model of $\mathbb{R}$T extends to a unique model of $\mathbb{R}$T.

2.2.5 Lemma. *A substructure $\mathcal{A}$ of $\mathcal{N} \models \mathbb{R}\mathrm{T}$ extends to a model $\mathcal{M} \models \mathbb{R}\mathrm{T}$ such that any embedding of $\mathcal{A}$ into $\mathcal{W} \models \mathbb{R}\mathrm{T}$ extends to an embedding of $\mathcal{M}$ into $\mathcal{W}$.*

Proof. Let (N, d, p) be the underlying $\mathbb{R}$-tree of $\mathcal{N}$. Let A be the underlying metric space of the substructure $\mathcal{A}$. Note that $p^{\mathcal{M}} = p^{\mathcal{A}} \in A$. Then, $\overline{E_A}$ is a closed subtree of N, and since N is complete, $\overline{E_A}$ is complete. Let $\mathcal{M}$ be the L_p-structure corresponding to $\overline{E_A}$ with basepoint $p^{\mathcal{A}}$. Then $\mathcal{M} \models \mathbb{R}\mathrm{T}$. If $\mathcal{W} \models \mathbb{R}\mathrm{T}$, it is straightforward to extend an embedding $\mathcal{A} \to \mathcal{W}$ to an embedding of $\mathcal{M}$. First, extend it isometrically to each $[a, b]$ for $a, b \in A$, then, extend to the closure. □

2.2.6 Theorem. *The L_p-theory $\mathbb{R}\mathrm{T}$ has amalgamation over substructures. That is, if $\mathcal{M}_0$, $\mathcal{M}_1$ and $\mathcal{M}_2$ are substructures of models of $\mathbb{R}\mathrm{T}$ and $f_1 \colon \mathcal{M}_0 \to \mathcal{M}_1$, $f_2 \colon \mathcal{M}_0 \to \mathcal{M}_2$ are embeddings, then there exists a model $\mathcal{N}$ of $\mathbb{R}\mathrm{T}$ and embeddings $g_i \colon \mathcal{M}_i \to \mathcal{N}$ such that $g_1 \circ f_1 = g_2 \circ f_2$.*

Proof. Let $\mathcal{M}_0, \mathcal{M}_1, \mathcal{M}_2$ and the f_i for $i = 1, 2$ be as above. Let (M_i, d_i, p_i) be the underlying metric space of $\mathcal{M}_i$ for $i = 0, 1, 2$. By Lemma 2.2.5 we may assume that each of the $\mathcal{M}_i$ is actually a model of $\mathbb{R}\mathrm{T}$, and therefore M_i is path connected for $i = 0, 1, 2$. In addition, to simplify notation we may assume $M_0 \subseteq M_1$, $M_0 \subseteq M_2$, and $M_1 \cap M_2 = M_0$, that the f_i are inclusion maps and that $p_0 = p_1 = p_2$. Let $N = M_1 \cup M_2$ and define $d \colon N \times N \to \mathbb{R}$ by:

- if $x, y \in M_i$ for $i = 1$ or $i = 2$, then define $d(x, y) = d_i(x, y)$;
- if $x \in M_1$ and $y \in M_2$, define

$$d(y, x) = d(x, y) = \inf\{d_1(x, z) + d_2(z, y) \mid z \in M_0\}.$$

This is a standard construction for the amalgamation of metric spaces, and it is straightforward to show the function d is a metric on N. Note that in the second case of the definition, if one of x or y is in M_0, then the infimum will be realized by setting $z = x$ or $z = y$. Thus the two cases of the definition agree. Also, for $i = 0, 1, 2$, the metric d_i is equal to the restriction of d to $M_i \subseteq N$.

Let $\mathcal{N}$ be the L_p-structure corresponding to (N, d, p) for the (N, d) defined above, where $p = p_0$. Define $g_i \colon M_i \to N$ by $g_i(x) = x$ for $i = 1, 2$. Then, $g_i(p_0) = p_i$ for $i = 1, 2$. So, g_i are embeddings of $\mathcal{M}_i$ into $\mathcal{N}$ for $i = 1, 2$. It is clear that $g_1 \circ f_1 = g_2 \circ f_2$, since on the level of underlying spaces, both maps are just inclusion maps. Because it is the L_p-structure corresponding to a metric space, $\mathcal{N}$ satisfies the axioms from (1)-(5) of Definition 1.8.6 by Lemma 1.8.7. Thus, to show $\mathcal{N} \models \mathbb{R}\mathrm{T}$, it suffices to show that N is an $\mathbb{R}$-tree by Theorem 2.2.3.

Claim: For $x \in M_1$ and $y \in M_2$

$$d(x, y) = d(x, e_y) + d(e_y, y)$$

where e_y is the closest point to y in M_0.

Proof of claim: Because $e_y \in M_0$, by the definition of d we know,

$$d(x, y) \leq d_1(x, e_y) + d_2(e_y, y) = d(x, e_y) + d(e_y, y).$$

Towards contradiction, assume $d(x, y) < d(x, e_y) + d(e_y, y)$. Then there exists $z \in M_0$ such that

$$d(x, z) + d(z, y) = d_1(x, z) + d_2(z, y) < d_1(x, e_y) + d_2(e_y, y) = d(x, e_y) + d(e_y, y).$$

Because $z \in M_0$ and therefore $e_y \in [z, y] \subseteq M_2$ we conclude

$$d(z, y) = d_2(e_y, z) + d_2(y, e_y) = d(e_y, z) + d(y, e_y).$$

Therefore

$$d(x, z) + d(e_y, z) + d(y, e_y) = d(x, z) + d(z, y) < d(x, e_y) + d(e_y, y)$$

which implies

$$d(x, z) + d(e_y, z) < d(x, e_y).$$

This contradicts the triangle inequality for d. Therefore, it must be true that

$$d(x, y) = d(x, e_y) + d(e_y, y).$$

So, the claim is proved.

By our construction, $N = M_1 \cup M_2$. Let $x, y \in N$. If $x \in M_1$ and $y \in M_2$, then since $d(x, y) = d(x, e_y) + d(e_y, y)$, we may concatenate $[x, e_y]$ and $[e_y, y]$ to get a geodesic segment between x and y in N. Denote this geodesic segment by $[x, y]$. We must show $[x, y]$ is the unique arc connecting x and y in N. Assume $\alpha : [0, t] \to N$ is a continuous map with $\alpha(0) = x$ and $\alpha(t) = y$. For convenience let α also refer to the image of α. Towards contradiction, assume that setwise, $\alpha \neq [x,y]$. We conclude that either $\alpha \cap M_1 \neq [x,y] \cap M_1$ or $\alpha \cap M_2 \neq [x,y] \cap M_2$ because if both of these were equal it would contradict our assumption. Without loss of generality, say $\alpha \cap M_1 \neq [x,y] \cap M_1$. Let $a, b \in M_1$ such that $a \in [x,y] \setminus \alpha$ and $b \in \alpha \setminus [x,y]$. Now, $\gamma = [x,a] \cup [a,b]$ is an arc in M_1 from x to b. If we take α restricted to the interval $[0, \alpha^{-1}(b)]$ we get another arc α' in M_1 from x to b. Since $a \in \gamma$ but is not on α' we know these two arcs are distinct. Then, $g_1^{-1}(\gamma)$ and $g_1^{-1}(\alpha')$ are distinct arcs in M_1, which can't happen because M_1 is an $\mathbb{R}$-tree. This is our contradiction. So, we can't have more than one arc from x to y in N.

If $x, y \in M_1$, then since M_1 is an $\mathbb{R}$-tree there is a geodesic segment $[x,y]$ in $M_i \subseteq N$ and moreover this is the only path between x and y contained in M_1. Towards contradiction, suppose there is an arc α between x and y in N not equal to $[x,y]$. As before let α denote both the map and its image in N. Since α is a different arc, $\alpha \cap M_1 \neq [x,y] \cap M_1$, and we may proceed as above. If $x, y \in M_2$ the same argument works, just switch the subscripts 1 and 2. Thus, between any $x, y \in N$ there is a unique arc, and that arc is a geodesic segment. Therefore N is an $\mathbb{R}$-tree. □

Intuitively, it is clear that one may build an $\mathbb{R}$-tree by "gluing" copies of intervals in $\mathbb{R}$ and other $\mathbb{R}$-trees together. The following lemmas capture these ideas.

2.2.7 Lemma. *Let Λ be a linearly ordered set. If $(\mathcal{M}_\lambda \mid \lambda \in \Lambda)$ is a chain of models of $\mathbb{R}\mathrm{T}$, then the union of this chain is a model of $\mathbb{R}\mathrm{T}$.*

Proof. Note that $\mathbb{R}\mathrm{T}$ is an $\forall\exists$-theory. Therefore, the conclusion is immediate from Proposition 1.7.4. □

This lemma implies that the union of any increasing chain of $\mathbb{R}$-trees is an $\mathbb{R}$-tree. In [6, Lemma 2.1.14], this was proved for countable increasing chains of $\mathbb{R}$-trees, and the proof there actually also works for chains of arbitrary length.

2.2.8 Lemma. *Let $\mathcal{M}$ be a model of $\mathbb{R}\mathrm{T}$ with underlying $\mathbb{R}$-tree (M, d, p). Let $(a_i \mid i \in I)$ for an ordinal I be a list of elements of M, which are not necessarily distinct. For each $i \in I$, let $\mathcal{T}_i = (T_i, d_i, p_i) \models \mathbb{R}\mathrm{T}$. Then we may construct an extension $\mathcal{N} \models \mathbb{R}\mathrm{T}$ of $\mathcal{M}$, such that there is a copy of T_i isometrically embedded in $(N \setminus M) \cup \{a_i\}$ via an isometry taking p_i to a_i for all $i \in I$. Moreover, if $i \neq j$ and $a_i \neq a_j$, then the images of T_i and T_j in N are disjoint, and if $i \neq j$ but $a_i = a_j$, then the images of T_i and T_j intersect only at $\{a_i\}$.*

Proof. Assume the situation described in the hypotheses. Define $\mathcal{M}_0$ to be the result of amalgamating $\mathcal{M}$ with $\mathcal{T}_0$ over a substructure consisting of a point. The point in M is a_0, and the point in T_0 is p_0. Then, by the proof of Lemma 2.2.6, there is a copy of T_0 isometrically embedded in $(M_0 \setminus M) \cup \{a_0\}$ via an isometry taking p_0 to a_0. Note that this copy of T_0 does not intersect M except at a_0.

For $i \in I$, define $\mathcal{M}_{i+1}$ to be the result of amalgamating $\mathcal{M}_i$ with $\mathcal{T}_{i+1}$ over the point $a_{i+1} \in M \subset M_i$, and the point p_{i+1} in T_{i+1}. Then, by the proof of Lemma 2.2.6 there is a copy of T_{i+1} isometrically embedded in $(M_{i+1} \setminus M_i) \cup \{a_{i+1}\} \subseteq (M_{i+1} \setminus M) \cup \{a_{i+1}\}$. Note that if a_{i+1} is equal to a_j for some $j < i+1$, then the images of T_{i+1} and T_j overlap only at $a_{i+1} = a_j$. If $a_{i+1} \neq a_j$ for all $j < i+1$, then the image of T_{i+1} is disjoint from the images of the T_j for $j < i+1$, because all those images are contained in $M \setminus \{a_{i+1}\}$. If $i \in I$ is a limit ordinal, first let $\widehat{\mathcal{M}_i}$ be the union of the chain of models $(\mathcal{M}_j \mid j < i)$. Then, define $\mathcal{M}_i$ to be the result of amalgamating $\widehat{\mathcal{M}_i}$ with $\mathcal{T}_i$ over the point $p_i \in T_i$ and the point $a_i \in \widehat{\mathcal{M}_i}$. Each $\mathcal{M}_j$ is an extension of $\mathcal{M}$, and an extension of $\mathcal{M}_i$ for all $i < j$. Thus, we get an increasing chain of models extending $\mathcal{M}$ indexed by I. Let $\mathcal{N}$ be the completion of $\bigcup_{i \in I} \mathcal{M}_i$. Then $\mathcal{N} \models \mathbb{R}\mathrm{T}$ by Lemma 2.2.7 and clearly $\mathcal{N}$ extends $\mathcal{M}$. $\square$

2.2.9 Lemma. *There exists an $\mathbb{R}$-tree N such that every point in N has degree three.*

Proof. Let N_0 be the $\mathbb{R}$-tree $\mathbb{R}$ with basepoint 0. Let $\mathcal{N}_0$ be the corresponding L_p-structure, which is a model of $\mathbb{R}\mathrm{T}$. Let $\mathcal{M} = \mathcal{N}_0$ and let $(a_i \mid i \in I)$ be an enumeration of N_0. For each i, let the $\mathbb{R}$-tree T_i be $\mathbb{R}^{\geq 0}$ with basepoint 0, and let $\mathcal{T}_i$ be the corresponding model of $\mathbb{R}\mathrm{T}$. Define $\mathcal{N}_1$ to be the extension of $\mathcal{N}_0$ that results from the application of Lemma 2.2.8 and let (N_1, d, p) be its underlying $\mathbb{R}$-tree. There are now three branches in N_1 at each point of N_0. However, the points on the rays we added (with the exception of the basepoints) all still have only two branches. Let $\{a'_i \mid i \in I'\}$ be an enumeration of $N_1 \setminus N_0$ and apply Lemma 2.2.8, now with and $\mathcal{M} = \mathcal{N}_1$, and the same $\mathcal{T}_i$, but now defined for all $i \in I'$. Call the resulting extension $\mathcal{N}_2$. In the underlying $\mathbb{R}$-tree N_2 of $\mathcal{N}_2$ there are exactly three branches in N_2 out of each point of N_1. Generally, for $n \geq 1$, enumerate $N_n \setminus N_{n-1}$ and apply Lemma 2.2.8 to get an extension $\mathcal{N}_{n+1}$ of $\mathcal{N}_n$ such that in the $\mathbb{R}$-tree N_{n+1}, all the points in N_n have exactly three branches and all the points in $N_{n+1} \setminus N_n$ have exactly two branches. Define $\mathcal{N} = \bigcup_{n \in \omega} \mathcal{N}_n$. This will be a model of $\mathbb{R}\mathrm{T}$ such that there are exactly three branches at every point in its underlying metric space N. $\square$

2.2.10 Remark. The lemma above just is one example of how amalgamation may be used to construct an example of an $\mathbb{R}$-tree with a specific branching pattern. Obviously, the construction could be modified to produce $\mathbb{R}$-trees with quite complicated patterns of branch points of varying degrees. Another construction of an $\mathbb{R}$-tree with three branches at every point is mentioned in Example 2.4.2.

In the rest of this thesis, we will not always explicitly refer to the amalgamation construction whenever we want to construct $\mathbb{R}$-trees. We may use the more intuitive language which describes "adding" rays or intervals, or "gluing" things together at a point.

2.3 Richly branching $\mathbb{R}$-trees

In this section we find the model companion of $\mathbb{R}$T.

2.3.1 Definition. Let $h > 0$. An $\mathbb{R}$-tree (M, d) is *h-richly branching* if the set

$$B_h := \{b \in M | \text{ at } b \text{ there are at least three branches of height } \geq h\}$$

is dense in (M, d). If (M, d) is h-richly branching for some h, we say it is *richly branching.*

2.3.2 Lemma. *Suppose (M, d) is an h-richly branching $\mathbb{R}$-tree for some $h > 0$. Let $a \in M$ and let β be a branch at a. Then β has infinite height.*

Proof. Assume (M, d) is an h-richly branching $\mathbb{R}$-tree, let $a \in M$ and β a branch at a. It suffices to show that for any $r \in \mathbb{R}^{>0}$, there exists $b \in \beta$ such that $d(a, b) \geq r$. Let $r \in \mathbb{R}^{>0}$. Let k be the smallest integer larger than $\frac{r}{h}$. Let $\epsilon = \frac{h}{2k}$. Let $a_1 \in B_h \cap \beta$ be such that $d(a, a_1) \geq h - \epsilon$. Such an a_1 exists by the density of B_h in M, the path-connectedness of β and the fact that the height of β is at least h. There are at least three branches at a_1 of length at least h. Let β_1 be a branch at a_1 that does not contain a. So setwise, β_1 is contained in β. Let $a_2 \in B_h$ be a point in β_1 with $d(a_1, a_2) \geq h - \epsilon$. Note that $a_1 \in [a, a_2]$ by how we chose a_2. We may find a_2 for the same reasons we may find a_1. Continue like this to get $a = a_0, a_1, a_2, ..., a_{k+1} \in \beta$ such that $d(a_i, a_{i+1}) \geq h - \epsilon$ for all $0 \leq i \leq k$ and $a_i \in [a_{i-1}, a_{i+1}]$ for all $1 \leq i \leq k$. Using Lemma 2.1.2 we know that

$$d(a_0, a_{k+1}) = d(a_0, a_1) + d(a_1, a_2) + ... + d(a_{k-1}, a_k) + d(a_k, a_{k+1}).$$

Then

$$\begin{aligned} d(a,a_{k+1}) &= \big(\sum_{i=0}^{k-1} d(a_i,a_{i+1})\big) + d(a_k,a_{k+1}) \geq k(h-\epsilon)+(h-\epsilon) \\ &= kh - k\epsilon + (h-\epsilon) \geq r - \frac{h}{2} + (h - \frac{h}{2k}) \geq r. \end{aligned}$$

Let $b = a_{k+1}$. Then, $d(a,b) \geq r$. ☐

This lemma has the following immediate corollary.

2.3.3 Corollary. *An $\mathbb{R}$-tree (M,d) is h-richly branching for some $h > 0$ if and only if it is h-richly branching for every $h > 0$.*

Next we give axioms for the class of richly branching $\mathbb{R}$-trees.

2.3.4 Definition. Given $n \in \mathbb{N}$ let $\widetilde{\varphi}_n(x)$ be the L_p-formula

$$\inf_{2n,y_1} \inf_{2n,y_2} \inf_{2n,y_3} \max_{1\leq i<j\leq 3} \{|(d(x,y_i)+d(x,y_j)) - d(y_i,y_j)|, |d(x,y_i) - \frac{1}{2}|\}$$

where the variable x is of sort $M^{(n)}$. Let φ_n be the closed formula $\sup\limits_{n,x} \widetilde{\varphi}_n$.

2.3.5 Definition. Let $\mathrm{rb}\mathbb{R}\mathrm{T} = \mathbb{R}\mathrm{T} \cup \{\varphi_n = 0 \mid n \in \mathbb{N}\}$.

2.3.6 Theorem. *The L_p-theory* $\mathrm{rb}\mathbb{R}\mathrm{T}$ *axiomatizes the class of models of T_p whose underlying metric spaces are richly branching $\mathbb{R}$-trees.*

Proof. Let $\mathcal{M} \models T_p$ be such that its underlying metric space (M,d,p) is a richly branching $\mathbb{R}$-tree. Then $\mathcal{M} \models \mathbb{R}\mathrm{T}$ by Theorem 2.2.3. Also, for every $h > 0$ the set B_h is dense in M. Let $n \in \mathbb{N}$ and let $a \in M^{(n)}$. Since $B_{1/2}$ is dense in M, the set $B_{1/2} \cap M^{(n)}$ is dense in $M^{(n)}$. For any $b \in B_{1/2} \cap M^{(n)}$, there exist $c_1, c_2, c_3 \in M^{(2n)}$ that are each distance $\frac{1}{2}$ from b and each on a different branch out of b. These exist by the definition of $B_{1/2}$, and by Lemma 2.1.2 they witness the fact that $\widetilde{\varphi}_n(b)^{\mathcal{M}} = 0$ for $b \in B_{1/2} \cap M^{(n)}$. Now, the uniform continuity of $\widetilde{\varphi}_n$ and the density of $B_{1/2} \cap M^{(n)}$ in $M^{(n)}$ imply that $\widetilde{\varphi}_n(a)^{\mathcal{M}} = 0$ for any $a \in M^{(n)}$. Therefore the condition $\varphi_n = 0$ is true in $\mathcal{M}$. Since n was arbitrary, we conclude that $\mathcal{M} \models \mathrm{rb}\mathbb{R}\mathrm{T}$.

Now, assume the L_p-structure $\mathcal{M} \models \mathrm{rb}\mathbb{R}\mathrm{T}$. Assume $\frac{1}{16} > \epsilon > 0$ and $a \in M$. Let $\mathcal{O}_\epsilon(a)$ denote the open ball of radius ϵ in M centered at a. Let n be large enough so that $\mathcal{O}_\epsilon(a) \subseteq M^{(n)}$. By our assumption, $\widetilde{\varphi}_n(a)^{\mathcal{M}} = 0$. We will show that $B_{1/4} \cap \mathcal{O}_\epsilon(a) \neq \emptyset$.

Because $\widetilde{\varphi}_n(a)^{\mathcal{M}} = 0$, we know there exist distinct $b_1, b_2, b_3 \in M^{(2n)} \subseteq M$ such that for $i \neq j \in \{1, 2, 3\}$

$$|d(a, b_i) - \frac{1}{2}| \leq \epsilon \text{ and } 2\,\mathrm{dist}(a, [b_i, b_j]) = |d(a, b_i) + d(a, b_j) - d(b_i, b_j)| \leq \epsilon.$$

Let $z \in [b_1, b_2]$ be the closest point to a on $[b_1, b_2]$. Then, there are two cases.

Case I: The point z is also the closest point to b_3 on $[b_1, b_2]$. Then we know that

$$d(z, a) = \mathrm{dist}(a, [b_1, b_2]) \leq \frac{\epsilon}{2}.$$

We also know that

$$d(z, b_i) \geq d(a, b_i) - d(z, a) \geq (\frac{1}{2} - \epsilon) - d(z, a) \geq (\frac{1}{2} - \epsilon) - \frac{\epsilon}{2} = \frac{1}{2} - \frac{3\epsilon}{2} > \frac{1}{4}$$

for $i = 1, 2, 3$, since we assumed $\epsilon < \frac{1}{16}$. It is also clear that $z \in [b_1, b_2]$ (since it is the closest point to a on that segment), and $z \in [b_2, b_3]$ and $z \in [b_1, b_3]$ (since z is the closest point to b_3 on $[b_1, b_2]$, which means the paths from b_3 to b_1 and from b_3 to b_2 must go through z.) So, there are at least three distinct branches at z. Therefore in this case, the point z is in $B_{1/4} \cap \mathcal{O}_\epsilon(a)$.

Case II: The point z is not the closest point to b_3 on $[b_1, b_2]$. Then let y be the closest point to b_3 on $[b_1, b_2]$. Note that $y \in [b_1, b_2]$, and $y \in [b_1, b_3]$ and $y \in [b_2, b_3]$. Also, $z \in [a, y]$.

Claim: The point y is either the closest point to a on $[b_1, b_3]$ or the closest point to a on $[b_2, b_3]$.

Proof of claim: We know either $y \in [z, b_1]$ or $y \in [z, b_2]$. Assume $y \in [z, b_1]$. Then $y \in [a, b_1]$, since $[a, b_1]$ must include the segment $[z, b_1]$. Because $z \in [a, y]$ and $y \in [z, b_3]$ we know $[a, b_3] \subseteq [a, y] \cup [z, b_3]$, and $[a, y] \cap [z, b_3] = [y, z]$. Thus $[a, b_3] = [a, z] \cup [z, y] \cup [y, b_3]$, and y is on $[a, b_3]$. So $y \in [a, b_1] \cap [a, b_3] \cap [b_1, b_3]$, making it the point on $[b_1, b_3]$ closest to a. The proof in the case where $y \in [z, b_2]$ is exactly analogous to the above, with the conclusion that y is the closest point to a on $[b_2, b_3]$. So the claim is proved.

Now, assume for convenience that y is the closest point to a on $[b_1, b_3]$. So, we know that

$$d(a, y) = \mathrm{dist}(a, [b_1, b_3]) \leq \frac{\epsilon}{2}.$$

We also know that

$$d(y,b_i) \geq d(a,b_i) - d(y,a) \geq (\frac{1}{2} - \epsilon) - d(y,a) \geq (\frac{1}{2} - \epsilon) - \frac{\epsilon}{2} = \frac{1}{2} - \frac{3\epsilon}{2} > \frac{1}{4}.$$

Therefore, $y \in B_{1/4} \cap \mathcal{O}_\epsilon(a)$. (The proof that $y \in B_{1/4} \cap \mathcal{O}_\epsilon(a)$ in the case where y is the closest point to a on $[b_2, b_3]$ goes exactly the same way.) Thus, $B_{1/4} \cap \mathcal{O}_\epsilon(a) \neq \emptyset$. Our $\frac{1}{16} > \epsilon > 0$ was arbitrary, and therefore, $B_{1/4}$ is dense in M. Therefore the underlying metric space of $\mathcal{M}$ is a richly branching $\mathbb{R}$-tree. □

2.3.7 Lemma. *Assume* $\mathcal{M} \models$ rb$\mathbb{R}$T.

1. *Let $\mathcal{M}$ be ω-saturated. Then for any $m \in \mathbb{N}$, there are at least m branches at any point.*
2. *Let κ be an infinite cardinal. If $\mathcal{M}$ is κ-saturated, then there are at least κ branches at every point.*

Proof. Let $\mathcal{M} \models$ rb$\mathbb{R}$T.

Proof of (1): Assume $\mathcal{M}$ is ω-saturated. For $m, n \in \mathbb{N}$, let $\sigma_{n,m}(x)$ be the following L_p-formula, where x is of sort n:

$$\inf_{2n,y_1} ..., \inf_{2n,y_m} \max_{1\leq i<j\leq m} \{|(d(x,y_i) + d(x,y_j)) - d(y_i,y_j)|, |d(x,y_i) - \frac{1}{2}|\}.$$

If the condition $\sup_{n,x} \sigma_{n,m}(x) = 0$ is true in $\mathcal{M}$, then, by 1.7.3, for any $a \in M$ there are points $b_1, ..., b_m \in M^{(2n)}$ so that

$$|(d(a,b_i) + d(a,b_j)) - d(b_i,b_j)| = 0, \quad \text{and} \quad |d(a,b_i) - \frac{1}{2}| = 0.$$

By Lemma 2.1.2, the points $b_1, ..., b_m$ are distinct and in separate branches at a. Then there must be at least m branches at a.

Thus it suffices to show each $\sup_{n,x} \sigma_{n,m}(x) = 0$ is true in $\mathcal{M}$. Let $a \in M^{(n)}$ and let $\frac{1}{2} > \epsilon > 0$. Since $\mathcal{M} \models$ rb$\mathbb{R}$T we know the axiom $\varphi_n = 0$ is true in $\mathcal{M}$, and therefore $\widetilde{\varphi}_n^{\mathcal{M}}(a) = 0$. By ω-saturation, we know that the quantifiers in this formula are realized exactly. Thus, by Lemma 2.1.2 there are at least 3 branches of height at least $\frac{1}{2}$ at *every* point. Let $c_1, ..., c_m$ be such that $c_i \in [c_{i-1}, c_{i+1}]$ and $d(c_i, a) = \frac{\epsilon}{2^i}$. For each $i = 1, ..., m$ let b_i be a point on a branch at c_i that does not contain any of the c_j for $i \neq j$. Such a branch exists at every c_i because each of them has at least 3 branches, and by how we chose them, only two of those branches will contain any

c_j for $j \neq i$. Note that this makes c_i the closest point to a on $[b_i, b_j]$. Let $d(b_i, c_i) = \frac{1}{2}$. The distance

$$d(a, b_i) = d(a, c_i) + d(c_i, b_i) < \frac{1}{2} + \frac{\epsilon}{2^i}$$

and clearly $d(a, b_i) > d(c_i, b_i) = \frac{1}{2}$, so $|d(a, b_i) - \frac{1}{2}| < \epsilon$. Also,

$$|(d(a, b_i) + d(a, b_j)) - d(b_i, b_j)| = 2\,\mathrm{dist}(a, [b_i, b_j]) = 2d(a, c_i) = \frac{\epsilon}{2^i} < \epsilon.$$

Since $\frac{1}{2} > \epsilon > 0$ was arbitrary, we conclude $\sup\limits_{n,x} \sigma^{\mathcal{M}}_{n,m}(x) = 0$

Proof of (2): Assume $\mathcal{M}$ is κ-saturated and towards contradiction assume there is a point $b \in M$ such that there are $< \kappa$ branches at b. Let n be large enough so that $b \in M^{(n)}$. Enumerate all the branches at b by $\{\beta_i \mid i < \alpha\}$ for some $\alpha < \kappa$. Let $a_i \in \beta_i$ be the point on β_i with $d(a_i, b) = \frac{1}{2}$, and let B be the set of parameters $\{b\} \cup \{a_i \mid i < \alpha\}$. Note that $|B| < \kappa$. Let Σ be the following set of $L_p(B)$ conditions: for each $i < \alpha$ and $\epsilon > 0$

$$E_i(x) := \max\{|(d(b, x) + d(b, a_i)) - d(a_i, x)|, |d(b, x) - \frac{1}{2}|\} = 0$$

where x is a variable from sort $M^{(2n)}$. Take a finite subset Γ of Σ and let $G = \{i \mid E_i(x) \in \Gamma\}$. So, G is finite. By the first part of this lemma there are at least $|G|+1$ branches at b. Therefore, Γ is satisfiable. Simply take x on a branch β such that $\beta \neq a_i$ for any $i \in G$. Now, by the saturation assumption we know there is an x that satisfies Σ. This x is on a branch not mentioned in our list of branches. This is a contradiction. Therefore there must be at least κ branches at every point in M. $\square$

2.4 Example: Universal $\mathbb{R}$-trees

This section presents some examples of richly branching $\mathbb{R}$-trees.

2.4.1 Definition. Let μ be a cardinal. An $\mathbb{R}$-tree M is called *μ-universal* if, for any $\mathbb{R}$-tree N with $\leq \mu$ branches at every point, there is an isometric embedding of N into M.

2.4.2 Example. ([10, Lemma 2.1.1]) In [10], the authors construct examples of complete, homogeneous μ-universal $\mathbb{R}$-trees A_μ for $\mu \geq 2$. These spaces are unique up to isometry in that any complete homogeneous $\mathbb{R}$-tree with μ branches at each point is isometric to A_μ. Clearly, for $\mu \geq 3$, these spaces A_μ are examples of richly branching $\mathbb{R}$-trees.

A group is *hyperbolic* if its Cayley graph is a δ-hyperbolic metric space for some $\delta > 0$ (see [5, Chapter III.Γ Definition 2.2.1]). A *non-elementary* hyperbolic group is one that does not have a cyclic subgroup of finite index.

2.4.3 Definition. Let (M, d, p) be a family of pointed metric spaces. Let U be a non-principal ultrafilter on $\mathbb{N}$ and let $(\nu_m)_{m \in \mathbb{N}}$ be a sequence of positive integers such that $\lim_{n \to \infty} \nu_n = \infty$. The asymptotic cone of (M, d, p) with respect to $(\nu_m)_{m \in \mathbb{N}}$ and U is the ultraproduct of pointed metric spaces $\prod_U (M, \frac{d}{\nu_m}, p_m)$. Denote the asymptotic cone of (M, d, p) with respect to $(\nu_m)_{m \in \mathbb{N}}$ and U by $A_{U,(\nu_m)}(M, d, p)$.

The asymptotic cone of a finitely generated group is defined to be the asymptotic cone of its Cayley graph. It is a fact that the asymptotic cone of a hyperbolic group is an $\mathbb{R}$-tree and is homogeneous. If the hyperbolic group is non-elementary, then its asymptotic cone is actually isomorphic to A_{2^ω} (see [9]).

2.4.4 Proposition. *([10] or [9]) Let G be a non-elementary hyperbolic group. Let U be a non-principal ultrafilter on $\mathbb{N}$. By [9, Proposition 3.1.1] the asymptotic cone of G is a richly branching $\mathbb{R}$-tree.*

2.5 The model companion of $\mathbb{R}$T

In this section we to build up to the proof that the theory rb$\mathbb{R}$T of richly branching $\mathbb{R}$-trees is the model companion of $\mathbb{R}$T.

2.5.1 Definition. Let $\mathcal{M} \models T_p$ with underlying metric space (M, d, p), and $a = a_1, ..., a_k \in M$. For convenience assume $a_1 = p$. Let $b = b_1, ..., b_n \in M$. For $y_1, ..., y_n$ from the appropriate sorts, define the partial type $D_b^{\mathcal{M}}(y_1, ..., y_n/a)$ to be:

$$\{|d(a_l, y_j) - d^{\mathcal{M}}(a_l, b_j)| = 0, |d(y_i, y_j) - d^{\mathcal{M}}(b_i, b_j)| = 0 \mid i, j = 1, ..., n;\ l = 1, ..., k\}.$$

2.5.2 Lemma. *Let $\mathcal{M} \models T_p$ with underlying metric space (M, d, p), and $a = a_1, ..., a_k \in M$. For convenience assume $a_1 = p$. Let $b = b_1, ..., b_n \in M$. If $c_1, ..., c_n \in M$ are such that $(c_1, ..., c_n) \models D_b^{\mathcal{M}}(y_1, ..., y_n/a)$, then for any quantifier free formula $\varphi(x_1, ..., x_k, y_1, ..., y_n)$ we have $\varphi(a_1, ..., a_k, b_1, ..., b_n)^{\mathcal{M}} = \varphi(a_1, ..., a_k, c_1, ..., c_n)^{\mathcal{M}}$.*

Proof. This is proved by an easy induction on the definition of quantifer free formulas. □

The following lemma is the main tool for proving Theorem 2.5.4.

2.5.3 Lemma. *Assume $\mathcal{M}$ is an ω_1-saturated model of* rb$\mathbb{R}$T*. Let K be a non-empty, finitely generated $\mathbb{R}$-tree with basepoint p. For any $e \in M$ and any countable collection $\{\beta_i \mid i \in \omega\}$ of branches at e, there exists an isometric embedding f of K into M such that $f(p) = e$ and $f(k) \cap \beta_i = \{e\}$ for all $i \in \omega$.*

Proof. Assume $\mathcal{M}$ is an ω_1-saturated model of rb$\mathbb{R}$T. Let K be a non-empty, finitely generated $\mathbb{R}$-tree with basepoint p. By Lemma 2.1.10, there is a minimal set of generators for K, namely, the set of endpoints of K. We will prove this lemma by induction on the size of this minimal generating set.

Base Case: Assume K has one generator, namely the basepoint p. So, $K = \{p\}$. Let $e \in M$, and define $f : K \to M$ by $f(p) = e$. This isometric embedding clearly satisfies $f(K) \cap \beta = \{e\}$ for every branch at e. So the lemma is true when K has one generator.

Inductive Step: Let K have $n > 1$ generators. Assume that if K' is an $\mathbb{R}$-tree that has fewer than n generators and basepoint p', then for any $e \in M$ and any collection $\{\beta_i \mid i \in \omega\}$ of branches at e, there exists an isometric embedding f' of K' into M sending p' to e such that $f'(K') \cap (\cup_{i\in\omega}\beta_i) = \{e\}$. Let $e \in M$ and take an arbitrary collection $\{\beta_i \mid i \in \omega\}$ of branches at e. There must be generators b_1, b_2 of K so that $p \in [b_1, b_2]$. List these two, plus the rest of the generators as $b_1, ..., b_n$.

Case I: If $n = 2$, then we just have 2 generators b_1 and b_2 and $K = [b_1, b_2]$. First, assume p is a generator (say it is b_1.) Find a branch β at e that does not intersect any branch from $\{\beta_i \mid i \in \omega\}$ (except at e). This branch exists by Lemma 2.3.7 because $\mathcal{M}$ is ω_1-saturated and hence has uncountably many branches at e. The branch has infinite height by Lemma 2.3.2, so we may find a point $c \in \beta$ such that $d(e, c) = d(p, b_2)$. Both $[e, c] \subseteq M$ and $[p, b_2]$ are isometric to the interval $[0, d(e, c)]$ in $\mathbb{R}$. So, there is an isometry sending p to e with the properties we wanted. Now, assume p is not a generator. Using Lemma 2.3.7, find two branches at e that only intersect the members of $\{\beta_i \mid i \in \omega\}$ at e. On one of these branches let c_1 be the point that is $d(p, b_1)$ away from e. On the other, let c_2 be the point that is $d(p, b_2)$ away from e. Let f be the isometric embedding sending $K = [b_1, b_2]$ to $[c_1, c_2]$, with $f(b_1) = c_1$ and $f(b_2) = c_2$. Then clearly $f(p) = e$ and $f(K)$ intersects $\{\beta_i \mid i \in \omega\}$ only at e.

Case II: If $n \geq 3$, let K' be the $\mathbb{R}$-tree generated by $b_1, ..., b_{n-1}$ and let $f' : K' \to M$ be an isometric embedding such that $f'(p) = e$ and $f'(K') \cap (\cup_{i\in\omega}\beta_i) = \{e\}$. Note that K' is a closed subtree of K. Let a be the closest point in K' to b_n. Look at $f'(a) \in M$. Find a branch at $f'(a)$

that only intersects $f'(K')$ at a. We know this branch has infinite height. So, we may let c be a point on this branch with distance $d(a, b_n)$ from $f'(a)$, and extend f' by sending the segment $[a, b_n] \in K$ to the segment $[f'(a), c]$ isometrically. Let f be this extension of f'. Clearly f is an isometric embedding which sends p to e and $f(K) \cap (\cup_{i\in\omega}\beta_i) = \{e\}$. $\square$

2.5.4 Theorem. *The L_p-theory* rb$\mathbb{R}$T *is the model companion of* $\mathbb{R}$T.

Proof. Clearly any model of rb$\mathbb{R}$T is a model of $\mathbb{R}$T. If $\mathcal{M}$ is a model of $\mathbb{R}$T, $\mathcal{M}$ can be extended to a model of rb$\mathbb{R}$T. For example, use Lemma 2.2.8 to glue in a copy of the 3-branching $\mathbb{R}$-tree from Lemma 2.2.9 at every point in M. By the definition of model companion (Definition 1.7.10), it remains to show that rb$\mathbb{R}$T is model complete. We will show the equivalent condition given in Proposition 1.7.9: if every $\mathcal{M} \models$ rb$\mathbb{R}$T, is an existentially closed model of rb$\mathbb{R}$T, then rb$\mathbb{R}$T is model complete. Let $\mathcal{N} \models$ rb$\mathbb{R}$T be an extension of $\mathcal{M}$. We may assume $\mathcal{M}$ and $\mathcal{N}$ are ω_1-saturated. (This is because we may consider the structure $\big((M, d, p), (N, d, p), \iota\big)$ where ι is the embedding from $\mathcal{M}$ to $\mathcal{N}$, and take an ω_1-saturated extension of that structure. If we can verify the definition of existentially closed in that setting, it will be true of $\mathcal{M}$ and $\mathcal{N}$.) Let $a_1, ..., a_k \in M$.

Claim: For any $b_1, ..., b_n \in N$, there exist $c_1, ..., c_n \in M$ with $c_1, ..., c_n \models D_b^{\mathcal{N}}(y_1, ..., y_n/a)$.

Proof of claim: Let $b_1, ..., b_n \in N$. Recall that $E_a \subseteq M$ is the subtree generated by $a_1, ..., a_k$. Define an equivalence relation on $b_1, ..., b_n$ by $b_i \sim b_j$ if b_i and b_j have the same closest point in E_a. Let $A_1, ..., A_m$ be the equivalence classes of this equivalence relation, and for $j = 1, ..., m$ let e_j be the unique closest point in E_a common to the members of A_j. For each $j = 1, ..., m$ let K_j be the $\mathbb{R}$-tree generated by $A_j \cup \{e_j\}$, with basepoint $p_j = e_j$. Note that each K_j is closed and for $i \neq j$, $K_i \cap K_j = \emptyset$. For each $j = 1, ..., m$, by Lemma 2.5.3, there is an isometric embedding $f_j \colon K_j \to M$ sending p_j to e_j such that $f_j(K_j)$ does not intersect E_a except at e_j. Note that e_j is the unique closest point in E_a for every point in $f_j(K_j)$.

Let f be the union of the functions f_j for $j = 1, ..., m$. If b_i and b_j are both in A_l, then $d(b_i, b_j) = d(f_l(b_i), f_l(b_j)) = d(f(b_i), f(b_j))$. If b_i and b_j are in $A_l \neq A_m$ respectively, then

$$\begin{aligned} d(b_i, b_j) &= d(b_i, e_l) + d(e_l, e_m) + d(e_m, b_j) \\ &= d(f_l(b_i), f_l(e_l)) + d(e_l, e_m) + d(f_m(e_m), f_m(b_j)) \\ &= d(f(b_i), e_l) + d(e_l, e_m) + d(e_m, f(b_j)) \\ &= d(f(b_i), f(b_j)). \end{aligned}$$

Therefore the function f is an isometric embedding from $\bigcup_{j=1}^{m} K_j$ to M. Let $c_i = f(b_i)$ for all $i \in \{1, ..., n\}$. Then clearly $d(b_i, b_j) = d(f(b_i), f(b_j)) = d(c_i, c_j)$ for all $i, j \in \{1, ..., n\}$. Now let $l \in \{1, ..., k\}$ and $i \in \{1, ..., n\}$. Let e_j be the closest point to b_i in E_a. Then

$$d(a_l, b_i) = d(a_l, e_j) + d(e_j, b_i) = d(a_l, e_j) + d(f(e_j), f(b_i)) = d(a_l, e_j) + d(e_j, c_i) = d(a_l, c_i).$$

Thus, the claim is true.

Now, it follows by Lemma 2.5.2 that for any quantifier free formula $\varphi(x_1, ..., x_k, y_1, ..., y_n)$ and any $b_1, ..., b_n \in N$ there exist $c_1, ..., c_n \in M$ such that

$$\varphi(a_1, ..., a_k, b_1, ..., b_n)^{\mathcal{N}} = \varphi(a_1, ..., a_k, c_1, ..., c_n)^{\mathcal{N}} = \varphi(a_1, ..., a_k, c_1, ..., c_n)^{\mathcal{M}}.$$

So, if $\epsilon > 0$ and $b_1, ..., b_n \in N$ such that

$$|\varphi(a_1, ..., a_k, b_1, ..., b_n)^{\mathcal{N}} - \inf_{s_1, y_1} ... \inf_{s_n, y_n} \varphi(a_1, ..., a_k, y_1, ..., y_n)^{\mathcal{N}}| \leq \epsilon$$

we know we can find $c_1, ..., c_n \in M$ with

$$\varphi(a_1, ..., a_k, c_1, ..., c_n)^{\mathcal{M}} = \varphi(a_1, ..., a_k, b_1, ..., b_n)^{\mathcal{N}}.$$

This implies

$$\inf_{s_1, y_1} ... \inf_{s_n, y_n} \varphi(a_1, ..., a_k, y_1, ..., y_n)^{\mathcal{M}} \leq \inf_{s_1, y_1} ... \inf_{s_n, y_n} \varphi(a_1, ..., a_k, y_1, ..., y_n)^{\mathcal{N}}.$$

Also, the fact that $M \subseteq N$ implies

$$\inf_{s_1, y_1} ... \inf_{s_n, y_n} \varphi(a_1, ..., a_k, y_1, ..., y_n)^{\mathcal{M}} \geq \inf_{s_1, y_1} ... \inf_{s_n, y_n} \varphi(a_1, ..., a_k, y_1, ..., y_n)^{\mathcal{N}}.$$

We conclude that given any inf-formula $\inf_{s_1, y_1} ... \inf_{s_n, y_n} \varphi(x_1, ..., x_k, y_1, ..., y_n)$ and any $a_1, ..., a_k \in N$ we know

$$\inf_{s_1, y_1} ... \inf_{s_n, y_n} \varphi(a_1, ..., a_k, y_1, ..., y_n)^{\mathcal{M}} = \inf_{s_1, y_1} ... \inf_{s_n, y_n} \varphi(a_1, ..., a_k, y_1, ..., y_n)^{\mathcal{N}}.$$

Therefore $\mathcal{M}$ is an existentially closed model of rb$\mathbb{R}$T. □

2.6 Properties of the model companion

In this section various properties of rb$\mathbb{R}$T are presented. The theory rb$\mathbb{R}$T has quantifier elimination and is complete and stable. It is not categorical in any cardinal.

2.6.1 Lemma. *The L_p-theory* rb$\mathbb{R}$T *has quantifier elimination.*

Proof. By Theorem 2.5.4, Theorem 2.2.6 and Proposition 1.7.13. □

2.6.2 Corollary. *The L_p-theory* rb$\mathbb{R}$T *is complete.*

Proof. Observe that we may embed the structure $\{p\}$ consisting of just a basepoint in any model of rb$\mathbb{R}$T. As in first order logic, this fact along with quantifier elimination implies that rb$\mathbb{R}$T is complete. □

2.6.3 Lemma. *Let $M \models \mathbb{R}\mathrm{T}$, and let $b, c \in M$. Let $A \subseteq M$ and define $A' = A \cup \{p\}$. Then $\mathrm{tp}_{\mathcal{M}}(b/A) = \mathrm{tp}_{\mathcal{M}}(c/A)$ if and only if b and c have the same unique closest point $e \in \overline{E_{A'}}$ and $d(b, e) = d(c, e)$.*

Proof. Assume the situation described in the hypotheses. For the forward direction, assume $\mathrm{tp}_{\mathcal{M}}(b/A) = \mathrm{tp}_{\mathcal{M}}(a/A)$. Then we know $d(b, a) = d(c, a)$ for all $a \in \overline{E_{A'}}$, which implies b and c must have the same unique closest point $e \in \overline{E_{A'}}$. Moreover, it is clear that $d(b, e)$ must equal $d(c, e)$. For the other direction, assume b and c have the same unique closest point $e \in \overline{E_{A'}}$ and that $d(b, e) = d(c, e)$. Since rb$\mathbb{R}$T has quantifier elimination, it suffices to show that the quantifier-free types of c and b over A are the same. To show the quantifier-free types over A are the same, it suffices to show $d(a, b) = d(a, c)$ for all $a \in A$ and $d(p, b) = d(p, c)$. This follows easily from our assumptions, since for any $a \in A'$, the point e must be on both $[a, b]$ and $[a, c]$. □

2.6.4 Theorem. *The theory* rb$\mathbb{R}$T *is stable. Indeed when κ is an infinite cardinal,* rb$\mathbb{R}$T *is κ-stable if and only if κ satisfies $\kappa^\omega = \kappa$.*

Proof. Let κ be an infinite cardinal. Let $\mathcal{M} \models$ rb$\mathbb{R}$T be κ^+-saturated, with underlying $\mathbb{R}$-tree (M, d, p). First, assume $\kappa = \kappa^\omega$. Let $|A| = \kappa$. Then

$$|E_A| \leq |A \times A| 2^\omega = \kappa^2 2^\omega \leq \kappa^\omega 2^\omega = \kappa^\omega = \kappa.$$

Counting possible 1-types using Lemma 2.6.3 shows

$$|S_1(A)| \leq |\overline{E_A} \times \mathbb{R}^{\geq 0}| = |\overline{E_A}|2^\omega \leq |E_A|^\omega 2^\omega = \kappa^\omega 2^\omega = \kappa^\omega = \kappa.$$

Thus the rb$\mathbb{R}$T is κ-stable.

For the other direction, note that given any infinite κ we may construct, via a tree construction, a subset A of M with $|A| = \kappa$ and $|\overline{E_A}| = \kappa^\omega$. Now assume that $\kappa^\omega > \kappa$, and let A be such a subset. For each $e \in \overline{E_A}$, choose b_e on a branch out of e that intersects $\overline{E_A}$ only at e with $d(e, b_e) = 1$. We may always find such a branch provided our model is saturated enough. The set $\{b_e \mid e \in \overline{E_A}\}$ has cardinality κ^ω, and for any $e \neq f$ in $\overline{E_A}$ it is straightforward to show $d(\text{tp}(b_e/A), \text{tp}(b_f/A)) \geq 2$. Since $\kappa^\omega > \kappa$, this implies the theory is not κ-stable. □

Let κ be a cardinal so that $\kappa = \kappa^\omega$ and $\kappa > 2^\omega$. Let U be a κ-universal domain for $rb\mathbb{R}$T. A subset of U is *small* if its cardinality is $< \kappa$.

2.6.5 Definition. Let A, B and C be small subsets of U. Say A is *-independent from B over C, denoted $A \mathop{\smile\!\!\!\!\!\mid}^*_C B$, if and only if for all $a \in A$ we have $\text{dist}(a, \overline{E_{B \cup C}}) = \text{dist}(a, \overline{E_C})$.

2.6.6 Lemma. *$A \mathop{\smile\!\!\!\!\!\mid}^*_C B$ if and only if for all $a \in A$ the closest point to a in $\overline{E_{B \cup C}}$ is the same as the closest point to a in $\overline{E_C}$.*

Proof. Assume $A \mathop{\smile\!\!\!\!\!\mid}^*_C B$. Take an arbitrary $a \in A$. Let e_1 be the unique closest point to a in $\overline{E_{B \cup C}}$ and e_2 the unique closest point to a in $\overline{E_C}$. We assumed $\text{dist}(a, \overline{E_{B \cup C}}) = \text{dist}(a, \overline{E_C})$, which implies $d(a, e_1) = d(a, e_2)$. Since $e_2 \in \overline{E_C} \subseteq \overline{E_{B \cup C}}$, we know $e_1 \in [a, e_2]$ by Lemma 2.1.1. Therefore, $e_1 = e_2$. Since a was arbitrary, we know this holds for all $a \in A$. For the other direction, assume for all $a \in A$ the closest point to a in $\overline{E_{B \cup C}}$ is the closest point to a in $\overline{E_C}$. Then, clearly $\text{dist}(a, \overline{E_{B \cup C}}) = \text{dist}(a, \overline{E_C})$ for all $a \in A$. □

2.6.7 Theorem. *The $\mathop{\smile\!\!\!\!\!\mid}^*$ independence relation is the model theoretic independence relation for $rb\mathbb{R}$T.*

Proof. Note, in what follows, we will abbreviate unions such as $B \cup C$ as BC. We will show $\mathop{\smile\!\!\!\!\!\mid}^*$ satisfies all the properties of a stable independence relation on a universal domain of a stable theory as given in [1, Theorem 14.12]. Then, by [1, Theorem 14.14], we know $\mathop{\smile\!\!\!\!\!\mid}^*$ is the model theoretic independence relation for the stable theory rb$\mathbb{R}$T.

(1) Invariance under automorphisms

Any automorphism σ satisfies $\sigma(\overline{E_A}) = \overline{E_{\sigma(A)}}$ and is distance preserving.

(2) Symmetry: if $A \mathop{\smile\!\!\!\!\!\!|}^{*}_{C} B$, then $B \mathop{\smile\!\!\!\!\!\!|}^{*}_{C} A$.

Assume $A \mathop{\smile\!\!\!\!\!\!|}^{*}_{C} B$. This means for all $a \in A$ we have that the closest point in $\overline{E_{BC}}$ to a is $e_a \in \overline{E_C}$. Thus, by Lemma 2.1.2, for any $a \in A$, for any $y \in \overline{E_{BC}}$ we have $[a, y] \cap \overline{E_C} \neq \emptyset$. It follows that for any $x \in \overline{E_A}$, for any $y \in \overline{E_{BC}}$ there exists a point of $\overline{E_C}$ on $[x, y]$. Let $b \in B$. Then for any $x \in \overline{E_A}$ there is a point of $\overline{E_C}$ on $[x, b]$. It follows that the closest point in $\overline{E_A}$ to any $b \in B$ is in $\overline{E_C}$.

(3) Transitivity: $A \mathop{\smile\!\!\!\!\!\!|}^{*}_{C} BD$ if and only if $A \mathop{\smile\!\!\!\!\!\!|}^{*}_{C} B$ and $A \mathop{\smile\!\!\!\!\!\!|}^{*}_{BC} D$.

We know

$$E_C \subseteq E_{BC} \subseteq E_{BCD}$$

which implies

$$\operatorname{dist}(a, \overline{E_C}) \geq \operatorname{dist}(a, \overline{E_{BC}}) \geq \operatorname{dist}(a, \overline{E_{BCD}}).$$

Therefore $\operatorname{dist}(a, \overline{E_{BCD}}) = \operatorname{dist}(a, \overline{E_C})$ if and only if

$$\operatorname{dist}(a, \overline{E_{BC}}) = \operatorname{dist}(a, \overline{E_C}) \text{ and } \operatorname{dist}(a, \overline{E_{BCD}}) = \operatorname{dist}(a, \overline{E_{BC}}).$$

Hence

$$A \mathop{\smile\!\!\!\!\!\!|}^{*}_{C} BD \text{ if and only if } A \mathop{\smile\!\!\!\!\!\!|}^{*}_{C} B \text{ and } A \mathop{\smile\!\!\!\!\!\!|}^{*}_{BC} D.$$

(4) Finite character: $A \mathop{\smile\!\!\!\!\!\!|}^{*}_{C} B$ if and only if $a \mathop{\smile\!\!\!\!\!\!|}^{*}_{C} B$ for all finite tuples $a \in A$.

This is clear from the definition.

(5) Extension: for all A, B, C we can find A' such that $\operatorname{tp}(A/C) = \operatorname{tp}(A'/C)$ and $A' \mathop{\smile\!\!\!\!\!\!|}^{*}_{C} B$.

By finite character and compactness, it suffices to show this statement when A is a finite tuple. Let $e \in \overline{E_C}$ be the unique point closest to $\overline{E_A} = E_A$. Let $\beta < \kappa$ be the cardinality of $\overline{E_B}$. Then there are at most β branches in $\overline{E_B}$ at any point of $\overline{E_B}$. Using Lemma 2.3.7, we may modify the proof of Lemma 2.5.3 so that we avoid any given collection of branches of size $< \kappa$. Since A is finite, we may use this modified Lemma 2.5.3 to embed a copy of $\overline{E_A} = E_A$ on branches at e that do not intersect $\overline{E_B}$ except at e. The image of A under this embedding gives us A'.

(6) Local Character: if $a = a_1, ..., a_m$ is a finite tuple, there is a countable $B_0 \subseteq B$ such that $a \mathop{\smile\!\!\!\!\!\!|}^{*}_{B_0} B$.

Let e_i be the closest point of $\overline{E_B}$ to a_i for $i = 1, ..., m$. Let B_i be a countable subset of B such that e_i is an element of $\overline{E_{B_i}}$. Let $B_0 = \bigcup_i^m B_i$.

(7) Stationarity: if $\mathrm{tp}(A/M) = \mathrm{tp}(A'/M)$, $A \mathop{\smile\!\!\!\!\!\!\mid}^*_M B$, and $A' \mathop{\smile\!\!\!\!\!\!\mid}^*_M B$, then $\mathrm{tp}(A/BM) = \mathrm{tp}(A'/BM)$, where M is a small submodel of U.

By quantifier elimination, $\mathrm{tp}(A/BM)$ is determined by $\{\mathrm{tp}(a/BM) \mid a \in A\}$ plus the information $\{d(a_1, a_2) \mid a_1, a_2 \in A\}$. These distances $\{d(a_1, a_2) \mid a_1, a_2 \in A\}$ are fixed by $\mathrm{tp}(A/M)$. Thus, it suffices to show the conclusion in the case when $A = \{a\}$ and $A' = \{a'\}$. If a or a' is in M the conclusion is obvious, so assume $a, a' \notin M$. The type of a (or a') over BM is determined by two parameters, the unique point in $\overline{E_{BM}}$ that is closest to a, and the distance from a to that point. Since $a \mathop{\smile\!\!\!\!\!\!\mid}^*_M B$, it follows that the closest point in $M = \overline{E_M}$ to a is the same as the closest point in $\overline{E_{BM}}$ to a, and the same is true for a'. Since $\mathrm{tp}(a/M) = \mathrm{tp}(a'/M)$, we know a and a' have the same closest point e in M and $d(a, e) = d(a', e)$. Since e is also the closest point in $\overline{E_{BM}}$ to a and a', we know that $\mathrm{tp}(a/BM) = \mathrm{tp}(a'/BM)$ by Lemma 2.6.3. □

2.6.8 Remark. The proof of stationarity above does not use the fact that M is a model, it works for any set M.

2.6.9 Notation. Now that we know $\mathop{\smile\!\!\!\!\!\!\mid}^*$ is model theoretic independence for rb$\mathbb{R}$T, we will simply denote it by $\mathop{\smile\!\!\!\!\!\!\mid}$.

2.6.10 Lemma. *Let $\mathcal{M} \models$ rb$\mathbb{R}$T with underlying $\mathbb{R}$-tree (M, d, p). Let α be an infinite cardinal and assume there exists $a \in M$ with degree α. Then the density character of M is at least α.*

Proof. Assume the situation described in the hypotheses. Note that the density character of M must be at least ω. Every branch at a has infinite extent, so on each branch out of a we may find b so that $d(a, b) = 1$. Let $(b_i \mid i < \alpha)$ be the collection of all these points. Note that for $i \neq j$, $a \in [b_i, b_j]$ and therefore $d(b_i, b_j) = 2$. For each $i < \alpha$, let O_i be an open ball of radius $\frac{1}{2}$ centered at b_i. Then O_i is contained in the same branch at a as b, because $a \notin [b, x]$ for every point $x \in O_a$. Therefore, these O_i form a collection of disjoint open sets in M with cardinality α. Thus, the density character of M must be at least α. □

2.6.11 Theorem. *The theory* rb$\mathbb{R}$T *is* not *ω-categorical.*

Proof. Any isomorphism g between models $\mathcal{M}$ and $\mathcal{N}$ of rb$\mathbb{R}$T will be a homeomorphism on the underlying $\mathbb{R}$-trees, and therefore g must preserve branching.

Using Lemma 2.2.8, build a separable model $\mathcal{M}$ of rb$\mathbb{R}$T so that each branch point in the underlying $\mathbb{R}$-tree M has degree three. Let N_0 be the $\mathbb{R}$-tree $\mathbb{R}$ with basepoint 0. Let $\mathcal{N}_0$ be the corresponding L_p-structure, which is a model of $\mathbb{R}$T. Let $A_0 = \{a_i^0 \mid i \in \omega\}$ be a countable dense subset of N_0. For each i, let the $\mathbb{R}$-tree T_i be $\mathbb{R}^{\geq 0}$ with basepoint 0, and let $\mathcal{T}_i$ be the corresponding model of $\mathbb{R}$T. Note that T_i is separable. Define $\mathcal{N}_1$ to be the extension of $\mathcal{N}_0$ that results from the application of Lemma 2.2.8 and let (N_1, d, p) be its underlying $\mathbb{R}$-tree. There are now three branches at each point in $A_0 \subseteq N_1$. Moreover, because N_1 is equal to the countable union $\left(\bigcup_{i\in\omega} T_i\right) \cup N_0$, and each T_i is separable, we know that N_1 is separable. The points on the rays we added (with the exception of the basepoints) all still have only two branches. So, we must iterate this construction.

Let $A_1 = \{a_i^1 \mid i \in \omega\}$ be a countable, dense subset of $N_1 \setminus N_0$ and apply Lemma 2.2.8, to $\mathcal{N}_1$ to add the rays T_i to each point in A_1. Call the resulting extension $\mathcal{N}_2$. In the underlying $\mathbb{R}$-tree N_2 of $\mathcal{N}_2$ there are exactly three branches out of each point in $A_1 \cup A_0$, which is a countable, dense subset of N_1. Note that $\mathcal{N}_2$ is still separable, since it is a countable union of separable spaces. Generally, for $n \geq 1$, let $A_n = \{a_i^n \mid i \in \omega\}$ be a countable dense subset of $N_n \setminus N_{n-1}$ and apply Lemma 2.2.8 to get an extension $\mathcal{N}_{n+1}$ of $\mathcal{N}_n$ such that in the $\mathbb{R}$-tree N_{n+1}, all the points in $\bigcup_{i=0}^n A_i$ have exactly three branches and all the points in $N_{n+1} \setminus \bigcup_{i=0}^n A_i$ have exactly two branches. Define $\mathcal{N} = \bigcup_{n\in\omega} \mathcal{N}_n$. This will be a model of $\mathbb{R}$T such that there are exactly three branches at every point in $\bigcup_{i=0}^\infty A_i$, which is a countable, dense subset of the underlying space N of $\mathcal{N}$. Therefore, $\mathcal{N}$ is separable.

Alternatively, we may modify the construction above to get a separable model $\mathcal{M}$ of rb$\mathbb{R}$T so that every branch point in its underlying $\mathbb{R}$-tree M has degree four. Simply let each T_i be a copy of $\mathbb{R}$ with basepoint 0. The separable models $\mathcal{M}$ and $\mathcal{N}$ are not homeomorphic, and therefore cannot be isomorphic. □

2.6.12 Theorem. *Let $\kappa > \omega$ be a cardinal. The theory* rb$\mathbb{R}$T *is* not *κ-categorical.*

Proof. Let $\kappa > \omega$ be a cardinal. Using Lemma 2.2.8, we could construct a model $\mathcal{M}$ so that the set of branch points in its underlying $\mathbb{R}$-tree M is dense and of size κ, and there are κ-many branches at each branch point.

Alternatively, we could construct a model $\mathcal{N}$ so that in the underlying $\mathbb{R}$-tree N of $\mathcal{N}$ the set of branch points is dense and of size κ, there are κ branches at $p^{\mathcal{N}}$, but every other branch point has degree three. By Lemma 2.6.10 both $\mathcal{M}$ and $\mathcal{N}$ have density character at least κ. It follows from the construction using Lemma 2.2.8, $\mathcal{M}$ and $\mathcal{N}$ both have density character at most κ. Clearly $\mathcal{M}$ and $\mathcal{N}$ are not homeomorphic. Thus, they cannot be isomorphic. □

CHAPTER 3

HYPERBOLIC ISOMETRIES OF $\mathbb{R}$-TREES

3.1 Introduction to isometries of $\mathbb{R}$-trees

In this chapter we study classes of $\mathbb{R}$-trees equipped with isometries. This section presents some necessary background (parts of which will not be used until Chapter 4). Given an $\mathbb{R}$-tree M and a geodesic segment $[a, b] \subseteq M$, we will use the notation $(a, b]$ to mean $[a, b] \setminus \{a\}$, and likewise for $[a, b)$.

By an isometry of an $\mathbb{R}$-tree M, we mean a surjective distance preserving function from $M \to M$. Isometries of $\mathbb{R}$-trees fall into two categories. If an isometry f of an $\mathbb{R}$-tree M has a fixed point it is called *elliptic*, otherwise it is *hyperbolic*. The quantity $||f|| := \inf_{x \in M} d(x, f(x))$ is called the *translation distance* of f.

3.1.1 Lemma. *(See [7, 1.3]) Let f be an isometry of an $\mathbb{R}$-tree M. If $||f|| = 0$, then f is elliptic. If $||f|| > 0$, then f is hyperbolic and acts as a translation along an axis, which is a copy of $\mathbb{R}$ in M. The points on this axis are moved by exactly distance $||f||$.*

For a hyperbolic isometry of an $\mathbb{R}$-tree f, let A_f denote the axis of f.

3.1.2 Lemma. *(See [7, 1.3]) Let M be an $\mathbb{R}$-tree and let $f\colon M \to M$ be a hyperbolic isometry of M. Let A_f be the axis of f. Then for all $x \in M$,*

$$\operatorname{dist}(x, A_f) = \frac{d(x, f(x)) - ||f||}{2}.$$

Let $f\colon M \to M$ be an isometry of an $\mathbb{R}$-tree M. For any such f and M, define $f^0(x) = x$. Let $a \in M$. Define the *f-order of a* to be the cardinality of the orbit of a under f. When it is clear what function we are referring to, we will simply call this the *order of a*. For any $m \in \mathbb{N}$, define

$$\operatorname{fix}(f^m) := \{x \in M \mid d(x, f^m(x)) = 0\}.$$

Note that $\mathrm{fix}(f^m)$ a closed subtree of M (see [6, Chapter 3, Lemma 1.1].) Note also that if $a \in \mathrm{fix}(f^m)$, then the f-order of a must divide m. Recall from Section 2.2 that $\mu(x, y)$ denotes the midpoint of $[x, y]$, and that μ is a function uniformly 0-definable in all models of T. Define

$$\mu_m(x) := \mu(x, f^m(x))$$

which is the midpoint of $[x, f^m(x)]$ for $m \in \mathbb{N}^{>0}$. Define $\mu_0(x) = x$.

3.1.3 Lemma. *Let M be an $\mathbb{R}$-tree and let f be an elliptic isometry of M. Let $m \in \mathbb{N}$ and $x \in M$. Then $\mu_m(x)$ is the unique point in $\mathrm{fix}(f^m)$ that is closest to x.*

Proof. This follows directly from [6, Chapter 3, Lemma 1.1], applied to the isometry f^m. □

3.2 The signatures L_s and L_s-structures

In this section we define the continuous signatures L_s we use when studying $\mathbb{R}$-trees with an isometry. For each s, we also give an L_s-theory T_s such that models of T_s have unique underlying spaces and functions up to isomorphism.

Let $s \in \mathbb{N}^{>0}$. Let L_s be the signature L_p from page 16, plus a family of function symbols $(f_n \mid n \in \mathbb{N})$, where the arity of f_n is $(n; n+s)$ and f_n has modulus of uniform continuity $\Delta(\epsilon) = \epsilon$, for each $n \in \mathbb{N}$. Let (M, d, p, f) be such that (M, d, p) is a pointed metric space, $f : M \to M$ is a function that satisfies the modulus of uniform continuity $\Delta(\epsilon) = \epsilon$ and $d(p, f(x)) \leq n + s$ for all $x \in B_n(p)$. We define the *L_s-structure corresponding to* (M, d, p, f) to be the L_p structure corresponding to (M, d, p), together with the interpretation $f_n^{\mathcal{M}} = f \restriction B_n(p) \colon M^{(n)} \to M^{(n+s)}$.

Let T_s be the L_s-theory equal to the theory T_p from Section 1.8 together with the axioms

$$\sup_{m,x} d\big(f_n(I_{m,n}(x)), I_{m+s,n+s}(f_m(x))\big) = 0$$

with $m, n \in \mathbb{N}$ such that $m \leq n$.

3.2.1 Lemma. *Let (M, d, p) be a pointed $\mathbb{R}$-tree, and let $f \colon M \to M$ be a function satisfying $\Delta(\epsilon) = \epsilon$ and $d(p, f(x)) \leq n + s$ for all $x \in B_n(p)$. Then the L_s-structure $\mathcal{M}$ corresponding to (M, d, p, f) is a model of T_s.*

Proof. By Lemma 1.8.8 we know $\mathcal{M} \models T_p$. By definition, $f_n^{\mathcal{M}} = f \restriction B_n(p)$, and $I_{m,n}^{\mathcal{M}}$ is the inclusion map between $M^{(n)} = B_n(p)$ and $M^{(m)} = B_m(p)$ for all $m \leq n \in \mathbb{N}$. By assumption,

$d(p, f(x)) \leq n + s$ for all $x \in B_n(p)$. Therefore

$$\mathcal{M} \models \sup_{m,x} d\big(f_n(I_{m,n}(x)), I_{m+s,n+s}(f_m(x))\big) = 0.$$

□

3.2.2 Corollary. *Let (M, d, p) be a pointed $\mathbb{R}$-tree, and let $f\colon M \to M$ be an isometry such that $d(p, f(p)) \leq s$. Then the L_s-structure $\mathcal{M}$ corresponding to (M, d, p, f) is a model of T_s.*

Proof. Because f is an isometry, it will satisfy the modulus of uniform continuity $\Delta(\epsilon) = \epsilon$. Since $d(p, f(p)) \leq s$, for all $x \in B_n(p)$ we have $d(p, f(x)) \leq d(p, f(p)) + d(f(p), f(x)) = d(p, f(p)) + d(p, x) \leq n + s$. □

3.2.3 Lemma. *Let $\mathcal{M} = \Big((M^{(n)}, d^{(n)}), p_n, I_{m,n}, f_n\Big)$ be an L_s-structure that is a model of T_s. Let (M, d, p) be the underlying metric space of $\mathcal{M}$. Then there exists a unique function $f\colon M \to M$ such that $\mathcal{M}$ is isomorphic to the L_s-structure corresponding to (M, d, p, f).*

Proof. Let $\mathcal{M} = \Big((M^{(n)}, d^{(n)}), p_n, I_{m,n}, f_n\Big) \models T_s$. Let (M, d, p) be the underlying metric space for $\mathcal{M}$. Let $B_n(p)$ denote the closed ball of radius n centered at p in (M, d, p). By Lemma 1.8.9, we may assume $M^{(n)} = B_n(p)$ and $I^{\mathcal{M}}_{m,n}$ is the inclusion of $B_m(p)$ in $B_n(p)$.

For all $m \leq n$

$$\mathcal{M} \models \sup_{m,x} d(f_n(I_{m,n}(x)), I_{m+s,n+s}(f_m(x))) = 0$$

implies $f^{\mathcal{M}}_n \circ I^{\mathcal{M}}_{m,n}(x) = I^{\mathcal{M}}_{m+s,n+s} \circ f^{\mathcal{M}}_m(x)$ for all $x \in M^{(m)}$. This implies $f^{\mathcal{M}}_m = f^{\mathcal{M}}_n \restriction B_m(p)$. Take the union of the functions $f^{\mathcal{M}}_n$ to get $f\colon M \to M$ such that f restricted to $B_n(p)$ is equal to $f^{\mathcal{M}}_n$ for all $n \in \mathbb{N}$. □

For $\mathcal{M} \models T_s$, we call the unique (M, d, p, f) guaranteed by Lemma 3.2.3 the *underlying metric space and function* of $\mathcal{M}$. If the function is an isometry, we call (M, d, p, f) the *underlying metric space and isometry* of $\mathcal{M}$.

3.3 Theories of hyperbolic isometries of $\mathbb{R}$-trees

In this section we axiomatize certain classes of $\mathbb{R}$-trees with a hyperbolic isometry and find model companions for those theories. Let $s \in \mathbb{N}^{>0}$ and L_s as defined on page 51. Let $r \in \mathbb{R}^{>0}$ with $s \geq r$. We will use the abbreviation $||f_n||$ for the L_s-formula $\inf\limits_{n,x} d(x, f_n(x))$.

3.3.1 Definition. Let $\mathcal{K}_{r,s}$ be the class of $\mathcal{M} \models T_s$ with underlying metric space and function (M, d, p, f) so that (M, d, p) is a pointed $\mathbb{R}$-tree, and $f\colon M \to M$ is a hyperbolic isometry such that $d(p, f(p)) \leq s$ and $r \leq ||f||$.

3.3.2 Definition. Let $\mathrm{H}\mathbb{R}\mathrm{T}_{r,s}$ be the L_s-theory consisting of the following L_s-conditions:

1. the axioms of T_s and $\mathbb{R}\mathrm{T}$;

2. for $n \in \mathbb{N}$ the axiom
$$\sup_{n,x} \sup_{n,y} |d(x,y) - d(f(x), f(y))| = 0;$$

3. for $n \in \mathbb{N}$ the axiom
$$\sup_{n+s,y} \min\{n \dot{-} d(f(p), y), \quad \inf_{n,x} d(f(x), y)\} = 0;$$

4. for $n \in \mathbb{N}$ the axiom
$$\sup_{n,x}(r \dot{-} d(x, f(x))) = 0;$$

5. the axiom
$$d(p, f_1(p)) \dot{-} s = 0.$$

3.3.3 Lemma. *The class $\mathcal{K}_{r,s}$ is exactly the class of L_s-structures that are models of* $\mathrm{H}\mathbb{R}\mathrm{T}_{r,s}$.

Proof. Let $\mathcal{M} = \Big((M^{(n)}, d^{(n)}), p_n, I_{m,n}, f_n\Big) \models \mathrm{H}\mathbb{R}\mathrm{T}_{r,s}$ and let (M, d, p, f) be its underlying metric space and function. Recall that $f_n^{\mathcal{M}} = f \upharpoonright B_n(p)$, and $M^{(n)} = B_n(p)$ for each $n \in \mathbb{N}$. We know that (M, d, p) is a pointed $\mathbb{R}$-tree by Theorem 2.2.3. The axioms in item (2) above guarantee that f is an isometry, and the axioms in (3) guarantee that the image of $f_n^{\mathcal{M}}$ is dense in $M^{(n+s)}$ for every $n \in \mathbb{N}$. Then since $M^{(n)}$ and $M^{(n+s)}$ are complete, we conclude that each $f_n^{\mathcal{M}}$ is surjective. Therefore, f is a surjective isometry from M to M. The axioms in (4) guarantee that $r \leq ||f_n||$ for each $n \in \mathbb{N}$, which implies that $r \leq ||f||$. The axiom in (5) implies that $d(p, f(p)) \leq s$.

Let $\mathcal{M} = \Big((M^{(n)}, d^{(n)}), p_n, I_{m,n}, f_n\Big) \in \mathcal{K}_{r,s}$. The axioms in (1) are satisfied in $\mathcal{M}$ by Corollary 3.2.2 and Theorem 2.2.3. The axioms in (2) are true in $\mathcal{M}$ because f is an isometry, and the axioms in (3) are true in $\mathcal{M}$ because f is onto. The axiom in (4) is true in $\mathcal{M}$ because f is hyperbolic with $||f|| \geq r$. That $d^{(m)}(p, f_1(p)) \dot{-} s = 0$ follows directly from $d(p, f(p)) \leq s$, so axiom (5) is true in $\mathcal{M}$. Therefore, $\mathcal{M} \models \mathrm{H}\mathbb{R}\mathrm{T}_{r,s}$. □

For a model $\mathcal{M}$ of $\mathrm{H\mathbb{R}T}_{r,s}$, we refer to the underlying metric space and function (M, d, p, f) of $\mathcal{M}$ as the *underlying $\mathbb{R}$-tree and (hyperbolic) isometry* of $\mathcal{M}$.

3.3.4 Theorem. *The L_s-theory $\mathrm{H\mathbb{R}T}_{r,s}$ has the amalgamation property over substructures of models. That is, if $\mathcal{M}_0$, $\mathcal{M}_1$ and $\mathcal{M}_2$ are substructures of models of $\mathrm{H\mathbb{R}T}_{r,s}$ and $\varphi_1\colon \mathcal{M}_0 \to \mathcal{M}_1$, $\varphi_2\colon \mathcal{M}_0 \to \mathcal{M}_2$ are embeddings, then there exists a model $\mathcal{N}$ of $\mathrm{H\mathbb{R}T}_{r,s}$ and embeddings $g_i\colon \mathcal{M}_i \to \mathcal{N}$ such that $g_1 \circ \varphi_1 = g_2 \circ \varphi_2$.*

Proof. Let (M_0, d_0, p_0, f_0) be the underlying $\mathbb{R}$-tree and isometry of $\mathcal{M}_0$ and likewise for (M_1, d_1, p_1, f_1) and (M_2, d_2, p_2, f_2). As in the proof of Lemma 2.2.6, we may assume that $M_0 \subseteq M_1$, and $M_0 \subseteq M_2$ with $M_0 = M_1 \cap M_2$ and the φ_i are inclusion maps for $i = 1, 2$. In particular, this means $\mathcal{M}_0 \subseteq \mathcal{M}_1$ and the isometry $f_0 = f_1 \restriction M_0$. Likewise, $\mathcal{M}_0 \subseteq \mathcal{M}_2$ and $f_0 = f_2 \restriction M_0$. Therefore, $f_1(x) = f_0(x) = f_2(x)$ for all $x \in M_0$.

Let the L_p-structure $\mathcal{N}' \models \mathbb{R}\mathrm{T}$ with underlying metric space (N, d, p), and the embeddings $g_i\colon \mathcal{M}_i \to \mathcal{N}'$ be the result of applying Lemma 2.2.6 to the reducts of $\mathcal{M}_0$, $\mathcal{M}_1$ and $\mathcal{M}_2$ to L_p. Recall from the proof of Lemma 2.2.6 that $N = M_1 \cup M_2$ where $M_1 \cap M_2 = M_0$, and the embeddings g_i are the inclusion maps of M_i into N for $i = 1, 2$. Also, $p_0 = p_1 = p_2 = p$. Recall that $d(x, y) = d_i(x, y)$ if $x, y \in M_i$, and $d(x, y) = d_i(x, e_y) + d_j(y, e_y)$ for $x \in M_i$ and $y \in M_j$ if $i \neq j$, where e_y is the closest point in M_0 to y. Moreover, d_i is equal to the restriction of d to M_i for $i = 0, 1, 2$.

Define f on N by: $f = f_1$ on $M_1 \subseteq N$, and $f = f_2$ on $M_2 \subseteq N$. This is well defined because f_1 and f_2 are equal on $M_1 \cap M_2 = M_0$. We know $f \circ g_1 = g_1 \circ f_1$ and $f \circ g_2 = g_2 \circ f_2$, because g_i is the inclusion of M_i into N for $i = 1, 2$.

It remains to show that f is an isometry of N. Let $x, y \in N$. If $x, y \in M_i$ for one of $i = 1, 2$, then $d(f(x), f(y)) = d(f_i(x), f_i(y)) = d(x, y)$. If $x \in M_1 \setminus M_0$ and $y \in M_2 \setminus M_0$, then $f(x) \in M_1 \setminus M_0$ and $f(y) \in M_2 \setminus M_0$. Let e_y be the closest point to y in M_0. Then the closest point to $f(y) = f_2(y)$ in $M_0 \subseteq M_2 \subseteq N$ is $f_2(e_y) = f(e_y) = f_1(e_y)$. This implies,

$$
\begin{aligned}
d(f(x), f(y)) &= d(f(x), f(e_y)) + d(f(y), f(e_y)) \\
&= d_1(f_1(x), f_1(e_y)) + d_2(f_2(y), f_2(e_y)) \\
&= d_1(x, e_y) + d_2(y, e_y) \\
&= d(x, e_y) + d(y, e_y).
\end{aligned}
$$

Therefore f is an isometry.

Note that since p and $f(p)$ are both in M_0, we have $d(p, f(p)) = d_0(p_0, f_0(p_0)) \leq s$. Note also that for every $x \in M$, if $x \in M_i$, then $f(x) \in M_i$ for $i = 1, 2$. This implies that $d(x, f(x)) = d_i(x, f_i(x)) \geq r$, and thus, $||f|| \geq r > 0$. So, (N, d, p) is a pointed $\mathbb{R}$-tree and $f \colon N \to N$ is a hyperbolic isometry such that $d(p, f(p)) \leq s$ and $r \leq ||f||$. Therefore, the L_s-structure $\mathcal{N}$ corresponding to (N, d, p, f) is in the class $K_{r,s}$ by Definition 3.3.1. By Lemma 3.3.3, we know that $\mathcal{N} \models \mathrm{H}\mathbb{R}\mathrm{T}_{r,s}$. Since $f \circ g_1 = g_1 \circ f_1$ and $f \circ g_2 = g_2 \circ f_2$, we know the g_i give embeddings $g_1 \colon \mathcal{M}_1 \to \mathcal{N}$ and $g_2 \colon \mathcal{M}_2 \to \mathcal{N}$ of L_s-structures . Lastly, it is clear that $g_1 \circ \varphi_1 = g_2 \circ \varphi_2$, because on the level of the underlying $\mathbb{R}$-trees, these functions are inclusion maps.

In picturing this situation, it may help to recall that M_0 is closed under f_0 and f_0^{-1}, which implies that the axis A_0 of f_0 is contained in M_0. Since $f_0 = f_1 \upharpoonright M_0$, the axis A_1 of f_1 must equal A_0, and likewise the axis A_2 of f_2 must equal A_0. □

3.4 Model companions: hyperbolic case

For this section, let $s \in \mathbb{N}^{>0}$ and let $r \in R^{>0}$ such that $r \leq s$. In this section, we build up to the definition of the theory $\mathrm{rbH}\mathbb{R}\mathrm{T}_{r,s}$, and show it is the model companion of the theory $\mathrm{H}\mathbb{R}\mathrm{T}_{r,s}$.

3.4.1 Definition. Let $\mathcal{M} \models T_s$ with underlying metric space and function (M, d, p, f). Let $a = a_1, ..., a_k \in M$ and $b = b_1, ..., b_n \in M$, and for convenience let $a_1 = p$. For $y_1, ..., y_n$ from the appropriate sorts, define the partial type $D^{\mathcal{M}}_{f,b}(y_1, ..., y_n/a)$ to consist of:

- for $m, l \in \mathbb{N}$ and $i = 1, ..., k$, $j = 1, ..., n$ the condition

$$|d(f^m(a_i), f^l(y_j)) - \big(d(f^m(a_i), f^l(b_j))\big)^{\mathcal{M}}| = 0;$$

- for $m, l \in \mathbb{N}$ and $i, j = 1, ..., n$ the condition

$$|d(f^m(y_i), f^l(y_j)) - \big(d(f^m(b_i), f^l(b_j))\big)^{\mathcal{M}}| = 0.$$

3.4.2 Lemma. *Let $\mathcal{M} \models \mathrm{H}\mathbb{R}\mathrm{T}_{r,s}$ with underlying metric space and function (M, d, p, f). Let $a = a_1, ..., a_k \in M$, $b = b_1, ..., b_n \in M$, and $c_1, ..., c_n \in M$, and let $a_1 = p$. If $(c_1, ..., c_n) \models D^{\mathcal{M}}_{f,b}(y_1, ..., y_n/a)$, then for any quantifier free L_s-formula $\varphi(x_1, ..., x_k, y_1, ..., y_n)$*

$$\varphi(a_1, ..., a_k, b_1, ..., b_n)^{\mathcal{M}} = \varphi(a_1, ..., a_k, c_1, ..., c_n)^{\mathcal{M}}.$$

Proof. The proof is analogous to that of Lemma 2.5.2. □

3.4.3 Lemma. *Assume (M, d, p) is a pointed $\mathbb{R}$-tree and let f be a hyperbolic isometry of M. Let $a_1, ..., a_k \in M$ for $k \geq 1$, and assume $p = a_1$. Let A be the set*

$$\bigcup_{m \in \mathbb{Z}} \{f^m(a_i) \mid i = 1, ..., k\}$$

Then,

1. *E_A is closed under f and f^{-1}.*
2. *the axis A_f of f is contained in E_A;*
3. *$E_A = \overline{E_A}$;*
4. *for all $x \in E_A$, the degree of x in E_A (the number of branches at a in E_A) is at most countable;*

Proof. Let M, the points $a_1, ..., a_k$ and the set A be as above.

Proof of (1): If $x \in E_A$, then $x \in [f^m(a_i), f^n(a_j)]$ for some $m, n \in \mathbb{Z}$ and some $i, j \in \{1, ..., k\}$. Then $f(x) \in [f^{m+1}(a_i), f^{n+1}(a_j)]$ and $f^{-1}(x) \in [f^{m-1}(a_i), f^{n-1}(a_j)]$, which are also clearly contained in E_A.

Proof of (2): The segment $[a_1, f(a_1)] \subseteq E_A$ must contain the segment $[e, f(e)]$ where e is the point on the axis A_f closest to a_1. Also, $A_f = \bigcup_{m \in \mathbb{Z}} f^m([e, f(e)])$. It follows by part (1) that $A_f \subseteq E_A$.

Proof of (3): Assume there is a sequence $(x_n)_{n=1}^{\infty}$ of points in E_A converging to a point y in M. For all $i = 1, ..., k$, let e_i be the point on the axis closest to a_i. Then, $f^m(e_i)$ is the closest point on the axis to $f^m(a_i)$ for any $m \in \mathbb{Z}$. So for any $f^m(a_i)$ and any $f^l(a_j)$ with $l, m \in \mathbb{Z}$ and $i, j \in \{1, ..., k\}$ we know

$$[f^m(a_i), f^l(a_j)] = [f^m(a_i), f^m(e_i)] \cup [f^m(e_i), f^n(e_j)] \cup [f^l(a_j), f^l(e_j)].$$

Therefore, $E_A = A_f \cup \bigcup_{i=1,..,k} \bigcup_{m \in \mathbb{Z}} [f^m(a_i), f^m(e_i)]$, since each $[f^m(e_i), f^n(e_j)]$ is contained in A_f.

If $y \in A_f$, then $y \in E_A$. If $y \notin A_f$, then eventually the sequence $(x_n)_{n=1}^{\infty}$ is bounded away from A_f, that is, there exists $\delta > 0$ so that $\text{dist}(x_n, A_f) > \delta$ for all $n \in \mathbb{N}^{>0}$. Therefore, there is a tail of $(x_n)_{n=1}^{\infty}$ contained in $\bigcup_{i=1,..,k} \bigcup_{m \in \mathbb{Z}} [f^m(a_i), f^m(e_i)]$. We may replace the original sequence with this subsequence. Now, because there are only finitely many choices for i, some subsequence of $(x_n)_{n=1}^{\infty}$ must be contained in $\bigcup_{m\in\mathbb{Z}}[f^m(a_i), f^m(e_i)]$ for some fixed $i \in \{1, ..., k\}$. Again, replace the sequence with this subsequence. Now, for $m \neq l \in \mathbb{Z}$ the sets $[f^m(a_i), f^m(e_i)]$ and $[f^l(a_i), f^l(e_i)]$ have distance at least $||f||$ from one another. This implies that eventually, the elements of the sequence $(x_n)_{n=1}^{\infty}$ must be in one of them. Each $[f^m(a_i), f^m(e_i)]$ is closed, and therefore $y \in E_A$.

Proof of (4): This follows from the fact that E_A is generated by a countable set of points. □

3.4.4 Lemma. *Assume $\mathcal{N} \models \mathbb{H}\mathbb{R}\mathrm{T}_{r,s}$. Let (N, d, p, f) be the underlying metric space and hyperbolic isometry of $\mathcal{N}$. Let $\mathcal{T}$ be a non-empty closed subtree of N such that $\mathcal{T}$ is closed under f and f^{-1}. Let $b_1, ..., b_n \in N \setminus \mathcal{T}$. Let B be the set*

$$\bigcup_{m\in\mathbb{Z}} \{f^m(b_i) \mid i = 1, ..., n\}.$$

Define an equivalence relation on B by: $x \sim y$ if and only if x and y have the same unique closest point in $\mathcal{T}$. Let $\{G_h\}_{h=1}^{\infty}$ be the equivalence classes of $\sim$. For convenience and without loss of generality, assume $b_1, ..., b_n \in \bigcup_{h=1}^{n} G_h$. For each h, let e_h be the unique closest point in $\mathcal{T}$ common to all members of G_h. Then,

1. *$f^m(b_i)$ and $f^l(b_i)$ are in* different *equivalence classes for any $m \neq l$ and any $i = 1, ..., n$;*
2. *for all $h \in \mathbb{N}$ the equivalence class G_h has cardinality at most n;*
3. *the set $\{G_h\}_{h=1}^{\infty}$ of equivalence classes of $\sim$ is closed under f and f^{-1};*
4. *for $h \geq n+1$, there is some $j \in \{1, ..., n\}$ and some $m \in \mathbb{N}$ such that $f^m(G_j) = G_h$.*

Proof. Assume the situation described in the hypotheses. For any $m \in \mathbb{Z}$, if $x \in G_h$, then $f^m(e_h)$ is the closest point to $f^m(x)$ in $\mathcal{T}$, because $\mathcal{T}$ is closed under the isometries f and f^{-1}.

Proof of (1): Towards contradiction, assume there exist $l, m \in \mathbb{Z}$ and $i \in \{1, ..., n\}$ so that $f^l(b_i)$ and $f^m(b_i)$ are in the same equivalence class G_h. Then $f^{m-l}(e_h)$ is the closest point to $f^{m-l}(f^l(b_i)) = f^m(b_i)$ in $\mathcal{T}$. So $f^{m-l}(e_h) = e_h$, which contradicts that f is hyperbolic.

Proof of (2): Towards contradiction, assume there is an $h \in \mathbb{N}$ so that G_h has cardinality $\geq n+1$. Each point in G_h is of the form $f^m(b_i)$ for $i = 1, ..., n$ and $m \in \mathbb{Z}$. By the pigeonhole principle, since there are $> n$ distinct points in G_h there must be two that are different images of the same b_i. This contradicts (1).

Proof of (3): Let G_h be an equivalence class and let $x \in G_h$. Let G_j be the equivalence class of $f(x)$. Then, $e_j = f(e_h)$. This implies $G_j = f(G_h)$. Therefore, $f(G_h)$ is also an equivalence class of $\sim$. The argument for the closure of $\{G_h\}_{h=1}^{\infty}$ under f^{-1} is analogous.

Proof of (4): Let $h \geq n+1$. Let c the minimum of $\{i \mid \exists m \in \mathbb{Z},\ f^m(b_i) \in G_h\}$. Let G_j be the equivalence class of b_c. Since $b_1, ..., b_n \in \bigcup_{h=1}^{n} G_h$, we conclude $j \in \{1, ..., n\}$. The fact that $f^m(b_c) \in G_h$ implies that $f^m(e_j) = e_h$ and thus $f^m(G_j) = G_h$.

□

3.4.5 Theorem. *Let $\mathcal{M} \models \mathrm{H}\mathbb{R}\mathrm{T}_{r,s}$. Then $\mathcal{M}$ is an existentially closed model of $\mathrm{H}\mathbb{R}\mathrm{T}_{r,s}$ if and only the underlying $\mathbb{R}$-tree of $\mathcal{M}$ is richly branching.*

Proof. Let $\mathcal{M} \models \mathrm{H}\mathbb{R}\mathrm{T}_{r,s}$ and let (M, d, p, f) be the underlying $\mathbb{R}$-tree and isometry of M. For the forward direction, assume $\mathcal{M}$ is an existentially closed model. We know that (M, d, p) is a pointed $\mathbb{R}$-tree. We will show the reduct of $\mathcal{M}$ to L_p is a model of rb$\mathbb{R}$T, which implies M is richly branching. It suffices to verify that the conditions $\varphi_n = 0$ defined in 2.3.4 are true in $\mathcal{M}$. Let $n \in \mathbb{N}$ and $a \in M^{(n)}$. Then, $\widetilde{\varphi}_n(a)$ is an inf-formula with a parameter from $\mathcal{M}$. Let $\mathcal{N}$ be an extension of $\mathcal{M}$ with at least 3 branches of infinite height at every point in $\bigcup\{f^m(a) \mid m \in \mathbb{Z}\}$. To find such an extension, use Lemma 2.2.8 to add 3 new rays R_1^b,R_2^b, and R_3^b at each $b \in \bigcup\{f^m(a) \mid m \in \mathbb{Z}\}$ so that $R_i^b \cap R_j^b = \{b\}$ for $i \neq j$. This gives us an L_p-structure $\mathcal{N}$ extending the reduct of $\mathcal{M}$ to L_p. Let (N, d, p) be the underlying $\mathbb{R}$-tree of $\mathcal{N}$. Extend f to a hyperbolic isometry of N by defining $f(R_i^b) = R_i^{f(b)}$ isometrically for $i = 1, 2, 3$. Let $\mathcal{N}$ be the L_s-structure corresponding to (N, d, p, f). Because $\mathcal{M}$ is existentially closed, we know

$$\widetilde{\varphi}_n^{\mathcal{M}}(a) = \widetilde{\varphi}_n^{\mathcal{N}}(a).$$

We also know $\widetilde{\varphi}_n^{\mathcal{N}}(a) = 0$ because there are at least 3 branches of infinite height at a. Therefore $\widetilde{\varphi}_n^{\mathcal{M}}(a) = 0$. The number n and $a \in M^{(n)}$ were arbitrary. Thus for all $n \in \mathbb{N}$ we know $\varphi_n = \sup_{n,x} \widetilde{\varphi}(x) = 0$, and therefore the L_p-structure corresponding to the underlying metric space of $\mathcal{M}$ is a model of rb$\mathbb{R}$T. So, M must be a richly branching $\mathbb{R}$-tree.

Now, for the other direction assume that the underlying $\mathbb{R}$-tree (M, d, p) of $\mathcal{M}$ is richly branching. As in the proof of Theorem 2.5.4, we may assume $\mathcal{M}$ is ω_1-saturated without loss of generality. Since $\mathcal{M}$ is ω_1-saturated, its reduct to L_p is also ω_1-saturated. Also as in the proof of Theorem 2.5.4, to show $\mathcal{M}$ is existentially closed, by Lemma 3.4.2 it suffices to show: for any $a_1, ..., a_k \in M$, any extension $\mathcal{N} \models \mathrm{H}\mathbb{R}\mathrm{T}_{r,s}$ with underlying $\mathbb{R}$-tree (N, d, p) and any $b_1, ..., b_n \in$

N, there exist $c_1, ..., c_n \in M$ such that $c_1, ..., c_n$ satisfy the partial type $D^{\mathcal{N}}_{f,b}(x_1, ..., x_n/a)$.

Let $a_1, ..., a_k \in M$, and let $\mathcal{N} \models \mathrm{H}\mathbb{R}\mathrm{T}_{r,s}$ be an extension of $\mathcal{M}$. Let $b_1, ..., b_n \in N$. We may assume that none of the b_i are in M, because if so, we may let $c_i = b_i$ for that i and then add it and its images under f and f^{-1} to our list of parameters. So, assume $b_1, ..., b_n \in N \setminus M$. Let A be the set $\bigcup_{m\in\mathbb{Z}}\{f^m(a_i) \mid i = 1, ..., k\}$. Let $\mathcal{T}=E_A$. By Lemma 3.4.3, $\mathcal{T}$ is a non-empty closed subtree of N which is closed under f and f^{-1}. Define

$$B = \bigcup_{m\in\mathbb{Z}} \{f^m(b_i) \mid i = 1, ..., n\},$$

and define the equivalence relation $\sim$ on B with equivalence classes G_h as in Lemma 3.4.4. For $h \in \mathbb{N}$, let K_h be the subtree of N generated by $G_h \cup \{e_h\}$, with basepoint $p_h = e_h$. Each K_h is finitely generated, since $|G_h| \leq n$ by Lemma 3.4.4, part (2). Every point in K_h has e_h as its closest point in E_A. Also, if $K_h \cap K_l \neq \emptyset$, then $h = l$. If $x \in K_h$ and $y \in K_l$ for $h \neq l$, then $d(x, y) = d(x, e_h) + d(e_h, e_l) + d(y, e_l)$ by Lemma 2.1.3.

Next, we build an isometric embedding

$$g\colon \bigcup_{h=1}^{\infty} K_h \to M.$$

Step 1: By Lemma 3.4.3, any point in E_A has at most countable degree. So, we use Lemma 2.5.3 to find an isometric embedding $g_1\colon K_1 \to M$ sending p_1 to e_1 whose image does not intersect E_A except at e_1. Define the partial function g on $\bigcup_{l\in\mathbb{Z}} f^l(K_1)$ by:

- $g = g_1$ on K_1;
- for all $x \in K_1$, for all $l \in \mathbb{Z}$, define $g(f^l(x)) = f^l(g_1(x)) = f^l(g(x))$. This defines g on $f^l(K_1)$ for every $l \in \mathbb{Z}$.

Clearly g is isometric embedding on K_1, and g fixes e_1. Let $m, l \in \mathbb{Z}$ and $y \in f^m(K_1)$, $z \in f^l(K_1)$. Let $y' = f^{-m}(y) \in K_1$ and $z' = f^{-l}(z) \in K_1$. Then,

$$d(y, z) = d(y, f^m(e_1)) + d(f^m(e_1), f^l(e_1)) + d(f^l(e_1), z) \tag{3.1}$$

$$= d(y', e_1) + d(f^m(e_1), f^l(e_1)) + d(e_1, z') \tag{3.2}$$

$$= d(g(y'), e_1) + d(f^m(e_1), f^l(e_1)) + d(e_1, g(z')) \tag{3.3}$$

$$= d(f^m(g(y')), f^m(e_1)) + d(f^m(e_1), f^l(e_1)) + d(f^l(e_1), f^l(g(z'))) \tag{3.4}$$

$$= d(f^m(g(y')), f^l(g(z'))) \tag{3.5}$$

$$= d(g(f^m(y')), g(f^l(z'))) = d(g(y), g(z)) \tag{3.6}$$

Line (3.1) is true because $f^m(e_1)$ and $f^l(e_1)$ are the closest points in E_A to y and z respectively. Line (3.2) follows because f^{-m} and f^{-l} are isometries. Line (3.3) is true because g is an isometry when restricted to K_1 and line (3.4) follows because f^m and f^l are isometries. The function g fixes e_1, so e_1 is the closest point in E_A to $g(y)$, and therefore $f^m(e_1)$ is the closest point in E_A to $f^m(g(y))$. Thus, (3.5) follows from (3.4). Lastly, (3.6) is true by the definitions of g and the points y' and z'. So g is an isometry on $\bigcup_{l\in\mathbb{Z}} f^l(K_1)$. Also, it follows from the fact that E_A is closed under f and f^{-1} that the image of g intersects E_A only at the points $f^l(e_1)$ for $l \in \mathbb{Z}$. The isometry g commutes with f and f^{-1} by definition.

Step 2: Extend g to $\left(\bigcup_{l\in\mathbb{Z}} f^l(K_1)\right) \bigcup \left(\bigcup_{l\in\mathbb{Z}} f^l(K_2)\right)$. There are two cases.

Case I: If there are any points in K_2 at which g has already been defined, then $K_2 = f^l(K_1)$ for some $l \in \mathbb{Z}$. In this case g is already defined on K_2 and on all of its images under f and f^{-1}.

Case II: If there are no points in K_2 at which g is defined, use Lemma 2.5.3 to find an isometric embedding $g_2 \colon K_2 \to M$ sending p_2 to e_2 whose image does not intersect E_A except at e_2 and whose image does not intersect the image of g we have constructed so far. We extend our definition of g as follows:

- on K_2 we set $g = g_2$;
- for all $x \in K_2$, for all $l \in \mathbb{Z}$, define $g(f^l(x)) = f^l(g_2(x))$.

To check g is still an isometry, let $m, l \in \mathbb{Z}$ and $y \in f^m(K_1)$, $z \in f^l(K_2)$. Let $y' = f^{-m}(y) \in K_1$ and $z' = f^{-l}(z) \in K_2$. Then, $f^m(e_1)$ is the closest point in E_A to y and $f^l(e_2)$ is the closest point in E_A to z. Therefore

$$d(y,z) = d(y, f^m(e_1)) + d(f^m(e_1), f^l(e_2)) + d(f^l(e_2), z) \tag{3.7}$$

$$= d(y', e_1) + d(f^m(e_1), f^l(e_2)) + d(e_2, z') \tag{3.8}$$

$$= d(g(y'), e_1) + d(f^m(e_1), f^l(e_1)) + d(e_2, g(z')) \tag{3.9}$$

$$= d(f^m(g(y')), f^m(e_1)) + d(f^m(e_1), f^l(e_2)) + d(f^l(e_2), f^l(g(z'))) \tag{3.10}$$

$$= d(f^m(g(y')), f^l(g(z'))) \tag{3.11}$$

$$= d(g(f^m(y')), g(f^l(z'))) = d(g(y), g(z)) \tag{3.12}$$

Line (3.7) follows from Lemma 2.1.3, and line (3.8) comes from applying the isometries f^{-m} and f^{-l}. Then line (3.9) is true because g is an isometric embedding on K_2 and $g(e_2) = e_2$. Then f^m and f^l are applied, resulting in (3.10), and Lemma 2.1.3 gives (3.11). The last line comes from the definition of g. Therefore, the function g on $\left(\bigcup_{l\in\mathbb{Z}} f^l(K_1)\right) \bigcup \left(\bigcup_{l\in\mathbb{Z}} f^l(K_2)\right)$ is an isometric embedding which commutes with f and f^{-1}. Also, the image of g intersects E_A only at the points $f^l(e_1)$ or $f^l(e_2)$ for $l \in \mathbb{Z}$.

Step j**:** For $j \geq 3$, when g has been defined on $\bigcup_{i=1}^{j-1} \bigcup_{l\in\mathbb{Z}} f^l(K_i)$, we extend the domain to include $\bigcup_{l\in\mathbb{Z}} f^l(K_j)$. If g has already been defined at any point of K_j, then g is already defined on $\bigcup_{l\in\mathbb{Z}} f^l(K_j)$ as described above in Case I. If not, we find an isometry g_j defined on K_j and extend g as described in Case II. At each step, the new points from $\bigcup_{l\in\mathbb{Z}} f^l(K_j)$ in the domain of g all have different closest points in E_A as those from $\bigcup_{i=1}^{j-1} \bigcup_{l\in\mathbb{Z}} f^l(K_i)$ already in the domain. Using Lemma 2.1.3, we are able to show our extension is an isometric embedding as in Step 2. Therefore, after every step the partial function g is an isometric embedding and g commutes with f. Also, the image of g intersects E_A only at the points $\{f^l(e_h) \mid l \in \{1, ..., j\},\ h \in \{1, ..., n\}\}$.

By part (4) of Lemma 3.4.4, for $h \geq n+1$ there is some $j \in \{1, ..., n\}$ and some $m \in \mathbb{N}$ such that $f^m(K_j) = K_h$, and $f^m(e_j) = e_h$. Therefore, g is fully defined on $\bigcup_{h=1}^{\infty} K_h$ after n steps. Moreover, g is an isometric embedding and g commutes with f. Also, the image of g intersects E_A only at the points

$$\{f^l(e_h) \mid l \in \mathbb{Z}, h \in \{1, ..., n\}\}.$$

Let $c_i = g(b_i)$. Then since f commutes with g for all $x \in \bigcup_{h=1}^{\infty} K_h$ it follows that

$$f^m(c_i) = f^m(g(b_i)) = g(f^m(b_i)).$$

Therefore, for any $m, l \in \mathbb{N}$ and any $i, j \in \{1, ..., n\}$,

$$d(f^m(b_i), f^l(b_j)) = d(g(f^m(b_i)), g(f^l(b_j))) = d(f^m(c_i), f^l(c_j)).$$

In addition, if e_j is the closest point in E_A to $f^m(b_i)$, then for any point $x \in E_A$,

$$\begin{aligned} d(f^m(b_i), x) &= d(f^m(b_i), e_j) + d(e_j, x) && (3.13)\\ &= d(g(f^m(b_i)), g(e_j)) + d(e_j, x) && (3.14)\\ &= d(f^m(c_i), e_j) + d(e_j, x) && (3.15)\\ &= d(f^m(c_i), x). && (3.16) \end{aligned}$$

Lines (3.13) and (3.16) are true by Lemma 2.1.2 because e_j is the closest point in E_A to $f^m(b_i)$ and $f^m(c_i)$. Line (3.14) follows from (3.13) by an application of the isometry g and line (3.15) comes from (3.14) by the definition of c_i and the fact that $g(e_j) = e_j$. Therefore, we have found $c_1, ..., c_n \in M$ that satisfy the partial type $D_b^f(x_1, ..., x_n)$, and we are done.

Note that once we defined the subtrees K_h, the key to building g was the fact that given K_h for any $h = 1, ..., n$, there is an isometric embedding $g_h \colon K_h \to M$ sending p_h to e_h whose image does not intersect E_A except at e_h. Once we know that, the rest follows from Lemma 3.4.4 and from the fact that if $l \neq k$, then K_h and K_l have different closest points in E_A. □

3.4.6 Definition. Let the L_s-theory rbH$\mathbb{R}$T$_{r,s}$ consist of the axioms of H$\mathbb{R}$T$_{r,s}$ together with the axioms of rb$\mathbb{R}$T.

3.4.7 Theorem. *The L_s-theory* rbH$\mathbb{R}$T$_{r,s}$ *is the model companion of* H$\mathbb{R}$T$_{r,s}$.

Proof. Clearly any model of rbH$\mathbb{R}$T$_{r,s}$ is a model of H$\mathbb{R}$T$_{r,s}$. Any model $\mathcal{M}$ of H$\mathbb{R}$T$_{r,s}$ may be extended to a model $\mathcal{N}$ of rbH$\mathbb{R}$T$_{r,s}$ in the following manner. Let $M \models$ H$\mathbb{R}$T$_{r,s}$ with underlying $\mathbb{R}$-tree and isometry (M, d, p, f). Let X be the $\mathbb{R}$-tree we constructed in Lemma 2.2.9 where every point has degree 3, and choose a basepoint $q \in X$. Use Lemma 2.2.8 to glue in a copy of X at each point $a \in M$, so that q gets identified with a. Let (N, d, p) be the resulting extension of (M, d, p). Denote the copy of X at a by X_a. Then $X_a \cap M = \{a\}$, and $X_b \cap X_a = \emptyset$ for $b \neq a$ in M. Extend the isometry f from M to N by sending X_a to $X_{f(a)}$ in the obvious way, using the identity map on (X, q). Let $\mathcal{N}$ be the L_s-structure corresponding to (N, d, p, f). It follows that $\mathcal{N} \models$ rbH$\mathbb{R}$T$_{r,s}$.

It remains to show that rbH$\mathbb{R}$T$_{r,s}$ is model complete. Let $\mathcal{M} \models$ rbH$\mathbb{R}$T$_{r,s}$. Then, $\mathcal{M}$ is an existentially closed model of H$\mathbb{R}$T$_{r,s}$ by Theorem 3.4.5. Let $\mathcal{N} \models$ rbH$\mathbb{R}$T$_{r,s}$ be an extension of $\mathcal{M}$. Clearly $\mathcal{N}$ is also a model of H$\mathbb{R}$T$_{r,s}$. Then, since $\mathcal{M}$ is an existentially closed model of H$\mathbb{R}$T$_{r,s}$

we know for any inf-formula $\inf_{s_1,y_1} ... \inf_{s_n,y_n} \varphi(x_1, ..., x_k, y_1, ..., y_n)$ and any $a_1, ..., a_k \in M$

$$\inf_{s_1,y_1} ... \inf_{s_n,y_n} \varphi(a_1, ..., a_k, y_1, ..., y_n)^{\mathcal{N}} = \inf_{s_1,y_1} ... \inf_{s_n,y_n} \varphi(a_1, ..., a_k, y_1, ..., y_n)^{\mathcal{M}}.$$

Therefore, $\mathcal{M}$ is an existentially closed model of $\mathrm{rbH}\mathbb{R}\mathrm{T}_{r,s}$. Thus, any model $\mathcal{M}$ of $\mathrm{rbH}\mathbb{R}\mathrm{T}_{r,s}$ is an existentially closed model of $\mathrm{rbH}\mathbb{R}\mathrm{T}_{r,s}$, and therefore $\mathrm{rbH}\mathbb{R}\mathrm{T}_{r,s}$ is model complete. □

3.5 Properties of the model companions

For this section, fix $s \in \mathbb{N}^{>0}$ and $r \in \mathbb{R}$ such that $r \leq s$. In this section, we show that $\mathrm{rbH}\mathbb{R}\mathrm{T}_{r,s}$ has quantifier elimination, characterize the completions of $\mathrm{rbH}\mathbb{R}\mathrm{T}_{r,s}$, and discuss stability.

3.5.1 Lemma. *The L_s-theory $\mathrm{rbH}\mathbb{R}\mathrm{T}_{r,s}$ has quantifier elimination.*

Proof. By Theorem 3.4.7, Theorem 3.3.4 and Proposition 1.7.13. □

3.5.2 Theorem. *Let $t \in \mathbb{R}^{\geq 0}$ such that $r \leq t \leq s$, and let $q \in \mathbb{R}^{>0}$ such that $r \leq q \leq t$. Let the L_s-theory $\mathrm{rbH}\mathbb{R}\mathrm{T}^{q,t}_{r,s}$ be the theory $\mathrm{rbH}\mathbb{R}\mathrm{T}_{r,s}$, together with the axioms*

$$\big|||f_s|| - q\big| = 0$$

and

$$|d(p, f_s(p)) - t| = 0.$$

The theory $\mathrm{rbH}\mathbb{R}\mathrm{T}^{q,t}_{r,s}$ is a completion of $\mathrm{rbH}\mathbb{R}\mathrm{T}_{r,s}$.

Proof. Since $\mathrm{rbH}\mathbb{R}\mathrm{T}_{r,s}$ has quantifier elimination by Lemma 3.5.1, $\mathrm{rbH}\mathbb{R}\mathrm{T}^{q,t}_{r,s}$ has quantifier elimination by [1, Remark 13.4]. Let $\mathcal{M} \models \mathrm{rbH}\mathbb{R}\mathrm{T}^{q,t}_{r,s}$ with underlying $\mathbb{R}$-tree and isometry $(M, d_M, p^{\mathcal{M}}, f^{\mathcal{M}})$ and $\mathcal{N} \models \mathrm{rbH}\mathbb{R}\mathrm{T}^{q,t}_{r,s}$ with underlying $\mathbb{R}$-tree and isometry $(N, d_N, p^{\mathcal{N}}, f^{\mathcal{N}})$. The first new axiom guarantees that

$$||f^{\mathcal{M}}|| = q = ||f^{\mathcal{N}}||$$

because $[p^{\mathcal{M}}, f^{\mathcal{M}}(p^{\mathcal{M}})] \subseteq M^{(s)}$ must intersect the axis of $f^{\mathcal{M}}$, so $||f_s||^{\mathcal{M}} = ||f^{\mathcal{M}}||$, and likewise for $f^{\mathcal{N}}$. The second new axiom guarantees that

$$d_M(p^{\mathcal{M}}, f^{\mathcal{M}}(p^{\mathcal{M}})) = t = d_N(p^{\mathcal{N}}, f^{\mathcal{N}}(p^{\mathcal{N}})).$$

Then, the distance $\operatorname{dist}(p^{\mathcal{M}}, A_{f^{\mathcal{M}}})$ of $p^{\mathcal{M}}$ to the axis $A_{f^{\mathcal{M}}}$ of $f^{\mathcal{M}}$ is determined, since

$$\operatorname{dist}(p^{\mathcal{M}}, A_{f^{\mathcal{M}}}) = \frac{1}{2}\Big(d_M(p^{\mathcal{M}}, f^{\mathcal{M}}(p^{\mathcal{M}})) - ||f^{\mathcal{M}}||\Big).$$

Likewise, the distance $\operatorname{dist}(p^{\mathcal{N}}, A_{f^{\mathcal{N}}})$ of $p^{\mathcal{N}}$ to the axis of $f^{\mathcal{N}}$ is determined and must be equal to $\operatorname{dist}(p^{\mathcal{M}}, A_{f^{\mathcal{M}}})$. Then, for all $m \in \mathbb{Z}$

$$\begin{aligned} d_M(p^{\mathcal{M}}, (f^{\mathcal{M}})^m(p^{\mathcal{M}})) &= 2\operatorname{dist}(p^{\mathcal{M}}, A_{f^{\mathcal{M}}}) + |m| \cdot ||f^{\mathcal{M}}|| \\ &= 2\operatorname{dist}(p^{\mathcal{N}}, A_{f^{\mathcal{N}}}) + |m| \cdot ||f^{\mathcal{N}}|| \\ &= d_N(p^{\mathcal{N}}, (f^{\mathcal{N}})^m(p^{\mathcal{N}})). \end{aligned}$$

This implies that there is an isometry from the $f^{\mathcal{M}}$-orbit of $p^{\mathcal{M}}$ in M to the $f^{\mathcal{N}}$-orbit of $p^{\mathcal{N}}$ in N taking $p^{\mathcal{M}}$ to $p^{\mathcal{N}}$. Therefore, the substructure of $\mathcal{M}$ generated by $p^{\mathcal{M}}$ embeds in any model of $\text{rbH}\mathbb{R}\text{T}_{r,s}^{q,t}$. Then, because $\text{rbH}\mathbb{R}\text{T}_{r,s}^{q,t}$ has quantifier elimination, we know $\text{rbH}\mathbb{R}\text{T}_{r,s}^{q,t}$ is complete. □

It is not hard to see that any completion of $\text{rbH}\mathbb{R}\text{T}_{r,s}$ has an axiomatization of the form $\text{rbH}\mathbb{R}\text{T}_{r,s}^{q,t}$. The completions of $\text{rbH}\mathbb{R}\text{T}_{r,s}$ are exactly the L_s-theories $\operatorname{Th}(\mathcal{M})$ for $\mathcal{M} \models \text{rbH}\mathbb{R}\text{T}_{r,s}$. Given $\mathcal{M} \models \text{rbH}\mathbb{R}\text{T}_{r,s}$, choose $q = ||f^{\mathcal{M}}||$ and $t = d(p^{\mathcal{M}}, f^{\mathcal{M}}(p))$. Then, $\mathcal{M} \models \text{rbH}\mathbb{R}\text{T}_{r,s}^{q,t}$ and by Theorem 3.5.2 $\text{rbH}\mathbb{R}\text{T}_{r,s}^{q,t} = \operatorname{Th}(\mathcal{M})$. For the rest of this section, fix $t \in \mathbb{R}^{\geq 0}$ such that $r \leq t \leq s$, and $q \in \mathbb{R}^{>0}$ such that $r \leq q \leq t$.

Next, we turn to the question of the stability for each completion of $\text{rbH}\mathbb{R}\text{T}_{r,s}$. Given $\mathcal{M} \models \text{rbH}\mathbb{R}\text{T}_{r,s}^{q,t}$ and $A \subseteq M$, let $\mathcal{T}_A$ be the closed subtree of $\mathcal{M}$ generated by the set $\{f^m(a) \mid a \in A \cup \{p\},\ m \in \mathbb{Z}\}$.

3.5.3 Lemma. *Assume $\mathcal{M} \models \text{rbH}\mathbb{R}\text{T}_{r,s}$, and let $b, c \in M$ and $A \subseteq M$. Then $\operatorname{tp}_{\mathcal{M}}(b/A) = \operatorname{tp}_{\mathcal{M}}(c/A)$ if and only if b and c have the same unique closest point $e \in \mathcal{T}_A$ and $d(b,e) = d(c,e)$.*

Proof. Assume the situation described in the hypotheses. The forward direction is the same as in Lemma 2.6.3. For the other direction, assume b and c have the same unique closest point $e \in \mathcal{T}_A$ and that $d(b,e) = d(c,e)$. Since $\text{rbH}\mathbb{R}\text{T}_{r,s}^{q,t}$ has quantifier elimination, it suffices to show that the quantifier-free types of c and b over A are the same. To show the quantifier-free types are the same, by Lemma 3.4.2 it suffices to show $d(a, f^n(b)) = d(a, f^n(c))$ and $d(b, f^n(b)) = d(c, f^n(c))$

for all $n \in \mathbb{N}$, for all $a \in \mathcal{T}_A$. Let $a \in \mathcal{T}_A$ and $n \in \mathbb{N}$. Then,

$$\begin{aligned} d(a, f^n(b)) &= d(a, f^n(e)) + d(f^n(e), f^n(b)) && (3.17)\\ &= d(a, f^n(e)) + d(e, b) && (3.18)\\ &= d(a, f^n(e)) + d(e, c) && (3.19)\\ &= d(a, f^n(e)) + d(f^n(e), f^n(c)) && (3.20)\\ &= d(a, f^n(c)) && (3.21) \end{aligned}$$

for any $n \in \mathbb{N}$. Line (3.17) is true because $f^n(e)$ must be the closest point in $\mathcal{T}_A$ to $f^n(b)$. The next line follows because f^n is an isometry, and line (3.19) follows because $d(e, c) = d(e, b)$. Then, reverse those two steps with c instead of b. Also,

$$\begin{aligned} d(b, f^n(b)) &= d(b, e) + d(e, f^n(e)) + d(f^n(e), f^n(b)) && (3.22)\\ &= 2d(b, e) + d(e, f^n(e)) && (3.23)\\ &= 2d(c, e) + d(e, f^n(e)) && (3.24)\\ &= d(c, e) + d(e, f^n(e)) + d(f^n(e), f^n(c)) && (3.25)\\ &= d(c, f^n(c)) && (3.26) \end{aligned}$$

Line (3.22) is true by Lemma 2.1.3, and line (3.26) follows from (3.25) for the same reason. Therefore, $\text{tp}_{\mathcal{M}}(b/A) = \text{tp}_{\mathcal{M}}(c/A)$. □

3.5.4 Theorem. *The theory* $\text{rbH}\mathbb{R}\text{T}^{q,t}_{r,s}$ *is stable. Indeed when* κ *is an infinite cardinal,* $\text{rbH}\mathbb{R}\text{T}^{q,t}_{r,s}$ *is* κ*-stable if and only if* κ *satisfies* $\kappa^\omega = \kappa$.

Proof. The proof is the same as in Theorem 2.6.4, using Lemma 3.5.3 instead of Lemma 2.6.3. □

Let κ be a cardinal so that $\kappa = \kappa^\omega$ and $\kappa > 2^\omega$. Let U be a κ-universal domain for $\text{rbH}\mathbb{R}\text{T}^{q,t}_{r,s}$.

3.5.5 Definition. Let A, B and C be small subsets of U. Let $\widehat{A} := \{f^m(a) \mid a \in A \cup \{p\},\ m \in \mathbb{Z}\}$, and let $\widehat{B}$ and $\widehat{C}$ be defined analogously. Say A is ^{H}-independent from B over C, denoted $A \mathop{\smile\!\!\!\!\!\mid}^H_C B$, if and only if $\widehat{A} \mathop{\smile\!\!\!\!\!\mid}_{\widehat{C}} \widehat{B}$. That is, $A \mathop{\smile\!\!\!\!\!\mid}^H_C B$ if and only if for all $a \in \widehat{A}$ we have $\text{dist}(a, \overline{E_{\widehat{BC}}}) = \text{dist}(a, \overline{E_{\widehat{C}}})$.

3.5.6 Theorem. *The* $\mathop{\smile\!\!\!\!\!\!\!\!\mid}^{H}$ *independence relation is the model theoretic independence relation for* $\mathrm{rbH\mathbb{R}T}^{q,t}_{r,s}$.

Proof. The general line of argument is close to that of proof of Theorem 2.6.7, but in places we need lemmas from this chapter instead of those used in the proof of 2.6.7. As in that proof we use [1, Theorem 14.12] and show $\mathop{\smile\!\!\!\!\!\!\!\!\mid}^{H}$ has all seven of the properties required by that theorem.

Invariance under automorphisms is clear because any automorphism of a model of $\mathrm{rbH\mathbb{R}T}^{q,t}_{r,s}$ preserves $\mathop{\smile\!\!\!\!\!\!\!\!\mid}$ and satisfies $\sigma(\widehat{A}) = \widehat{\sigma(A)}$. Symmetry and transitivity follow from the definition of $\mathop{\smile\!\!\!\!\!\!\!\!\mid}^{H}$, and from the fact that $\mathop{\smile\!\!\!\!\!\!\!\!\mid}$ is symmetric and transitive. Finite character follows from the fact that $\mathop{\smile\!\!\!\!\!\!\!\!\mid}$ has finite character, and from the fact that $\widehat{A} = \bigcup_{a \in A \cup \{p\}} \{f^m(a) \mid m \in \mathbb{Z}\}$. Extension has much the same proof as in the proof of Theorem 2.6.7, but using Lemma 2.5.3 as in the proof of Theorem 3.4.5. For local character we use the same line of reasoning as in Theorem 2.6.7. The argument for stationarity uses Lemma 3.5.3 instead of Lemma 2.6.3. □

3.5.7 Theorem. *The theory* $\mathrm{rbH\mathbb{R}T}^{q,t}_{r,s}$ *is* not *ω-categorical.*

Proof. Since $L_p \subseteq L_s$ are both countable signatures, and $\mathrm{rb\mathbb{R}T}$ is the reduct of $\mathrm{rbH\mathbb{R}T}^{q,t}_{r,s}$ to L_p, this follows from the fact that $\mathrm{rb\mathbb{R}T}$ is not ω-categorical by [1, Proposition 12.13]. □

3.5.8 Theorem. *Let $\kappa > \omega$ be a cardinal. The theory* $\mathrm{rbH\mathbb{R}T}^{q,t}_{r,s}$ *is* not *κ-categorical.*

Proof. Let $\kappa > \omega$ be a cardinal. We construct non-isomorphic models of $\mathrm{rbH\mathbb{R}T}^{q,t}_{r,s}$, each with density character κ. First, use Lemma 2.2.8 to construct a model $\mathcal{N}$ of $\mathrm{rb\mathbb{R}T}$ with underlying space (N, d, p_N) such that there are κ-many branches at p_N, the set of branch points in N is of size $\leq \kappa$, and at each branch point $x \neq p_N$ there are at most ω-many branches. This N will have density character equal to κ.

Next, let $\mathcal{M} \models \mathrm{rbH\mathbb{R}T}^{q,t}_{r,s}$ be separable with underlying $\mathbb{R}$-tree and isometry (M, d, p, f). The degree of any point in a separable $\mathbb{R}$-tree is at most ω by Lemma 2.6.10. Build an extension W_1 of M by using Lemma 2.2.8 to add a copy of N at each point in $\{f^m(p) \mid m \in \mathbb{Z}\} \subset M$, identifying the basepoint $p_N \in N$ with the point. For each point $f^m(p)$ for $m \in \mathbb{Z}$, label the new copy of N by N_m. The $\mathbb{R}$-tree W_1 has density character at least κ by Lemma 2.6.10. It follows from the construction that W_1 has density character at most κ, therefore the density of W_1 is κ. Extend the isometry f to all of W_1 by defining $f(N_m) = N_{m+1}$ isometrically. Then, the L_s-structure $\mathcal{W}_1$ corresponding to (W_1, d, p, f) is a model of $\mathrm{rbH\mathbb{R}T}^{q,t}_{r,s}$.

To find the other non-isomorphic model, let $\widehat{W_2} \models \mathrm{rbH\mathbb{R}T}^{q,t}_{r,s}$ be κ-saturated with underlying $\mathbb{R}$-tree and isometry $(\widehat{W_2}, d_2, p_2, g)$. Then by Lemma 2.3.7 there are at least κ-many branches at every point in $\widehat{W_2}$. Choose $a \in \widehat{W_2}$ outside the orbit of p_2, and choose κ-many distinct rays at a. Choose κ-many distinct rays at p_2. Let $A_2 \subset \widehat{W_2}$ be the subspace consisting of a, p_2 and the chosen rays. Note that A_2 has density equal to κ. Apply the Downward Löwenheim-Skolem Theorem ([1, Proposition 7.3]) to $\widehat{\mathcal{W}_2}$ to get an elementary substructure $\mathcal{W}_2$ of density character $\leq \kappa$ which contains A_2. Then, $\mathcal{W}_2 \models \mathrm{rbH\mathbb{R}T}^{q,t}_{r,s}$ and has density character equal to κ by Lemma 2.6.10.

In W_1 only the points $\{f^m(p) \mid m \in \mathbb{Z}\}$ have degree κ, while the rest have at most degree ω. In W_2, the basepoint p_2 has degree κ and there is at least one point a outside the orbit of p_2 which has degree κ. Thus, $\mathcal{W}_1$ and $\mathcal{W}_2$ cannot be isomorphic. $\square$

CHAPTER 4

ELLIPTIC ISOMETRIES OF $\mathbb{R}$-TREES

For this chapter, we will need to refer to the background given in Section 3.1 and Section 3.2.

4.1 Theories of elliptic isometries on $\mathbb{R}$-trees

For this section, fix $s \in \mathbb{N}^{>0}$. In this section, we axiomatize classes of $\mathbb{R}$-trees with an elliptic isometry and find model companions for those theories. Let L_s be the signature defined on page 51.

4.1.1 Definition. Let $\mathcal{K}_{0,s}$ be the class of $\mathcal{M} \models T_s$ with underlying metric space and function (M, d, p, f) so that (M, d, p) is a pointed $\mathbb{R}$-tree, and $f\colon M \to M$ is an elliptic isometry such that $d(p, f(p)) \leq s$.

4.1.2 Definition. Let $\mathrm{E\mathbb{R}T}_s$ be the L_s-theory consisting of the following L_s-conditions:

1. the axioms of T_s and $\mathbb{R}\mathrm{T}$;
2. for $n \in \mathbb{N}$ the axiom
$$\sup_{n,x} \sup_{n,y} |d(x, y) - d(f_n(x), f_n(y))| = 0;$$
3. for $n \in \mathbb{N}$ the axiom
$$\sup_{n+s,y} \min\{d(f(p), y) \dot{-} n, \quad \inf_{n,x} d(f(y), x)\} = 0;$$
4. for $n \in \mathbb{N}$ such that $n \geq \frac{s}{2}$, the axiom
$$||f_n|| = 0;$$
5. the axiom $d^{(s)}(p, f_1(p)) \dot{-} s = 0$.

4.1.3 Lemma. *The class $\mathcal{K}_{0,s}$ is exactly the class of L_s-structures that are models of* $\mathrm{E\mathbb{R}T}_s$.

Proof. This proof is the same as the proof of Lemma 3.3.3, except for the axioms in (4). Assume $\mathcal{M} \in \mathcal{K}_{0,s}$. Then $\mu_1(p)$ is the closest fixed point to p. Since $\mu_1(p)$ is the midpoint of $[p, f(p)]$ and $d(p, f(p)) \leq s$, the axioms in (4) are true in $\mathcal{M}$. Conversely, the axioms in (4) imply that there must be a point in M fixed by f by Lemma 3.1.1. □

For a model $\mathcal{M}$ of $\mathrm{E\mathbb{R}T}_s$, we will refer to the underlying metric space and function (M, d, p, f) of $\mathcal{M}$ as the *underlying $\mathbb{R}$-tree and (elliptic) isometry* of $\mathcal{M}$.

4.1.4 Theorem. *The L_s-theory* $\mathrm{E\mathbb{R}T}_s$ *has the amalgamation property over substructures of models.*

Proof. Let (M_0, d_0, p_0, f_0) be the underlying $\mathbb{R}$-tree and isometry of $\mathcal{M}_0$ and likewise for (M_1, d_1, p_1, f_1) and (M_2, d_2, p_2, f_2). As in the proof of Lemma 2.2.6, we may assume that $M_0 \subseteq M_1$, and $M_0 \subseteq M_2$ with $M_0 = M_1 \cap M_2$ and the φ_i are inclusion maps for $i = 1, 2$. In particular, this means $\mathcal{M}_0 \subseteq \mathcal{M}_1$ and the isometry $f_0 = f_1 \restriction M_0$. Likewise, $\mathcal{M}_0 \subseteq \mathcal{M}_2$ and $f_0 = f_2 \restriction M_0$. Therefore, $f_1(x) = f_0(x) = f_2(x)$ for all $x \in M_0$.

Define (N, d, p) and f just as in the proof of Lemma 3.3.4. The isometry f of N is elliptic, since $f \restriction M_1 \subseteq M = f_1$, and f_1 had a fixed point. Therefore, the L_s-structure $\mathcal{N}$ corresponding to (N, d, p, f) is a model of $\mathrm{E\mathbb{R}T}_s$ that has the properties we want.

□

4.2 Orbits under an elliptic isometry

In this section, we provide lemmas which characterize the possible orbits of points in an $\mathbb{R}$-tree with an elliptic isometry.

Given an $\mathbb{R}$-tree M, an elliptic isometry f of M, and $a \in M$, the following lemmas describe the structure of the geodesic segment $[\mu_1(a), a]$ between a and the closest fixed point to a. (When we just say a point is *fixed* we mean fixed by f, that is, of f-order one.) That is, they describe how $[\mu_1(a), a]$ may be split up into subsegments of constant f-order, and how those segments must be arranged. This in turn determines the structure of the f-orbit of $a \in M$, that is, it determines all distances $d(f^k(a), f^l(a))$ for $k < l \in \mathbb{N}$.

4.2.1 Lemma. *Let $f\colon M \to M$ be an elliptic isometry of an $\mathbb{R}$-tree M. Let $a \in M$. Then*

1. *if $m, n \in \mathbb{N}^{>0}$ and $m|n$, then $\mu_n(a)$ is on $[\mu_m(a), a]$;*

2. *if a has finite f-order, $b \in [\mu_1(a), a]$ and b has f-order $m \in \mathbb{N}^{>0}$, then for all $x \in [\mu_1(a), b]$, the f-order of x divides m and for all $y \in [b, a]$, the f-order of y is a multiple of m;*

3. *if $b \in [\mu_1(a), a]$ and b has f-order ∞, then every $x \in [b, a]$ has f-order ∞;*

4. *if $\mu_n(a)$ has f-order n, then anything in $(\mu_n(a), a]$ has f-order ∞ or f-order kn for some $k \in \mathbb{N}$ with $k > 1$.*

Proof. Let $f\colon M \to M$ be an elliptic isometry of an $\mathbb{R}$-tree M. Let $a \in M$.

Proof of (1): Assume $m, n \in \mathbb{N}^{>0}$ and $m|n$. Then, $\mathrm{fix}(f^m) \subseteq \mathrm{fix}(f^n)$, which implies $\mu_m(a) \in \mathrm{fix}(f^n)$. By Lemma 3.1.3, $\mu_n(a)$ is the closest point in $\mathrm{fix}(f^n)$ to a. Therefore by Lemma 2.1.1, $\mu_n(a) \in [\mu_m(a), a]$.

Proof of (2): Assume a has finite f-order, $b \in [\mu_1(a), a]$ and b has f-order $m \in \mathbb{N}^{>0}$. Both $\mu_1(a)$ and b are in $\mathrm{fix}(f^m)$, which implies $[\mu_1(a), b] \subseteq \mathrm{fix}(f^m)$. Therefore, for any $x \in [\mu_1(a), b]$, the f-order of x divides m. Next, let $y \in [b, a]$. Then, $\mu_1(a) = \mu_1(y)$, because otherwise there would be a point of f-order one closer to a than $\mu_1(a)$. This implies $b \in [\mu_1(a), y] = [\mu_1(y), y]$. Then, if n is the f-order of y, we know $b \in [\mu_1(y), y] \subseteq \mathrm{fix}(f^n)$, which implies n is a multiple of m.

Proof of (3): Assume $b \in [\mu_1(a), a]$ and b has f-order ∞. Towards contradiction, assume there exists $x \in [b, a]$ with finite f-order n. Then, $[\mu_1(a), x] \subseteq \mathrm{fix}(f^n)$ which implies that $b \in \mathrm{fix}(f^n)$ and hence the f-order of b divides n. This is our contradiction.

Proof of (4): This follows from (2) and (3). Nothing in $(\mu_n(a), a]$ can be in $\mathrm{fix}(f^n)$ because $\mu_n(a)$ is the closest point in $\mathrm{fix}(f^n)$ to a. Therefore, if the order of the point is kn, we know $k > 1$. ☐

4.2.2 Definition. Let M be an $\mathbb{R}$-tree, $f\colon M \to M$ an elliptic isometry of M, and $a \in M$. Define

$$O(a) := \{t \in \mathbb{N} \mid \exists x \in [\mu_1(a), a] \text{ such that } x \text{ has } f\text{-order } t\}.$$

4.2.3 Lemma. *Let M be an $\mathbb{R}$-tree, $f\colon M \to M$ an elliptic isometry of M, and $a \in M$. Let $1 < t_1 < t_2 < \ldots$ be a list of the members of $O(a)$ in ascending order. Then, $\mu_t(a) \in [\mu_1(a), a]$ for all $t \in O(a)$, and for $j < l$ we have $\mu_{t_j}(a) \in [\mu_1(a), \mu_{t_l}(a)]$ and t_j divides t_l. Moreover, every* *point in* $(\mu_1(a), \mu_{t_1}]$ *has order* t_1*, and for* $j = 1, 2, ...$ *every point in* $(\mu_{t_j}(a), \mu_{t_{j+1}}(a)]$ *has order* t_{j+1}*. Additionally,*

1. *if the* f*-order of* a *is finite and* > 1*, then*
   - $O(a)$ *is finite and the largest element* t_k *of* $O(a)$ *is the order of* a *(which implies* $\mu_{t_k}(a) = a$*);*
   - $[\mu_1(a), a] = [\mu_1(a), \mu_{t_1}(a)] \cup (\mu_{t_1}(a), \mu_{t_2}(a)] \cup ... \cup (\mu_{t_{k-1}}(a), \mu_{t_k}(a)]$*, where this union is disjoint;*
2. *if the order of* a *is infinite, then* either $O(a)$ *is finite with largest element* t_k *and*
   - $[\mu_1(a), a] = [\mu_1(a), \mu_{t_1}(a)] \cup (\mu_{t_1}(a), \mu_{t_2}(a)] \cup ... \cup (\mu_{t_{k-1}}(a), \mu_{t_k}(a)] \cup (\mu_{t_k}(a), a]$ *where this union is disjoint and every point in* $(\mu_{t_k}(a), a]$ *has infinite order;*

   or $O(a)$ *is infinite and*
   - $\{\mu_{t_j}(a)\}_{j=1}^{\infty}$ *converges to a point* $\mu_\infty(a)$ *so that every point in* $[\mu_\infty(a), a]$ *has infinite order;*
   - $[\mu_1(a), a] = [\mu_1(a), \mu_{t_1}] \cup \left(\bigcup_{j=1}^{\infty} (\mu_{t_j}(a), \mu_{t_{j+1}}(a)] \right) \cup [\mu_\infty(a), a]$ *where this union is disjoint.*

Proof. Let M be an $\mathbb{R}$-tree, $f: M \to M$ an elliptic isometry of M, and $a \in M$. Let $1 < t_1 < t_2 < ...$ be a list of the members of $O(a)$ in ascending order. By Lemma 4.2.1 part (1) we have $\mu_t(a) \in [\mu_1(a), a]$ for all $t \in O(a)$. Then, it follows from part (4) of Lemma 4.2.1 that $\mu_1(a) = \mu_1(\mu_t(a))$ for each $t \in O(a)$. If $t \in O(a)$, then $\mu_t(a)$ actually has order t. Then, by the fact that $1 < t_1 < t_2 < ...$ and Lemma 4.2.1 part (2) it follows that for $j < l$ we have $\mu_{t_j}(a) \in [\mu_1(\mu_{t_l}(a)), \mu_{t_l}(a)] = [\mu_1(a), \mu_{t_l}(a)]$ and t_j divides t_l. The $t \in O(a)$ represent all of the possible finite orders of points on $[\mu_1(a), a]$, therefore, every point in $(\mu_1(a), \mu_{t_1}]$ has order t_1, and for all j, every point in $(\mu_{t_j}(a), \mu_{t_{j+1}}(a)]$ has order t_{j+1}.

Proof of (1): Assume the order of a is finite and > 1. By Lemma 4.2.1 part (2) any $t \in O(a)$ must divide the order of a. Then clearly $O(a)$ is finite and its largest element t_k is the order of a. That $[\mu_1(a), a] = [\mu_1(a), \mu_{t_1}(a)] \cup (\mu_{t_1}(a), \mu_{t_2}(a)] \cup ... \cup (\mu_{t_{k-1}}(a), \mu_{t_k}(a)]$ follows from the facts that $\mu_{t_j}(a) \in [\mu_1(a), \mu_{t_l}(a)]$ for $j < l$, and $\mu_{t_k}(a) = a$. The fact that this is a disjoint union is clear from an examination of the orders of the points in each piece.

Proof of (2): Assume the order of a is infinite.

Case I: Assume the set $O(a)$ is finite. Let t_k be the largest element of $O(a)$. Then, every point in $(\mu_{t_k}(a), a]$ must have infinite order. That $[\mu_1(a), \mu_{t_k}(a)] = [\mu_1(a), \mu_{t_1}(a)] \cup (\mu_{t_1}(a), \mu_{t_2}(a)] \cup ... \cup (\mu_{t_{k-1}}(a), \mu_{t_k}(a)]$ is clear by part (1). Then, it follows that $[\mu_1(a), a] = [\mu_1(a), \mu_{t_1}(a)] \cup (\mu_{t_1}(a), \mu_{t_2}(a)] \cup ... \cup (\mu_{t_k}(a), a]$ by the fact that $\mu_{t_j}(a) \in [\mu_1(a), \mu_{t_l}(a)]$ for $j < l$. The fact that this is a disjoint union is clear from an examination of the orders of the points in each piece.

Case II: Assume $O(a)$ is infinite. For $j < l$, we know $d(\mu_1(a), \mu_{t_j}(a)) < d(\mu_1(a), \mu_{t_l}(a))$ because $\mu_{t_j}(a) \in [\mu_1(a), \mu_{t_l}(a)]$. So, $\{d(\mu_1(a), \mu_{t_j}(a))\}_{j=1}^{\infty}$ is a strictly increasing sequence of positive real numbers, bounded above by $d(\mu_1(a), a)$. Therefore this sequence converges to some real number $R \leq d(\mu_1(a), a)$. Let $\mu_\infty(a)$ be the point on $[\mu_1(a), a]$ with distance R from $\mu_1(a)$. Then, $\{\mu_{t_j}(a)\}_{j \in \mathbb{N}}$ must converge to $\mu_\infty(a) \in [\mu_1(a), a]$ (since $[\mu_1(a), a]$ is an isometric copy of a real interval).

Claim: Every point in $[\mu_\infty(a), a]$ must have infinite order.

Proof of claim: Towards contradiction assume some point in $[\mu_\infty(a), a]$ has finite order. Then, $\mu_\infty(a)$ must have finite order, say it has order m. Then, $m \in O(a)$. But, $O(a)$ is an infinite collection of positive integers, and thus there must exist $t_j \in O(a)$ such that $t_j > m$. Then, $d(\mu_1(a), \mu_{t_j}(a)) > d(\mu_1(a), \mu_\infty(a)) = R$. But, by how we found R, we know $d(\mu_1(a), \mu_{t_j}(a)) < R$. This is a contradiction. Thus, every point in $[\mu_\infty(a), a]$ has infinite order. Note that $\mu_\infty(a)$ is the closest point of infinite order to $\mu_1(a)$ on $[\mu_1(a), a]$.

Since $\mu_{t_j}(a) \in [\mu_1(a), a]$ for all $j = 1, ..., k$ and $\mu_\infty(a) \in [\mu_1(a), a]$ it is clear that

$$[\mu_1(a), \mu_{t_1}(a)] \cup \big(\bigcup_{j=1}^{\infty} (\mu_{t_j}(a), \mu_{t_{j+1}}(a)]\big) \cup [\mu_\infty(a), a] \subseteq [\mu_1(a), a].$$

If $z \in [\mu_1(a), a]$, then if z has finite order it will be in

$$[\mu_1(a), \mu_{t_1}(a)] \cup \big(\bigcup_{j=1}^{\infty} (\mu_{t_j}(a), \mu_{t_{j+1}}(a)]\big)$$

If z has infinite order, it will be in $[\mu_\infty(a), a]$, since $\mu_\infty(a)$ is the closest point to $\mu_1(a)$ on $[\mu_1(a), a]$ of infinite order. Thus,

$$[\mu_1(a), a] = [\mu_1(a), \mu_{t_1}] \cup \big(\bigcup_{j=1}^{\infty} (\mu_{t_j}(a), \mu_{t_{j+1}}(a)]\big) \cup [\mu_\infty(a), a].$$

□

4.2.4 Lemma. *Let M be an $\mathbb{R}$-tree, $f\colon M \to M$ an elliptic isometry of M, and $a, b \in M$. If $O(a) = O(b)$ and $d(a, \mu_t(a)) = d(b, \mu_t(b))$ for all $t \in O(a)$, then $d(a, \mu_n(a)) = d(b, \mu_n(b))$ for all $n \in \mathbb{N}$.*

Proof. Let M be an $\mathbb{R}$-tree, $f\colon M \to M$ an elliptic isometry of M, and $a, b \in M$. Assume $O(a) = O(b)$ and $d(a, \mu_t(a)) = d(b, \mu_t(b))$ for all $t \in O(a)$. Let $n \in \mathbb{N}$. By Lemma 4.2.1 part (1) we know $\mu_n(a) \in [\mu_1(a), a]$. Let $t_k \in O(a)$ be the largest number in $O(a)$ that divides n. Then, $\mu_{t_k}(a) = \mu_n(a)$, because $\mu_{t_k}(a)$ will be the point in $\mathrm{fix}(f^n)$ that is closest to a. Since $O(a) = O(b)$, we conclude $\mu_{t_k}(b) = \mu_n(b)$. Now, $d(a, \mu_t(a)) = d(b, \mu_t(b))$ for all $t \in O(a)$ implies $d(a, \mu_n(a)) = d(b, \mu_n(b))$. □

4.2.5 Lemma. *Let M be an $\mathbb{R}$-tree, $f\colon M \to M$ an elliptic isometry of M, and $a, b \in M$. If $d(a, \mu_n(a)) = d(b, \mu_n(b))$ for all $n \in \mathbb{N}$, then for all $k, l \in \mathbb{Z}$*

$$d(f^k(a), f^l(a)) = d(f^k(b), f^l(b)).$$

Proof. Let M be an $\mathbb{R}$-tree, $f\colon M \to M$ an elliptic isometry of M, and $a, b \in M$. Assume $d(a, \mu_n(a)) = d(b, \mu_n(b))$ for all $n \in \mathbb{N}$. Let $k, l \in \mathbb{Z}$ with $k \leq l$. Then, $l - k \in \mathbb{N}$ and

$$d(f^k(a), f^l(a)) = d(a, f^{l-k}(a)) = 2d(a, \mu_{l-k}(a))$$

because f is an isometry and by the definition of μ_{l-k}. By our assumption, $2d(a, \mu_{l-k}(a)) = 2d(b, \mu_{l-k}(b))$. Therefore, $d(f^k(a), f^l(a)) = 2d(b, \mu_{l-k}(b)) = d(b, f^{l-k}(b)) = d(f^k(b), f^l(b))$. □

4.2.6 Remark. In the elliptic case, we need an analog of Lemma 3.4.3. Here, (M, d, p) is a pointed $\mathbb{R}$-tree and f an elliptic isometry of M, $a \in M$, and $A = \{f^m(a) \mid m \in \mathbb{Z}\}$. However, in the elliptic case the subtree $E_A = \bigcup_{m\in\mathbb{Z}} f^m([\mu_1(a), a])$ is *not* necessarily closed.

4.2.7 Lemma. *Assume (M, d, p) is a pointed $\mathbb{R}$-tree and f an elliptic isometry of M. Let $a_1, ..., a_k \in M$ for $k \geq 1$, and assume $p = a_1$. Let $A := \bigcup_{m\in\mathbb{Z}}\{f^m(a_i) \mid i = 1, ..., k\}$. Then,*

1. *$\overline{E_A}$ is closed under f and f^{-1};*
2. *for all $x \in \overline{E_A}$, the degree of x in $\overline{E_A}$ is at most 2^ω.*

Proof. Assume (M, d, p) is a pointed $\mathbb{R}$-tree and f an elliptic isometry of M. Let $a_1, ..., a_k \in M$ for $k \geq 1$, and assume $p = a_1$.

Proof of (1): It is clear that E_A must be closed under f and f^{-1}. If $y \in \overline{E_A} \setminus E_A$, then y is a limit point of a sequence $(x_n)_{n=1}^{\infty}$ of points from E_A. Then, because f is an isometry, $f(y)$ must be the limit of $(f(x_n))_{n=1}^{\infty}$. These points are all in E_A since E_A is closed under f, therefore $f(y) \in \overline{E_A}$. The proof for f^{-1} is the same.

Proof of (2): This follows from the fact that the cardinality of $\overline{E_A}$ is at most 2^{ω}. □

4.3 Model companions: elliptic case

For this section, fix $s \in \mathbb{N}$. In this section, we characterize the model companion of the theory $\mathrm{E\mathbb{R}T}_s$.

Next, we define an L_s-theory $\mathrm{rbE\mathbb{R}T}_s$ extending $\mathrm{E\mathbb{R}T}_s$ which will turn out to be the model companion of $\mathrm{E\mathbb{R}T}_s$. Informally, we want to add the following conditions when (M, d, p, f) is the underlying $\mathbb{R}$-tree and elliptic isometry of a model $\mathcal{M}$ of $\mathrm{E\mathbb{R}T}_s$, and $r \in \mathbb{R}^{>0}$, $h, l, m \in \mathbb{N}^{>0}$:

If $x \in M$ has order m, then there exist $y_1, ..., y_l \in M$ such that for all $i, j = 1, ..., l$ we have $d(x, y_i) = r$, $d(y_i, y_j) = 2r$ (if $i \neq j$), and every element of $(x, y_i]$ has order hm.

In order to express suitable forms of these conditions using L_s-conditions, we need to do some preliminary work.

To begin, we introduce an extension by definitions of $\mathrm{E\mathbb{R}T}_s$ in which the midpoint function $\mu(x_1, x_2)$ is available. Using it we also have at hand the functions $\mu_m(x) = \mu(x, f^m(x))$, in terms of which we can express the orders of elements. Recall from Lemma 2.2.4 that in models of $\mathbb{R}$T the distance $d(y, \mu(x_1, x_2))$ is given by the interpretation of certain L_p-formulas. Let $\mathcal{M} \models \mathbb{R}\mathrm{T}$ and let (M, d, p) be the underlying metric space of $\mathcal{M}$. Specifically, by Lemma 2.2.4, for $x_1, x_2, y \in M^{(n)}$ the distance $d(y, \mu(x_1, x_2))$ is given by the interpretation of the formula

$$\psi_n(x_1, x_2, y) = \max\{d(x_1, y) \dot{-} \frac{d(x_1, x_2)}{2}, d(x_2, y) \dot{-} \frac{d(x_1, x_2)}{2}\}$$

where all the variables are over $M^{(n)}$.

With this background, we introduce an extension by definitions $\mathrm{E\mathbb{R}T}_s(\mu)$ of $\mathrm{E\mathbb{R}T}_s$ in which the midpoint function is available. We denote the extended continuous signature by $L_s(\mu)$. In addition to L_s, it contains for each $n \in \mathbb{N}$ a binary function symbol $\mu^{(n)}$ of arity $(n, n; n)$ whose interpretation will be the restriction of the midpoint function to the sort $M^{(n)}$. The assigned modulus of uniform continuity for $\mu^{(n)}$ is $\Delta(\epsilon) = \frac{\epsilon}{2}$.

The theory $\mathrm{E\mathbb{R}T}_s(\mu)$ consists of $\mathrm{E\mathbb{R}T}_s$ together with the following $L_s(\mu)$-conditions that serve to define $\mu^{(n)}$: for each $n \in \mathbb{N}$

$$\sup_{n,x_1} \sup_{n,x_2} \sup_{n,y} \left| d(\mu^{(n)}(x_1, x_2), y) - \psi_n(x_1, x_2, y) \right| = 0.$$

Using Lemma 2.2.4, it is clear that each model $\mathcal{M} \models \mathrm{E}\mathbb{R}\mathrm{T}_s$ has a unique expansion to a model of $\mathrm{E}\mathbb{R}\mathrm{T}_s(\mu)$. We denote this expansion of $\mathcal{M}$ by $\mathcal{M}(\mu)$. If $\mathcal{M} \models \mathrm{E}\mathbb{R}\mathrm{T}_s$ and (M, d, p, f) is the underlying $\mathbb{R}$-tree and elliptic isometry of $\mathcal{M}$, then for each $n \in \mathbb{N}$, the interpretation of $\mu^{(n)}$ in $\mathcal{M}(\mu)$ is the restriction of the midpoint function of M to the sort $M^{(n)} = B_n(p)$.

For each $m \in \mathbb{N}^{>0}$ and $n \in \mathbb{N}$ we define

$$\mu_m^{(n)}(x) = \mu^{(n+ms)} \left(I_{n,n+ms}(x), f^m(x) \right)$$

where x is of sort n. This is an $L_s(\mu)$-term of arity $(n; n+ms)$. For $\mathcal{M}$ and (M, d, p, f) as above, the interpretation of $\mu_m^{(n)}$ in $\mathcal{M}(\mu)$ is the restriction of the function μ_m (giving the nearest fixed point of f^m) to $M^{(n)}$. As we have done in other similar situations, we will often write μ instead of $\mu^{(n)}$ and μ_m instead of $\mu_m^{(n)}$ in $L_s(\mu)$-formulas, leaving it up to the reader to assign the needed sort indices.

Our next result shows explicitly how certain $L_s(\mu)$-conditions can be simplified by eliminating occurrences of the function symbols μ_m. By applying this result several times, we can explicitly produce equivalent L_s-conditions for the axioms we add to $\mathrm{E}\mathbb{R}\mathrm{T}_s(\mu)$.

4.3.1 Lemma. *Assume $\mathcal{M} \models \mathrm{E}\mathbb{R}\mathrm{T}_s$ with underlying $\mathbb{R}$-tree and isometry (M, d, p, f). Let $\varphi(y, z_1, ..., z_k) = 0$ be an $L_s(\mu)$-condition where y is of sort n. Let $a \in M^{(n)}$, and $b_1, ..., b_k \in \mathcal{M}$ be from appropriate sorts. Let $m \in \mathbb{N}$. Then*

$$\left(\inf_{n+ms,y} \max\{\psi_{n+ms}(a, f^m(a), y), \varphi(y, b_1, ..., b_k)\} \right)^{\mathcal{M}(\mu)} = 0$$

if and only if

$$\varphi^{\mathcal{M}(\mu)}(\mu_m(a), b_1, ..., b_k) = 0.$$

Proof. Assume the hypotheses of the lemma. For the left to right direction, assume

$$\left(\inf_{n+ms,y} \max\{\psi_{n+ms}(a, f^m(a), y), \varphi(y, b_1, ..., b_)\} \right)^{\mathcal{M}(\mu)} = 0.$$

Then, for all $\epsilon > 0$ there exists $c \in M^{(n)}$ so that $\varphi^{\mathcal{M}(\mu)}(c, b_1, ..., b_k) \leq \epsilon$ and $\psi_{n+ms}^{\mathcal{M}(\mu)}(a, f^m(a), c) \leq$

ϵ, hence, $d(c, \mu_m(a)) \leq \epsilon$. Then by the uniform continuity of the function $\varphi^{\mathcal{M}(\mu)}$, we conclude $\varphi^{\mathcal{M}(\mu)}(\mu_m(a), b_1, ..., b_k) = 0$. In the other direction, the $\inf_{n,y}$ is clearly witnessed by $\mu_m(a)$. □

In the $L_s(\mu)$-conditions we will add to $\mathbb{ERT}_s(\mu)$, we need to use the sup quantifier over the set $\text{fix}(f^m) \cap M^{(n)}$. The next results show how to eliminate such quantifiers using formulas of $L_s(\mu)$.

4.3.2 Lemma. *For any $m \in \mathbb{N}$, the sets $\text{fix}(f^m)$ are uniformly definable in models of $\mathbb{ERT}_s$.*

Proof. Let $\mathcal{M} \models \mathbb{ERT}_s$ with underlying $\mathbb{R}$-tree and isometry (M, d, p, f). In M, the distance $\text{dist}(x, \text{fix}(f^m))$ is equal to the distance from x to $\mu_m(x)$ by Lemma 3.1.3. By definition, $\mu_m(x)$ is the midpoint of $[x, f^m(x)]$. Therefore, for $x \in M^{(n)}$ the predicate $\text{dist}(x, \text{fix}(f^m) \cap M^{(n)}) = d(x, \mu_m(x))$ is given by interpretation of the formula $\frac{1}{2}d(x, f^m(x))$ in $\mathcal{M}$. The same formula works for any model $\mathcal{M}$, therefore we have uniform definability. □

Now, consider any $L_s(\mu)$-formula $\psi(x)$. Using Lemma 4.3.2 and the proof of [1, Proposition 9.17] we can show that

$$\sup_{x \in \text{fix}(f^m) \cap M^{(n)}} \psi(x)$$

is the interpretation of an $L_s(\mu)$-formula in all models of $\mathbb{ERT}_s(\mu)$. Specifically, the proof of [1, Proposition 9.17] uses the fact that given a formula $\psi(x)$ with range interval $I_\psi = [0, b]$ and $n \in \mathbb{N}^{>0}$, there exists an increasing, continuous function $\alpha \colon [0, 2n] \to [0, b]$ with $\alpha(0) = 0$ such that $|\psi(y) - \psi(z)| \leq \alpha(d(y, z))$ holds in any $\mathcal{M}(\mu) \models \mathbb{ERT}_s(\mu)$ and any $y, z \in M^{(n)}$. Then, it is shown that in any $\mathcal{M}(\mu) \models \mathbb{ERT}_s(\mu)$, the predicate $\inf_{x \in \text{fix}(f^m) \cap M^{(n)}} \psi(x)$ is equal to the interpretation of the $L_s(\mu)$-formula

$$\inf_{n,x} \big(\psi(x) + \alpha\big(\frac{1}{2} d(x, f^m(x))\big)\big).$$

Lastly, the sup quantification is handled by noting

$$\sup_{x \in \text{fix}(f^m) \cap M^{(n)}} \psi(x) = b \dot{-} \Big(\inf_{x \in \text{fix}(f^m) \cap M^{(n)}} \big(b \dot{-} \psi(x)\big)\Big).$$

Finally, we are ready to extend $\mathbb{ERT}_s$ to the L-theory $\text{rb}\mathbb{ERT}_s$, which turns out to be its model companion.

4.3.3 Definition. For $r \in \mathbb{Q}, h, l, m, n \in \mathbb{N}^{>0}$, define $\widetilde{\varphi}_{r,h,l,m,n}(x)$ to be the $L_s(\mu)$-formula

$$\inf_{n+r, y_1, \ldots, y_l} \max\{|d(x, y_i) - r|, |2r - d(y_i, y_j)|, d(\mu_{km}(y_i), x), d(\mu_{hm}(y_i), y_i)\}$$

where the maximum is over all $1 \leq k \leq h-1$ and all $1 \leq i \neq j \leq l$, and x is a variable of sort $M^{(n)}$. Further, define $\varphi_{r,h,l,m,n}$ to be

$$\sup_{x \in \text{fix}(f^m) \cap M^{(n)}} \min\Big(\{d(x, f^k(x)) : k < m\} \cup \{\widetilde{\varphi}_{r,h,l,m,n}(x)\}\Big)$$

which we regard as an $L_s(\mu)$-formula as discussed above.

4.3.4 Definition. Let $\text{rbE}\mathbb{R}\text{T}_s(\mu)$ be the $L_s(\mu)$-theory whose axioms are the axioms of $\text{E}\mathbb{R}\text{T}_s(\mu)$ and $\text{rb}\mathbb{R}\text{T}$ together with all the conditions $\varphi_{r,h,l,m,n} = 0$ where $r \in \mathbb{Q};\ h, l, m, n \in \mathbb{N}^{>0}$. We also let $\text{rbE}\mathbb{R}\text{T}_s$ be the restriction of $\text{rbE}\mathbb{R}\text{T}_s(\mu)$ to L_s.

Note that the preceding discussion allows us to write explicit L_s-conditions that axiomatize $\text{rbE}\mathbb{R}\text{T}_s$, and yields the following description of its models:

4.3.5 Proposition. *1. Every model $\mathcal{M}$ of $\text{rbE}\mathbb{R}\text{T}_s$ has a unique expansion $\mathcal{M}(\mu)$ to a model of $\text{rbE}\mathbb{R}\text{T}_s(\mu)$, and every model of $\text{rbE}\mathbb{R}\text{T}_s(\mu)$ is of this form.*

2. Let $\mathcal{M}$ be an ω_1-saturated model of $\text{E}\mathbb{R}\text{T}_s$, with underlying $\mathbb{R}$-tree and elliptic isometry (M, d, p, f). Then $\mathcal{M} \models \text{rbE}\mathbb{R}\text{T}_s$ if and only if M is a richly branching $\mathbb{R}$-tree and for all $x \in M$ of order $m \geq 1$, all $r \in \mathbb{R}^{>0}$, and all $h, l \in \mathbb{N}$, there exist $y_1, \ldots, y_l \in M$ such that for all $i, j = 1, \ldots, l$ we have $d(x, y_i) = r$, $d(y_i, y_j) = 2r$ (if $i \neq j$) and every element of $(x, y_i]$ has order hm.

For more explanation of statement (2) in the proposition above, see the proof of 4.3.6 below. Next, we develop some technical machinery needed for the proof that $\text{rbE}\mathbb{R}\text{T}_s$ is the model companion of $\text{E}\mathbb{R}\text{T}_s$.

4.3.6 Lemma. *Let $\kappa \geq \omega_1$ be a cardinal. Assume $\mathcal{M}$ is a κ-saturated model of $\text{rbE}\mathbb{R}\text{T}_s$ with underlying $\mathbb{R}$-tree and isometry (M, d, p, f).*

1. For any $a \in M$ with order $m \in \mathbb{N}^{>0}$, any $r \in \mathbb{R}^{>0}$ and any $h \in \mathbb{N}$ there are at least κ many points b on distinct branches at a such that $d(a, b) = r$ and all the points in $(a, b]$ have order hm.

2. *For any* $a \in M$ *with order* $m \in \mathbb{N}^{>0}$*, and any* $r \in \mathbb{R}^{>0}$ *there are at least* κ*-many points* b *on distinct branches at* a *whose orbits are disjoint such that* $d(a,b) = r$*, and all the points in* $(a, b]$ *have infinite order.*

3. *For any* $a \in M$ *with order* ∞*, and any* $r \in \mathbb{R}^{>0}$ *there are at least* κ *many points* b *on distinct branches at* a *whose orbits are disjoint such that* $d(a,b) = r$*, and all the points in* $[a, b]$ *have infinite order.*

Proof. Assume the hypotheses of the lemma. Let $a \in M$ with order $m \in \mathbb{N}^{>0}$. Take $r \in \mathbb{R}^{>0}$ and $h, l \in \mathbb{N}$. Let n be large enough so that $a \in B_n(p) = M^{(n)}$. Then because $\mathcal{M} \models \varphi_{r,h,l,m,n} = 0$, and $a \in \text{fix}(f^m) \cap M^{(n)}$ we conclude

$$\min\Big(\{d^{\mathcal{M}}(a, f^k(a)) : k < m\} \cup \{\widetilde{\varphi}^{\mathcal{M}}_{r,h,l,m,n}(a)\}\Big) = 0.$$

Since the order of a is actually m, we know that $\min\{d^{\mathcal{M}}(a, f^k(a)) : k < m\} > 0$. Therefore, $\widetilde{\varphi}^{\mathcal{M}}_{r,h,l,m,n}(a)$ must equal 0.

Because $\mathcal{M}$ is κ-saturated for $\kappa \geq \omega_1$, we know that the inf in $\widetilde{\varphi}^{\mathcal{M}}_{r,h,l,m,n}(a)$ is satisfied exactly by Lemma 1.7.3. Therefore, there exist $b_1, ..., b_l \in M^{(n+r)}$ such that for all $i, j \in \{1, ..., l\}$ and all $k \in \{1, ..., h-1\}$:

$$|d(a, b_i) - r| = 0 \tag{4.1}$$

$$|2r - d(b_i, b_j)| = 0 \tag{4.2}$$

$$d(\mu_{km}(b_i), a) = 0 \tag{4.3}$$

$$d(\mu_{hm}(b_i), b_i) = 0. \tag{4.4}$$

The statements in (4.1) give that $d(a, b_i) = r$ for each $1 \leq i \leq l$. The statements in (4.1) and (4.2) imply by Lemma 2.1.2 that for $i \neq j$ the points b_i and b_j are on different branches a. For each $i = 1, ..., l$, statement (4.4) gives that $b_i \in \text{fix}(f^{hm})$, while the statements in (4.3) say that for all $1 \leq k \leq h-1$ the closest point to b_i in $\text{fix}(f^{km})$ is a. This implies that the order of every point in $(a, b_i]$ is greater than km for all $1 \leq k \leq h-1$. Therefore the order of b_i is exactly hm for $i = 1, ..., l$, and the order of every point in $(a, b_i]$ must be hm. Note that this gives us l-many such points b_i on distinct branches out of a, but it does not guarantee that their orbits do not intersect. To get l'-many such points b_i so that their orbits do not intersect, we can use this argument with $l \geq l'mh$.

Proof of (1): Let $h \in \mathbb{N}$. Towards contradiction, assume there are $< \kappa$ many points b on distinct branches at a such that $d(a, b) = r$ and all the points in $(a, b]$ have order hm. List them as $\{b_i\}_{i \in I}$ where $\text{card}(I) < \kappa$. Let A be the set $\{a\} \cup \{b_i\}_{i \in I}$. Note that this set also has cardinality $< \kappa$. The set of $L_s(A)$-conditions

$$\{|d(a, y) - r| = 0; |d(b_i, y) - 2r| = 0; d(\mu_{km}(y), a) = 0; d(\mu_{hm}(y), y) = 0 \mid i \in I,\ 1 \leq k \leq h - 1\}$$

is finitely satisfiable in $\mathcal{M}$. Therefore, by κ-saturation there is a point b that satisfies all these conditions at once. Since $d(\mu_{km}(b), a) = 0$ for all $1 \leq k \leq h - 1$ and $d(\mu_{hm}(b), b) = 0$, all points in $(a, b]$ must have order hm. Since $|d(a, b) - r| = 0$ and $|d(b_i, b) - 2r| = 0$, the point b must be on a different branch out of b than all the b_i. This is a contradiction.

Proof of (2): To see that there is at least one point b on a branch at a such that $d(a, b) = r$ and all the points in $(a, b]$ have order ∞, note that the set of $L_s(a)$ conditions

$$\{|d(a, y) - r| = 0; d(\mu_{km}(y), a) = 0; \mid k \in \mathbb{N}^{>0}\}$$

is finitely satisfiable in $\mathcal{M}$. So, by κ-saturation there is some b satisfying all of the L_s-conditions in the set at once. This b is on a branch at a such that $d(a, b) = r$ and all the points in $(a, b]$ have order ∞.

Now, towards contradiction, assume there are $< \kappa$ many points b on distinct branches at a with disjoint orbits, such that $d(a, b) = r$, and all the points in $(a, b]$ have order ∞. List them as $\{b_i\}_{i \in I}$ for $\text{card}(I) < \kappa$. Let A be the set $\{a\} \cup \bigcup_{l \in \mathbb{N}} \{f^l(b_i)\}_{i \in I}$. Note that A has cardinality $< \kappa$ The set of $L_s(A)$-conditions

$$\{|d(a, y) - r| = 0; |d(f^l(b_i), y) - 2r| = 0; d(\mu_{km}(y), a) = 0 \mid i \in I,\ k, l \in \mathbb{N}^{>0}\}$$

is finitely satisfiable in $\mathcal{M}$. Therefore, by κ-saturation there is a point b that satisfies all these conditions at once. This point b must be on a different branch at a from all the b_i, and it must be such that $\mu_{km}(b) = a$ for all $k \in \mathbb{N}^{>0}$. Therefore b has infinite order, which is a contradiction.

Proof of (3): Let $a \in M$ have order ∞. Because $\text{rb}\mathbb{R}\text{T} \subseteq \text{rbE}\mathbb{R}\text{T}_s$, we may use Lemma 2.3.7 to conclude that there are at least κ-many distinct branches at a, and each of these branches has infinite extent. So, we have κ-many branches at a different from the branch that contains $\mu_1(a)$. Then, for any $r \in \mathbb{R}^{>0}$, there exist κ-many points b on distinct branches at a with $d(a, b) = r$

and $a \in [\mu_1(a), b]$. It follows that $\mu_1(a) = \mu_1(b)$, and by Lemma 4.2.1 b must have infinite order. Moreover, these b must all have disjoint orbits. This is because if there are distinct b and b' meeting the criteria above so that $f^m(b) = b'$ for some $m \in \mathbb{Z}$, then since a is the midpoint of $[b, b']$, a must be fixed by f^m by Lemma 3.1.3. But, a had order ∞, so this is impossible. $\square$

In the next three lemmas, we build up the tools for showing that every highly saturated model of rbE$\mathbb{R}$T$_s$ is existentially closed as a model of rbE$\mathbb{R}$T$_s$. (From this fact it follows that rbE$\mathbb{R}$T$_s$ is the model companion of E$\mathbb{R}$T$_s$; see Theorem 4.3.13 below).

4.3.7 Lemma. *Assume $\mathcal{N} \models$ E$\mathbb{R}$T$_s$, and let (N, d, p, f) be the underlying $\mathbb{R}$-tree and elliptic isometry of $\mathcal{N}$. Let $b_1, ..., b_n \in N$, and $\mathcal{T}$ a non-empty closed subtree of N such that $\mathcal{T}$ is closed under f and f^{-1}. Let B be the set*

$$\bigcup_{m \in \mathbb{Z}} \{f^m(b_i) \mid i = 1, ..., n\}.$$

Define an equivalence relation on B by: $x \sim y$ if and only if x and y have the same unique closest point in $\mathcal{T}$. Let $\{G_h\}_{h=1}^{\infty}$ be the equivalence classes of $\sim$. Without loss of generality, assume $b_1, ..., b_n \in \bigcup_{h=1}^{n} G_h$. For each h, let e_h be the unique closest point in $\mathcal{T}$ common to all members of G_h. Then,

1. *the set of equivalence classes of $\sim$ is closed under f and f^{-1};*
2. *for each $h \geq n+1$, there is some $j \in \{1, ..., n\}$ and some $m \in \mathbb{Z}$ such that $f^m(G_j) = G_h$ and $f^m(e_j) = e_h$.*

Proof. The proofs are the same as for Lemma 3.4.4. $\square$

Let $\mathcal{N} \models$ rbE$\mathbb{R}$T$_s$ be an extension of $\mathcal{M} \models$ rbE$\mathbb{R}$T$_s$ where $\mathcal{M}$ is highly saturated. Let $\mathcal{T}$ be a closed subtree of $\mathcal{M}$. In the next lemma, given a point $b \in N$ we see how to embed the segment from b to $\mathcal{T}$ into $\mathcal{M}$ so that the structure of the orbit of b is preserved.

4.3.8 Lemma. *Let $\mathcal{M} \models$ rbE$\mathbb{R}$T$_s$ and $\mathcal{N} \models$ E$\mathbb{R}$T$_s$ an extension of $\mathcal{M}$ with (N, d, p, f) the underlying $\mathbb{R}$-tree and elliptic isometry of $\mathcal{N}$ and (M, d, p) the underlying $\mathbb{R}$-tree of $\mathcal{M}$. Assume $\mathcal{M}$ is κ-saturated for $\kappa \geq \omega_1$. Let $\mathcal{T}$ be a non-empty closed subtree of M that is closed under f and f^{-1}. Let $b \in N \setminus M$ and $e \in \mathcal{T}$ be the closest point in $\mathcal{T}$ to b. Take $\alpha < \kappa$ and $\{\beta_i \mid i \in \alpha\}$ a family of branches at e in M. Then, there is an isometric embedding g of $[e, b]$ into M so that $g(e) = e$ and $d(g(b), \mu_m(g(b))) = d(b, \mu_m(b))$ for all $m \in \mathbb{N}$, and the image of g does not intersect $\bigcup\{\beta_i \mid i \in \alpha\}$ except at e.*

Proof. Let the situation be as described in the statement of the lemma. Recall that $\mu_1(e)$ is the midpoint of $[e, f(e)]$. The $\mathbb{R}$-tree $\mathcal{T}$ is closed under f and convex in N, so $e \in \mathcal{T}$ implies $\mu_1(e) \in \mathcal{T}$. Therefore, $[\mu_1(e), e] \subseteq \mathcal{T}$, and so by Lemma 4.2.1 part (1), $\mu_m(e) \in \mathcal{T}$ for all $m \in \mathbb{N}$.

Assume e has infinite order. Then b must also have infinite order. To embed $[e, b]$ in M, apply Lemma 4.3.6 to find a point y on a branch out of e in M that does not intersect $\bigcup\{\beta_i \mid i \in \alpha\}$ except at e, such that $d(e, y) = d(e, b)$ and every point in $[e, y]$ has infinite order. We may avoid these α-many branches because there are at least κ-many branches out of e with the properties we need. Then, define $g(e) = e$ and $g(b) = y$ and extend g to $[e, b]$ isometrically. Every point on $[e, b]$ has infinite order. It follows that $\mu_m(e) = \mu_m(b)$ for all $m \in \mathbb{N}^{>0}$. Every point in $[g(e), g(b)] = [e, g(b)]$ has infinite order, so we know $\mu_m(g(b)) = \mu_m(g(e))$ for all $m \in \mathbb{N}^{>0}$. We conclude $\mu_m(g(b)) = \mu_m(g(e)) = \mu_m(e) = \mu_m(b)$ for all $m \in \mathbb{N}^{>0}$. Therefore, $d(g(b), \mu_m(g(b))) = d(b, \mu_m(b))$ for all $m \in \mathbb{N}^{>0}$.

Next assume both e and b have order one. Then all of $[e, b]$ has order one. Use Lemma 4.3.6 to find $y \in M$ on a branch out of e that does not intersect $\bigcup\{\beta_i \mid i \in \alpha\}$ except at e, with $d(e, y) = d(e, b)$ and so that every point in $[e, y]$ has order one. Define $g(e) = e$ and $g(b) = y$ and extend to $[e, b]$ isometrically. Note that since b and $g(b)$ both have order one, $\mu_m(b) = b$ and $\mu_m(g(b)) = g(b)$ for all $m \in \mathbb{N}^{>0}$. Therefore, $d(g(b), \mu_m(g(b))) = d(b, \mu_m(b))$ for all $m \in \mathbb{N}^{>0}$.

Now, assume $m \in \mathbb{N}^{>0}$ is the order of e and b has finite order $t_k > 1$. Then by Lemma 4.2.3 we know

$$[\mu_1(b), b] = [\mu_1(b), \mu_{t_1}(b)] \cup (\mu_{t_1}(b), \mu_{t_2}(b)] \cup ... \cup (\mu_{t_{k-1}}(b), \mu_{t_k}(b)]$$

where $1 < t_1 < t_2 < ... < t_k$ are the elements of $O(b)$, every point in $(\mu_1(b), \mu_{t_1}]$ has order t_1, and every point in $(\mu_{t_j}(b), \mu_{t_{j+1}}(b)]$ has order t_{j+1} for $j = 1, ..., k-1$.

Case I: Assume there is $t \in O(b)$ so that $m = t$ and $e = \mu_t(b)$. If $t = 1$, begin by defining g on $[\mu_1(b), \mu_{t_1}(b)] = [e, \mu_{t_1}(b)]$. Use Lemma 4.3.6 to find $y_1 \in M$ on a branch out of $e = \mu_1(b)$ that does not intersect $\bigcup\{\beta_i \mid i \in \alpha\}$ except at e, with $d(e, y_1) = d(e, \mu_{t_1}(b))$ and so that every point in $(e, y_1]$ is has order t_1. Define $g(e) = e$ and $g(\mu_1(b)) = y_1$ and extend g to all of $[e, \mu_{t_1}(b)]$ isometrically. Next, we extend the domain of g to include $(\mu_{t_1}(b), \mu_{t_2}(b)]$. Use Lemma 4.3.6 to find $y_2 \in M$ on a branch out of $y_1 = g(\mu_{t_1}(b))$ that does not intersect the previously defined image of g, with $d(y_1, y_2) = d(\mu_{t_1}(b), \mu_{t_2}(b))$ and so that every point in $(y_1, y_2]$ is has order t_2. Define $g(\mu_{t_2}(b)) = y_2$ and extend g to all of $[e, \mu_{t_1}(b)]$ isometrically. Repeat this process until you have embedded all of $[e, b]$. Note that, $\mu_t(b) \in [e, b]$ for all $t \in O(b)$, and by our construction, $d(e, \mu_t(b)) = d(g(e), \mu_t(g(b)))$ for all $t \in O(b)$. This implies $d(b, \mu_t(b)) = d(g(b), \mu_t(g(b)))$ for all

$t \in O(b)$. So by Lemma 4.2.4 we know that $d(g(b), \mu_m(g(b))) = d(b, \mu_m(b))$ for all $m \in \mathbb{N}^{>0}$.

If $t > 1$, then $t = t_j$ for some $j = 1, ..., k$. Begin by defining g on $[e, \mu_{t_{j+1}}(b)]$, which is equal to $[\mu_{t_j}(b), \mu_{t_{j+1}}(b)]$. Use Lemma 4.3.6 to find $y_{j+1} \in M$ on a branch out of $e = \mu_{t_j}(b)$ that does not intersect $\bigcup\{\beta_i \mid i \in \alpha\}$ except at e, with $d(e, y_{j+1}) = d(e, \mu_{t_{j+1}}(b))$ and so that every point in $(e, y_{j+1}]$ is has order t_{j+1}. Define $g(e) = e$ and $g(\mu_{t_{j+1}}(b)) = y_{j+1}$ and extend to $[e, \mu_{t_{j+1}}(b)]$ isometrically. Then continue as above to define g on $(\mu_{t_{j+2}}(b), \mu_{t_{j+3}}(b)] \cup ... \cup (\mu_{t_{k-1}}(b), \mu_{t_k}(b)]$, which is the rest of $[e, b]$. Note that $\mu_t(b) \in [e, b]$ for all $t \in O(b)$ such that $t > t_j$, and by our construction, $d(e, \mu_t(b)) = d(g(e), \mu_t(g(b)))$ for all $t \in O(b)$ such that $t > t_j$. This implies $d(b, \mu_t(b)) = d(g(b), \mu_t(g(b)))$ for all $t \in O(b)$ such that $t > t_j$. For $t \in O(a)$ such that $t \leq t_j$, we know that $\mu_t(b) = \mu_t(e) \in M$, and $e \in [\mu_t(b), b]$ and $d(e, g(b)) = d(e, b)$. It follows that $d(b, \mu_t(b)) = d(g(b), \mu_t(g(b)))$ for $t \in O(a)$ such that $t \leq t_j$. So by Lemma 4.2.4 we know that $d(g(b), \mu_m(g(b))) = d(b, \mu_m(b))$ for all $m \in \mathbb{N}^{>0}$.

Case II: Assume $m = 1$, but $e \neq \mu_1(b)$. Begin by defining g on $[e, \mu_1(b)]$. Use Lemma 4.3.6 to find $y_1 \in M$ on a branch out of e that does not intersect $\bigcup\{\beta_i \mid i \in \alpha\}$ except at e, with $d(e, y_1) = d(e, \mu_1(b))$ and so that every point in $(e, y_1]$ is has order one. Define $g(e) = e$ and $g(\mu_1(b)) = y_1$ and extend to $[e, \mu_1(b)]$ isometrically. Then, proceed as in Case I to embed the rest of $[e, b]$. Note that $\mu_t(b) \in [e, b]$ for all $t \in O(a)$, and by our construction, $d(e, \mu_t(b)) = d(g(e), \mu_t(g(b)))$ for all $t \in O(a)$. It follows that $d(b, \mu_t(b)) = d(g(b), \mu_t(g(b)))$ for $t \in O(a)$ such that $t \leq t_j$. So by Lemma 4.2.4 we know that $d(g(b), \mu_m(g(b))) = d(b, \mu_m(b))$ for all $m \in \mathbb{N}^{>0}$.

Case III: Assume $m > 1$ but $e \neq \mu_m(b)$. Say $m = t_j$. Begin by defining g on $[e, \mu_{t_j}(b)]$. Use Lemma 4.3.6 to find $y_j \in M$ on a branch out of e that does not intersect $\bigcup\{\beta_i \mid i \in \alpha\}$ except at e, with $d(e, y_j) = d(e, \mu_{t_j}(b))$ and so that every point in $(e, y_j]$ is has order t_j. Define $g(e) = e$ and $g(\mu_{t_j}(b)) = y_j$ and extend to $[e, \mu_{t_j}(b)]$ isometrically. Then, proceed as in Case I to embed the rest of $[e, b]$. In this case, we also may conclude that $d(g(b), \mu_m(g(b))) = d(b, \mu_m(b))$ for all $m \in \mathbb{N}^{>0}$.

Now, assume $m \in \mathbb{N}^{>0}$ is the order of e and that b has order ∞. There are two more cases under this assumption.

Case IV: Assume $O(a)$ is finite with largest element t_k. Then, by Lemma 4.2.3, $[\mu_1(b), b] = [\mu_1(b), \mu_{t_1}(b)] \cup (\mu_{t_1}(b), \mu_{t_2}(b)] \cup ... \cup (\mu_{t_k}(b), b]$ where every point in $(\mu_1(b), \mu_{t_1}]$ has order t_1, every point in $(\mu_{t_j}(b), \mu_{t_{j+1}}(b)]$ has order t_{j+1} for $j = 1, ..., k-1$, and every point in $(\mu_{t_k}(b), b]$ has infinite order. Because $\mu_{t_k}(b)$ has finite order, and $\mu_m(\mu_{t_k}(b)) = \mu_m(b)$ for all $m \in \mathbb{N}$ we may use the cases above to define g with the desired properties on $[e, \mu_{t_k}(b)]$. Then, by Lemma 4.3.6 part (3) we can find y_∞ such that $d(y_k, y_\infty) = d(\mu_{t_k}(b), b)$ and every point in $(y_k, y_\infty]$ has infinite order, on a branch in M out of $g(\mu_{t_k}(b)) = y_k$ that does not intersect $\bigcup\{\beta_i \mid i \in \alpha\} \setminus \{e\}$ and does not intersect the image of g built previously except at y_k. Define $g(b) = y_\infty$ and extend g to $[\mu_{t_k}(b), b]$ isometrically. This defines g on $[e, b]$ such that $\mu_m(g(b)) = g(\mu_m(b))$

Case V: Assume the order of b is infinite and $O(a)$ is infinite. Then by Lemma 4.2.3 we have $\{t_j\}_{j=1}^{\infty}$ with $1 < t_j$ for all j so that every point in $(\mu_1(b), \mu_{t_1}]$ has order t_1, and for all j, every point in $(\mu_{t_j}(b), \mu_{t_{j+1}}(b)]$ has order t_{j+1}. Also, $\{\mu_{t_j}(b)\}_{j=1}^{\infty}$ converges to a point $\mu_\infty(b)$ so that every point in $[\mu_\infty(b), b]$ has infinite order, and

$$[\mu_1(b), b] = [\mu_1(b), \mu_{t_1}(b)] \cup \Big(\bigcup_{j=1}^{\infty} (\mu_{t_j}(b), \mu_{t_{j+1}}(b)]\Big) \cup [\mu_\infty(b), b].$$

For each $j \in \mathbb{N}$, proceed as in Case II to find isometric embeddings $g_j \colon [e, \mu_{t_j}(b)] \to M$ such that $g_j \subseteq g_{j+1}$ and so that $d(e, g_j(\mu_{t_j}(b)) = d(e, \mu_{t_j}(b))$. Let g be the union of those embeddings. Then, g is an isometric embedding of $[e, \mu_\infty(b))$ into M.

The distances $d(e, g(\mu_{t_j}(b)))$ are all $\leq d(e, b)$ and hence form a bounded, increasing sequence of real numbers. Since

$$d\big(g(\mu_{t_j}(b)), g(\mu_{t_l}(b))\big) = \big|d\big(e, g(\mu_{t_j}(b))\big) - d\big(e, g(\mu_{t_l}(b))\big)\big|$$

it follows that $\{g(\mu_{t_j}(b))\}_{j\in\mathbb{N}}$ form a Cauchy sequence in M. Since M is complete, $\{g(\mu_{t_j}(b))\}_{j\in\mathbb{N}}$ converges to $z \in M$. Define $g(\mu_\infty(b)) = z$. We now have an isometric embedding of $[e, \mu_\infty(b)]$ into M, and the last step is to extend to all of $[e, b]$. Use Lemma 4.3.6 to find $y \in M$ with $d(\mu_\infty(b), b) = d(z, y)$ so that every point in $[z, y]$ has infinite order. Find this y on a branch out of $z = g(\mu_\infty(b))$ that does not intersect $\bigcup\{\beta_i \mid i \in \alpha\} \setminus \{e\}$ and does not intersect the previously defined image of g except at z. Set $g(b) = y$ and extend to $[z, y]$ isometrically. We may show that in this last case $d(g(b), \mu_m(g(b))) = d(b, \mu_m(b))$ is true for all $m \in \mathbb{N}$, since the corresponding statement was true for each g_j. □

In this next lemma, we extend our embedding from Lemma 4.3.8 to the subtree defined by the orbit of the segment.

4.3.9 Lemma. *Let $\mathcal{M} \models \text{rb}\mathbb{E}\mathbb{R}\text{T}_s$ and $\mathcal{N} \models \mathbb{E}\mathbb{R}\text{T}_s$ an extension of $\mathcal{M}$, with (N, d, p, f) the underlying $\mathbb{R}$-tree and elliptic isometry of $\mathcal{N}$, and (M, d, p) the underlying $\mathbb{R}$-tree of M. Assume $\mathcal{M}$ is κ-saturated for some cardinal $\kappa \geq \omega_1$. Let $\mathcal{T}$ be a non-empty closed subtree of M closed under f and f^{-1}. Let $b \in N \setminus M$ and $e \in M$ the closest point in $\mathcal{T}$ to b. Take $\alpha < \kappa$ be a cardinal, and $\{\beta_i \mid i \in \alpha\}$ a family of branches at e in M. Then, there is an isometric embedding g of $\bigcup_{m\in\mathbb{Z}} f^m([e, b])$ into M so that $g(e) = e$ and g commutes with f and so that the image of g does not intersect $\bigcup_{i\in\alpha}\{f^m(\beta_i) \mid m \in \mathbb{Z}\}$ except at the points $f^m(e)$ for $m \in \mathbb{Z}$.*

Proof. Assume the situation described in the hypotheses. Then, by Lemma 4.3.8 there exists an isometric embedding $g\colon [e, b] \to M$ so that $g(e) = e$, the image of g intersects $\bigcup_{i\in\alpha} \beta_i$ only at e, and $d(g(b), \mu_m(g(b))) = d(b, \mu_m(b))$ for all $m \in \mathbb{N}$. Then, by Lemma 4.2.5, for all $k, l \in \mathbb{Z}$

$$d(f^k(b), f^l(b)) = d(f^k(g(b)), f^l(g(b))).$$

Therefore we may extend g to an isometric embedding from $\bigcup_{m\in\mathbb{Z}} f^m([e, b]) = \bigcup_{m\in\mathbb{Z}}[e, f^m(b)])$ into M by defining $g(f^m(b)) = f^m(g(b))$ for each $m \in \mathbb{Z}$ and extending isometrically to the segments between. It is immediate from the definition of g that g commutes with f. The image of g does not intersect $\bigcup_{i\in\alpha}\{f^m(\beta_i) \mid m \in \mathbb{Z}\}$ except at the points $f^m(e)$ for $m \in \mathbb{Z}$, because the image of g can only intersect $\bigcup_{i\in\alpha} \beta_i$ at e and g commutes with f. $\square$

4.3.10 Lemma. *Let $\mathcal{M} \models \text{rb}\mathbb{E}\mathbb{R}\text{T}_s$ and $\mathcal{N} \models \mathbb{E}\mathbb{R}\text{T}_s$ an extension of $\mathcal{M}$ with (N, d, p, f) the underlying $\mathbb{R}$-tree and elliptic isometry of $\mathcal{N}$ and (M, d, p) the underlying $\mathbb{R}$-tree of $\mathcal{M}$. Assume $\mathcal{M}$ is κ-saturated for $\kappa > 2^\omega$. Let $a_1, ..., a_k \in M$, $b_1, ..., b_n \in N \setminus M$. Let the set B and $\{G_h\}_{h=1}^{\infty}$ and e_h be as in the hypotheses of Lemma 4.3.7, with the tree $\mathcal{T} = \overline{E_A}$. Define K_h to be the subtree of N generated by $G_h \cup \{e_h\}$. Let $K = K_h$ for some $h \in \mathbb{N}^{>0}$. Then, there exists an isometric embedding*

$$g\colon \bigcup_{m\in\mathbb{Z}} f^m(K) \to M$$

with $g(e_h) = e_h$ such that g commutes with f, and so that the image of g intersects $\overline{E_A}$ only at the points $f^m(e_h)$ for $m \in \mathbb{Z}$.

Proof. Assume the situation described in the hypotheses of the lemma. We may assume without loss of generality that $b_1, ..., b_k$ are exactly the members of B that are in K, where $k \leq n$. Let $e = e_h$ be the closest point in $\overline{E_A}$ common to all points in G_h. Then e is the closest point in $\overline{E_A}$ to any point from K.

Define $A_1 \subseteq A_2 \subseteq ... \subseteq A_k \subseteq \bigcup_{m\in\mathbb{Z}} f^m(K)$ as follows:

- $A_1 = \bigcup_{m\in\mathbb{Z}} f^m([e, b_1])$
- $A_2 = A_1 \cup \bigcup_{m\in\mathbb{Z}} f^m([z_2, b_2])$ where z_2 is the closest point to b_2 in $[e, b_1]$;
- $A_{i+1} = A_i \cup \bigcup_{m\in\mathbb{Z}} f^m([z_{i+1}, b_{i+1}])$ where z_{i+1} is the closest point to b_{i+1} in $[e_1, b_1] \cup [z_2, b_2] \cup ... \cup [z_i, b_i] \subseteq K$.

Then, $A_k = \bigcup_{m\in\mathbb{Z}} f^m(K)$, since K is generated by $b_1, ..., b_k$. We begin by defining g on $[e, b_1]$. By Lemma 4.3.9 there exists an isometric embedding g_1 of $A_1 = \bigcup_{m\in\mathbb{Z}} f^m([e, b_1])$ into M so that $g_1(e) = e$ and g_1 commutes with f. Moreover, we may find g_1 so that the image of g_1 only intersects $\overline{E_A}$ at e. This is because the degree of e in $\overline{E_A}$ is at most 2^ω by Lemma 4.2.7, so when we set $\mathcal{T} = \overline{E_A}$, we may let $\{\beta_i \mid i \in 2^\omega\}$ be the collection of branches at e that intersect $\overline{E_A}$ at points other than e. Define $g = g_1$ on A_1. Since we have just embedded A_1 in M, we now assume, for notational convenience, that $A_1 \subseteq M$ and that g_1 was an inclusion map. Let $\mathcal{T}_1 = \overline{E_A} \cup A_1$.

Claim: $\mathcal{T}_1$ is a closed subtree of M. Assume towards contradiction that there is a Cauchy sequence $(x_j)_{j=1}^\infty$ in $\mathcal{T}_1$ whose limit y is not in $\mathcal{T}_1$. Then, the Cauchy sequence must be bounded away from $\overline{E_A} \subseteq \mathcal{T}_1$, which means there exists a $\delta > 0$ and a point c in $[e, b_1]$ with distance $\delta/2$ from e so that eventually, the sequence $(x_j)_{j=1}^\infty$ is contained in $f^m([c, b_1]) \subseteq A_1$. But, these sets $f^m([c, b_1])$ for $m \in \mathbb{Z}$ all have distance at least δ from one another. Therefore, $(x_j)_{j=1}^\infty$ must eventually be in one of them. Each $f^m([c, b_1])$ is closed and hence y is in $A_1 \subseteq \mathcal{T}_1$, which is a contradiction. The sets $\overline{E_A}$ and A_1 are both closed under f and f^{-1}, thus the subtree $\mathcal{T}_1$ is closed under f and f^{-1}.

It is straightforward to see that z_2 is the closest point to b_2 in $\mathcal{T}_1$ (because the $b_1, ..., b_k$ were in K, and hence had the same closest point in $\overline{E_A}$). If $b_2 \in [e, b_1]$, then g is already defined on A_2, so we may assume $b_2 \notin A_1$. We extend the domain of g to $A_2 = A_1 \bigcup_{m\in\mathbb{Z}} f^m([z_2, b_2])$ as follows. Use Lemma 4.3.9 with the tree $\mathcal{T}_1$ and $\{\beta_i \mid i \in 2^\omega\}$ equal to the collection of branches at z_2 that intersect $\mathcal{T}_1$. Because $z_2 \in A_1$, which is generated by $\{b_1\}$, we know this collection is has size at most 2^ω. (If $z_2 = e$, then the size of the collection is 2^ω. If $z_2 \neq e$, then there are only two branches at z_2 that intersect $\mathcal{T}_1$.) We conclude that there exists an isometric embedding g_2 of $\bigcup_{m\in\mathbb{Z}} f^m([z_2, b_2])$ into M sending z_2 to z_2 so that g commutes with f. Since $[e, b_1] \cap \overline{E_A} = \{e\}$, and z_2 is on a branch at e that does not intersect $\overline{E_A}$ except at e, we know that the image of g_2 does not intersect $\overline{E_A}$, except possibly at $f^m(e)$ for $m \in \mathbb{Z}$. Because the domains of g_1 and g_2 were disjoint except at points $f^m(z_2)$ for $m \in \mathbb{Z}$ and $g_1(z_2) = z_2 = g_2(z_2)$ it follows that $g_1 \cup g_2$ is an isometric embedding of A_2 into M whose range intersects $\overline{E_A}$ only at $f^m(e)$. As before, since we have embedded A_2 in M, we may assume $A_2 \subseteq M$ and that g_2 was inclusion.

Let $\mathcal{T}_2 = \mathcal{T}_1 \cup A_2$. Then $\mathcal{T}_2$ is a closed subtree of M that is closed under f and f^{-1}. Then, z_3 is the closest point to b_3 in $\mathcal{T}_2$. We extend the domain of g to $A_3 = A_2 \bigcup_{m\in\mathbb{Z}} f^m([z_3, b_3])$ in the same way we extended from A_1 to A_2. Proceed to extend g in the same manner until you have extended it to $A_k = \bigcup_{m\in\mathbb{Z}} f^m(K)$. This g has the desired properties. □

In this next lemma, given $\mathcal{M} \models \mathrm{E\mathbb{R}T}_s$, for any $h \in \mathbb{N}$ we add new rays each point $a \in M$ of relative order h. That is, if a has order m, the new rays will have order hm, and if a has infinite order the new rays will also have infinite order.

4.3.11 Lemma. *Let $\mathcal{M} \models \mathrm{E\mathbb{R}T}_s$ with underlying $\mathbb{R}$-tree and isometry (M, d, p, f). Let $h \in \mathbb{N}^{>0}$. Then there exists an extension $\mathcal{N} \models \mathrm{E\mathbb{R}T}_s$ of $\mathcal{M}$ with underlying $\mathbb{R}$-tree and isometry (N, d, p, f) such that:*

- *if $a \in M$, and $m \in \mathbb{N}^{>0}$ is the order of a, then there is a new ray R_a at a contained in $\big(N \setminus M\big) \cup \{a\}$, such that the f-order of every point in $R_a \setminus \{a\}$ is equal to hm;*
- *if $a \in M$ has infinite order, there exists a new ray R at a such that every point on R has infinite order.*

Proof. Assume $\mathcal{M} \models \mathrm{E\mathbb{R}T}_s$ with underlying $\mathbb{R}$-tree and isometry (M, d, p, f). Let $h \in \mathbb{N}^{>0}$. Enumerate M by $\{a_i \mid i \in |M|\}$. If a_i has finite f-order, denote the f-order of a_i by m_i. Let T_i be an $\mathbb{R}$-tree made up of h copies of $\mathbb{R}^{\geq 0}$, which all intersect at 0 but are otherwise disjoint. (We could easily build such an $\mathbb{R}$-tree using Theorem 2.2.6.) Let $q_i = 0$ be the basepoint of T_i. Label the h rays of T_i by $R_1^{a_i}, R_2^{a_i}, ..., R_h^{a_i}$. Use Lemma 2.2.8 to find an extension $\mathcal{N}$ of $\mathcal{M}$ that glues in a new copy of T_i at a_i, identifying q_i and a_i. Let (N, d, p) be the underlying $\mathbb{R}$-tree of $\mathcal{N}$. Extend the isometry f on M to these new rays by setting $f(R_j^a) = R_{j+1}^{f(a)}$ for $j = 1, ..., h-1$, and $f(R_h^a) = R_1^{f(a)}$ for each $a \in M$ (mapping R_j^a isometrically to $R_{j+1}^{f(a)}$ by using the fact that they are both isometric to $\mathbb{R}^{\geq 0}$). It is straightforward to check that this extension of f is a well-defined isometry of N.

Let $a \in M$ have order m. To show that every point (except the basepoint) in each "new" ray in $\mathcal{N}$ has order hm, it suffices to show that every point in R_1^a has order hm. If $h = 1$, then

$f^m(R_1^a) = R_1^{f^m(a)} = R_1^a$, but for $k < m$ $f^k(R_1^a) = R_1^{f^k(a)} \neq R_1^a$. Therefore, every point on the R_1^a has order m. If $h = 2$, then $f^m(R_1^a) = R_2^{f^m(a)}$ is a ray out of $f^m(a) = a$ distinct from R_1^a, and

$$f^{2m}(R_1^a) = f^m(f^m(R_1^a)) = f^m(R_2^{f^m(a)}) = f^m(R_2^a) = R_1^{f^m(a)} = R_1^a.$$

For any $k < 2m$, with $k \neq m$, we know $f^m(a) \neq a$, so clearly $f^k(R_1^a) \neq R_1^a$. Therefore, the order of each point but the basepoint on R_1^a is $2m$. For arbitrary $h \in \mathbb{N}^{>0}$ the argument is analogous. Note that in this process, points of infinite order in $\mathcal{M}$ also get new rays added to them, so that every point on those new rays has infinite order. $\square$

4.3.12 Lemma. *Let $\mathcal{M} \models \mathrm{E\mathbb{R}T}_{r,s}$ with underlying $\mathbb{R}$-tree and isometry (M, d, p, f), $a = a_1, ..., a_k \in M$ and $b = b_1, ..., b_n \in M$ with $a_1 = p$. If $c_1, ..., c_n \in M$ satisfy the partial type $D_{f,b}^{\mathcal{M}}(y_1, ..., y_n/a)$ given in Definition 3.4.1, then for any quantifier free L_s-formula $\varphi(x_1, ..., x_k, y_1, ..., y_n)$ we have $\varphi(a_1, ..., a_k, b_1, ..., b_n)^{\mathcal{M}} = \varphi(a_1, ..., a_k, c_1, ..., c_n)^{\mathcal{M}}$.*

Proof. The proof is analogous to that of Lemma 2.5.2. $\square$

4.3.13 Theorem. *The L_s-theory* $\mathrm{rbE\mathbb{R}T}_s$ *is the model companion of* $\mathrm{E\mathbb{R}T}_s$.

Proof. Clearly any model of $\mathrm{rbE\mathbb{R}T}_s$ is also a model of $\mathrm{E\mathbb{R}T}_s$. Assume $\mathcal{M} \models \mathrm{E\mathbb{R}T}_s$ has underlying $\mathbb{R}$-tree and isometry (M, d, p, f). By iterating Lemma 4.3.11, we find an extension of $\mathcal{M}$ that is a model of $\mathrm{rbE\mathbb{R}T}_s$. Let $(n_i)_{i=1}^{\infty}$ be a sequence of positive integers such that every positive integer appears infinitely many times in the sequence. Let $\mathcal{M}_0 = \mathcal{M}$. For $i \geq 1$, let $\mathcal{M}_i$ be the result of applying Lemma 4.3.11 to M_{i-1} with $h = n_i$. This process creates an increasing chain $(\mathcal{M}_i \mid i \in \omega)$ of models of $\mathrm{E\mathbb{R}T}_s$. Let the L_s-structure $\mathcal{W}$ be the union of this chain. Then the underlying space W of $\mathcal{W}$ is an $\mathbb{R}$-tree by Lemma 2.2.7, and the isometry $f^{\mathcal{W}}$ on W is the union of the isometries $f^{\mathcal{M}_i}$. Therefore, $f^{\mathcal{W}}$ is elliptic, since it extends $f^{\mathcal{M}_0}$, and we know $d(p^{\mathcal{W}}, f^{\mathcal{W}}(p^{\mathcal{W}})) = d(p^{\mathcal{M}_0}, f^{\mathcal{M}_0}(p^{\mathcal{M}_0}))$. Let $a \in W$. Let m be the order of a and let $h, l \in \mathbb{N}^{>0}$. For large enough i, the number h has appeared at least l times in our sequence $(n_i)_{i=1}^{\infty}$. Therefore, at a with order $m \in \mathbb{N}^{>0}$, there are at least l-many rays so that every point but a along the ray is of order hm. So, for any $r \in R^{>0}$ we may find $b_1, ..., b_l$ on these rays with $d(a, b_i) = r$ for $i = 1, ..., l$ so that each point in $(a, b_i]$ has order hm. In addition, at any point $a \in M$ of order ∞, there are arbitrarily many rays at a such that every point on those rays has order ∞. This implies that $\mathcal{W} \models \mathrm{rbE\mathbb{R}T}_s$.

It remains to show $\mathrm{rbE\mathbb{R}T}_s$ is model complete. It suffices by Proposition 1.7.9 to show every model of $\mathrm{rbE\mathbb{RT}}_s$ is an existentially closed model of $\mathrm{rbE\mathbb{RT}}_s$. Let $\mathcal{M} \models \mathrm{rbE\mathbb{RT}}_s$. Let $\mathcal{N} \models \mathrm{rbE\mathbb{RT}}_s$ be an extension of $\mathcal{M}$ with underlying $\mathbb{R}$-tree and isometry (N, d, p, f). We may assume that $\mathcal{M}$ is κ-saturated for $\kappa > 2^\omega$. (The reason here is the same as the reason that we may assume ω_1-saturation in the proof of Theorem 2.5.4.) Since $\mathcal{M}$ is κ-saturated, its reduct to L_p is also κ-saturated. Also as in the proof of Theorem 2.5.4, to show $\mathcal{M}$ is existentially closed it suffices, by Lemma 3.4.2, to show: for any $a_1, ..., a_k \in M$, for any extension $\mathcal{N} \models \mathrm{H\mathbb{R}T}_{r,s}$, for any $b_1, ..., b_n \in N$ there exist $c_1, ..., c_n \in M$ such that $c_1, ..., c_n$ satisfy the partial type $D^{\mathcal{N}}_{f,b}(x_1, ..., x_n/a)$.

Let $a_1, ..., a_k \in M$, and let $b_1, ..., b_n \in N$. We may assume that none of the b_i are in M, because if so, we may let $c_i = b_i$ for that i and then add it and its images under f and f^{-1} to our set of parameters A. So, assume $b_1, ..., b_n \in N \setminus M$. Let $\mathcal{T} = \overline{E_A}$, and note by Lemma 4.2.7 that this makes $\mathcal{T}$ a closed subtree of N that is closed under f and f^{-1}. Define

$$B = \bigcup_{m \in \mathbb{Z}} \{f^m(b_i) \mid i = 1, ..., n\},$$

the equivalence relation $\sim$ on B and the classes G_h as in Lemma 4.3.7. For $h \in \mathbb{N}$, let K_h be the subtree of N generated by $G_h \cup \{e_h\}$. Every point in K_h has e_h as its closest point in $\overline{E_A}$. Also, if $K_h \cap K_l \neq \emptyset$, then every point in K_h and every point in K_l must have the same closest point in $\overline{E_A}$. This implies that if $K_h \cap K_l \neq \emptyset$, then $h = l$.

Lemma 4.3.10 implies that given K_h for any $h = 1, ..., n$, there is an isometric embedding $g_h \colon K_h \to M$ sending p_h to e_h whose image does not intersect $\overline{E_A}$ except at e_h and so that g_h commutes with f. Therefore, we may build an isometric embedding

$$g \colon \bigcup_{i=1}^{\infty} K_i \to M$$

exactly as in the proof of Theorem 3.4.5.

Let $c_i = g(b_i)$. Then since g commutes with f, for all $x \in \bigcup_{i=1}^{\infty} K_i$ we know

$$f^m(c_i) = f^m(g(b_i)) = g(f^m(b_i)).$$

Therefore, for any $m, l \in \mathbb{Z}$ and any $i, j \in \{1, ..., n\}$,

$$d(f^m(b_i), f^l(b_j)) = d(g(f^m(b_i)), g(f^l(b_j))) = d(f^m(c_i), f^l(c_j)).$$

In addition, for any point $x \in \overline{E_A}$,

$$\begin{aligned} d(f^m(b_i), x) &= d(f^m(b_i), e_j) + d(e_j, x) \\ &= d(g(f^m(b_i)), g(e_j)) + d(e_j, x) \\ &= d(f^m(c_i), e_j) + d(e_j, x) = d(f^m(c_i), x). \end{aligned}$$

Therefore, we have found $c_1, ..., c_n \in M$ that satisfy the partial type $D_b^f(x_1, ..., x_n)$. □

4.4 Properties of the model companions

4.4.1 Lemma. *The L_s-theory* $\mathrm{rbE\mathbb{R}T}_s$ *has quantifier elimination.*

Proof. By Theorem 4.3.13, Theorem 4.1.4 and Proposition 1.7.13. □

4.4.2 Theorem. *Let π be the partial type*

$$\{d(p, \mu_m(p)) = r_m \mid m \in \mathbb{N}\}$$

where $r_m \in \mathbb{R}^{\geq 0}$. Assume π is consistent with $\mathrm{rbE\mathbb{R}T}_s$ *and let* $\mathrm{rbE\mathbb{R}T}_s^\pi$ *be the L_s-theory consisting of the conditions in* $\mathrm{rbE\mathbb{R}T}_s \cup \pi$*. Then the L_s-theory* $\mathrm{rbE\mathbb{R}T}_s^\pi$ *is complete.*

Proof. Assume the situation given in the hypotheses. Let $\mathcal{M}$ and $\mathcal{N}$ be models of $\mathrm{rbE\mathbb{R}T}_s$. Consider the substructures generated by $p^\mathcal{M}$ and $p^\mathcal{N}$ in $\mathcal{M}$ and $\mathcal{N}$ respectively. If follows from Lemma 4.2.5 and the axioms in π that these substructures are isomorphic. Therefore, the substructure of $\mathcal{M}$ generated by $p^\mathcal{M}$ embeds in any model of $\mathrm{rbE\mathbb{R}T}_s^\pi$. Then, because $\mathrm{rbE\mathbb{R}T}_s^\pi$ has quantifier elimination, we know $\mathrm{rbE\mathbb{R}T}_s^\pi$ is complete. □

It is not hard to see that any completion of $\mathrm{rbE\mathbb{R}T}_s$ has an axiomatization of the form $\mathrm{rbE\mathbb{R}T}_s^\pi$. The completions of $\mathrm{rbE\mathbb{R}T}_s$ are exactly the L_s-theories $\mathrm{Th}(\mathcal{M})$ for $\mathcal{M} \models \mathrm{rbE\mathbb{R}T}_s$. To see this, given $\mathcal{M} \models \mathrm{rbE\mathbb{R}T}_s$, let

$$\pi = \{d(p, \mu_m(p)) = d^\mathcal{M}(p^\mathcal{M}, \mu_m(p^\mathcal{M})) \mid m \in \mathbb{N}\}.$$

Then, $\mathcal{M} \models \mathrm{rbE\mathbb{R}T}_s^\pi$ and by Theorem 4.4.2 we conclude $\mathrm{rbE\mathbb{R}T}_s^\pi$ is equivalent to $\mathrm{Th}(\mathcal{M})$.

In the rest of this section, fix a completion $\mathrm{rbE\mathbb{R}T}_s^\pi$ of $\mathrm{rbE\mathbb{R}T}_s$. Next, we turn to the question of the stability of $\mathrm{rbE\mathbb{R}T}_s^\pi$. Given $\mathcal{M} \models \mathrm{rbE\mathbb{R}T}_s^\pi$ and $A \subseteq M$, let $\mathcal{T}_A$ be the closed subtree of $\mathcal{M}$

generated by the set $\{f^m(a) \mid a \in A \cup \{p\},\ m \in \mathbb{Z}\}$. That is, $\mathcal{T}_A = \overline{E_{\{f^m(a) \mid a \in A \cup \{p\},\ m \in \mathbb{Z}\}}}$.

4.4.3 Lemma. *Let $\mathcal{M} \models \mathrm{rbE\mathbb{R}T}^\pi_s$, and let $b, c \in M$ and $A \subseteq M$. Then $\mathrm{tp}_\mathcal{M}(b/A) = \mathrm{tp}_\mathcal{M}(c/A)$ if and only if*

- *b and c have the same unique closest point in $e \in \mathcal{T}_A$ and $d(e, b) = d(e, c)$;*
- *$d(b, \mu_m(b)) = d(c, \mu_m(c))$ for all $m \in \mathbb{N}$.*

Proof. Assume the situation described in the hypotheses. For the forward direction assume $\mathrm{tp}_\mathcal{M}(b/A) = \mathrm{tp}_\mathcal{M}(c/A)$. Then for all $x \in \mathcal{T}_A$, $d(x, b) = d(x, c)$. (The distance to a point in the closure of a set is determined by distances to points in the set, and the type $\mathrm{tp}_\mathcal{M}(b/A)$ determines all distances to points in $E_{\{f^m(a) \mid a \in A \cup \{p\},\ m \in \mathbb{Z}\}}$.) Therefore, b and c must have the same closest point e in $\mathcal{T}_A$ and it must be that $d(e, b) = d(e, c)$. Also, for any $m \in \mathbb{N}$ we know

$$d(b, \mu_m(b)) = \frac{1}{2} d(b, f^m(b)) = \frac{1}{2} d(c, f^m(c)) = d(c, \mu_m(c)).$$

For the other direction, assume b and c have the same unique closest point in $e \in \mathcal{T}_A$ and $d(e, b) = d(e, c)$, and assume $d(b, \mu_m(b)) = d(c, \mu_m(c))$ for all $m \in \mathbb{N}$. Since $\mathrm{rbE\mathbb{R}T}^\pi_s$ has quantifier elimination, it suffices to show that the quantifier-free types of c and b over A are the same. To show the quantifier-free types are the same, by Lemma 4.3.12 it suffices to show $d(x, f^n(b)) = d(x, f^n(c))$ and $d(b, f^n(b)) = d(c, f^n(c))$ for all $n \in \mathbb{N}$, for all $x \in \mathcal{T}_A$. Since $d(b, \mu_m(b)) = d(c, \mu_m(c))$ for all $m \in \mathbb{N}$, Lemma 4.2.5 implies $d(b, f^n(b)) = d(c, f^n(c))$ for all $n \in \mathbb{N}$. If $x \in \mathcal{T}_A$ and $n \in \mathbb{N}$, then $f^n(e) \in [x, f^n(b)]$ and $f^n(e) \in [x, f^n(c)]$ because $f^n(e)$ must be the closest point in $\mathcal{T}_A$ to $f^n(b)$ and $f^n(c)$ (since $\mathcal{T}_A$ is closed under f and f^{-1}. Therefore, $d(x, f^n(b)) = d(x, f^n(e)) + d(f^n(e), f^n(b)) = d(x, f^n(e)) + d(f^n(e), f^n(c)) = d(x, f^n(c))$. This implies that $d(x, f^n(b)) = d(x, f^n(c))$ for all $x \in \mathcal{T}_A$ and all $n \in \mathbb{N}$. □

4.4.4 Theorem. *The L_s-theory $\mathrm{rbE\mathbb{R}T}^\pi_s$ is stable. Indeed, when κ is an infinite cardinal, $\mathrm{rbE\mathbb{R}T}^\pi_s$ is κ-stable if and only if κ satisfies $\kappa^\omega = \kappa$.*

Proof. By a counting argument as in Theorem 2.6.4, using Lemma 4.4.3 instead of Lemma 2.6.3. □

Let κ be a cardinal so that $\kappa = \kappa^\omega$ and $\kappa > 2^\omega$. Let U be a κ-universal domain for $\mathrm{rbE\mathbb{R}T}^\pi_s$.

4.4.5 Definition. Let A, B and C be small subsets of U. Let $\widehat{A} := \{f^m(a) \mid a \in A \cup \{p\},\ m \in \mathbb{Z}\}$, and let $\widehat{B}$ and $\widehat{C}$ be defined analogously. Say A is ^{E}-independent from B over C, denoted

$A \mathop{\smile\!\!\!\!\!\!\downarrow}^E_C B$, if and only if $\widehat{A} \mathop{\smile\!\!\!\!\!\!\downarrow}_{\widehat{C}} \widehat{B}$ in the sense of models of rb$\mathbb{R}$T. That is, $A \mathop{\smile\!\!\!\!\!\!\downarrow}^E_C B$ if and only if for all $a \in \widehat{A}$ we have $\operatorname{dist}(a, \overline{E_{\widehat{BC}}}) = \operatorname{dist}(a, \overline{E_{\widehat{C}}})$.

4.4.6 Theorem. *The* $\mathop{\smile\!\!\!\!\!\!\downarrow}^E$ *independence relation is the model theoretic independence relation for* rbE$\mathbb{R}$T$^\pi_s$.

Proof. Again here as in Theorem 3.5.6, the general line of argument is close to that of proof of Theorem 2.6.7, but in places we need lemmas from this chapter instead of those used in the proof of 2.6.7. As in that proof we use [1, Theorem 14.12], and show $\mathop{\smile\!\!\!\!\!\!\downarrow}^E$ has all seven of the properties required by that theorem.

Invariance under automorphisms is clear because any automorphism of a model $\mathcal{M}$ of rbE$\mathbb{R}$T$^\pi_s$ preserves $\mathop{\smile\!\!\!\!\!\!\downarrow}$ and satisfies $\sigma(\widehat{A}) = \widehat{\sigma(A)}$. Symmetry and transitivity follow from the definition of $\mathop{\smile\!\!\!\!\!\!\downarrow}^E$, and from the fact that $\mathop{\smile\!\!\!\!\!\!\downarrow}$ is symmetric and transitive. Finite character follows from the fact that $\mathop{\smile\!\!\!\!\!\!\downarrow}$ has finite character, and from the fact that $\widehat{A} = \bigcup_{a \in A \cup \{p\}} \{f^m(a) \mid m \in \mathbb{Z}\}$. Extension has much the same proof as in the proof of 2.6.7, but using Lemma 4.3.10 instead of Lemma 2.5.3. For local character we use the same line of reasoning as in Theorem 2.6.7. The argument for stationarity uses Lemma 4.4.3 instead of Lemma 2.6.3. ☐

4.4.7 Theorem. *The theory* rbE$\mathbb{R}$T$^\pi_s$ *is* not *ω-categorical.*

Proof. Since $L_p \subseteq L_s$ are both countable signatures, and rb$\mathbb{R}$T is the reduct of rbE$\mathbb{R}$T$^\pi_s$ to L_p, this follows from the fact that rb$\mathbb{R}$T is not ω-categorical by [1, Proposition 12.13]. ☐

4.4.8 Theorem. *Let $\kappa > \omega$ be a cardinal. The theory* rbE$\mathbb{R}$T$^\pi_s$ *is* not *κ-categorical.*

Proof. Let $\kappa > \omega$ be a cardinal. We construct non-isomorphic models of rbE$\mathbb{R}$T$^\pi_s$, each with density character κ. First, let $\mathcal{M}$ be a separable model of rbE$\mathbb{R}$T$^\pi_s$ with underlying $\mathbb{R}$-tree and isometry (M, d, p, f). The degree of any point in a separable $\mathbb{R}$-tree is at most ω by Lemma 2.6.10. Pick a point $q \in M$ fixed by f. Use Lemma 2.2.8 to add κ-many new distinct rays at q, all of which are fixed pointwise by f, and let $\widehat{\mathcal{W}_1}$ be this extension of $\mathcal{M}$.

Next, we use the procedure in the beginning of the proof of Theorem 4.3.13 to extend $\widehat{\mathcal{W}_1}$ to a model $\mathcal{W}_1'$ of rbE$\mathbb{R}$T$^\pi_s$. Let $(n_i)_{i=1}^\infty$ be a sequence of positive integers such that every positive integer appears infinitely many times in the sequence. Let $\mathcal{M}_0 = \widehat{\mathcal{W}_1}$. For $i \geq 1$, let $\mathcal{M}_i$ be the result of applying Lemma 4.3.11 to M_{i-1} with $h = n_i$. This process creates an increasing chain $(\mathcal{M}_i \mid i \in \omega)$ of models of E$\mathbb{R}$T$_s$. Let the L_s-structure $\mathcal{W}_1'$ be the union of this chain, with underlying $\mathbb{R}$-tree and isometry (W_1', d, p, f). As in the proof of Theorem 4.3.13, it is straightforward to see that $\mathcal{W}_1' \models \mathrm{rbE}\mathbb{R}\mathrm{T}_s^\pi$. Moreover, at each of the countably many steps in the construction above, we added at most finitely many branches at any given point. Therefore, the degree of every point $x \neq q$ in W_1' is at most ω.

Let $A_1 \subset W_1'$ be a subspace consisting of q and the κ-many rays we added at q. Note that A_1 has density equal to κ. Apply the Downward Löwenheim-Skolem Theorem ([1, Proposition 7.3]) to $\mathcal{W}_1'$ to get an elementary substructure $\mathcal{W}_1$ of density character $\leq \kappa$ which contains A_1. Because $\mathcal{W}_1$ contains a point of degree κ, we know $\mathcal{W}_1$ has density character equal to κ by Lemma 2.6.10. Also, $\mathcal{W}_1 \models \mathrm{rbE}\mathbb{R}\mathrm{T}_s^\pi$, and q is the only point in the underlying $\mathbb{R}$-tree of degree κ, every other point has degree at most ω.

To find the other non-isomorphic model, let $\widehat{\mathcal{W}_2} \models \mathrm{rbE}\mathbb{R}\mathrm{T}_s^\pi$ be κ-saturated with underlying $\mathbb{R}$-tree and isometry $(\widehat{W_2}, d_2, p_2, g)$. Then by Lemma 2.3.7 there are at least κ-many branches at every point in $\widehat{W_2}$. Choose $a \in \widehat{W_2}$ different from the basepoint, and choose κ-many distinct rays at a. Choose κ-many distinct rays at the basepoint p_2. Let $A_2 \subset \widehat{W_2}$ be the subspace consisting of a, p_2 and the chosen rays. Apply the Downward Löwenheim-Skolem Theorem ([1, Proposition 7.3]) to $\widehat{\mathcal{W}_2}$ to get an elementary substructure $\mathcal{W}_2$ of density character $\leq \kappa$ which contains A_2. Then, $\mathcal{W}_2 \models \mathrm{rbE}\mathbb{R}\mathrm{T}_s^\pi$ and has density character equal to κ by Lemma 2.6.10. There are at least two distinct points (p_2 and a) with degree κ in the underlying $\mathbb{R}$-tree W_2 of $\mathcal{W}_2$. In W_1 only the point q has degree κ, while the rest of the points have at most degree ω. Thus, $\mathcal{W}_1$ and $\mathcal{W}_2$ cannot be isomorphic. $\square$

REFERENCES

[1] Ben Yaacov I., Berenstein A,, Henson C.W., Usvyatsov A., Model theory for metric structures. In *Model Theory with Applications to Algebra and Analysis*, Vol. II, eds. Z. Chatzidakis, D. Macpherson, A. Pillay, and A. Wilkie, Lecture Notes series of the London Mathematical Society, No. 350, Cambridge University Press, 2008, 315–427.

[2] Ben Yaacov I., Usvyatsov A., Continuous first order logic and local stability. To appear in Trans. Amer. Math. Soc.

[3] Ben Yaacov I., Continuous first order logic for unbounded metric structures, submitted.

[4] Ben Yaacov I., Modular functionals and perturbations of Nakano spaces. J. Logic and Analysis 1 (2009), 1–42.

[5] Bridson M.R., Haefliger A., *Metric Spaces of Non-positive Curvature.* Grundlehren der mathematischen Wissenschaften 319, Springer, 1999.

[6] Chiswell, I. *Introduction to Λ-trees.* World Scientific, Singapore, 2001.

[7] Culler, M., Morgan, J.W., Group actions on $\mathbb{R}$-trees. Proc. London Math. Soc. 55 (1987), 571–604.

[8] van den Dries L., Wilkie A.J., Gromov's theorem on groups of polynomial growth and elementary logic. J. Algebra 89 (1984), 349–374.

[9] Drutu, C., Quasi-isometry invariants and asymptotic cones. International Journal of Algebra and Computation 12 (2002), 99–135.

[10] Dyubina, A., Polterovich, I. Explicit constructions of universal $\mathbb{R}$-trees and asymptotic geometry of hyperbolic spaces. Bull. London Math. Soc. 33 (2001), 727–734.

[11] Hodges, W., *A Shorter Model Theory.* Cambridge University Press, Cambridge, 1997.

[12] Morgan, J.W., Shalen, P. Valuations, trees, and degenerations of hyberbolic structures. I. Annals of Math., 120 (1984), 401–476.

[13] Roe, J., *Lectures on Coarse Geometry.* University Lecture Series, 31. American Mathematical Society, Providence, RI, 2003.

AUTHOR'S BIOGRAPHY

Sylvia Ellen Booth Carlisle was born February 13th, 1980 in St. Louis, Missouri. She grew up mainly in Iowa City, Iowa, where in her freshman year her geometry teacher encouraged her to stop sitting at the back of the classroom and join the high school math club. Sylvia graduated from Carleton College in 2002 and moved on to the doctoral program in mathematics at the University of Illinois, Urbana-Champaign. After graduating, she will be a tenure-track assistant professor at Eastern Illinois University in Charleston, Illinois.