FAKULTÄT FÜR !NFORMATIK

Infinitary Rewriting
–
Theory and Applications

MASTERARBEIT

zur Erlangung des akademischen Grades

Master of Science
Computational Logic MSc

im Rahmen des Studiums

Computational Logic

eingereicht von

Patrick Bahr
Matrikelnummer 0827963

an der
Fakultät für Informatik der Technischen Universität Wien

Betreuer: Ao. Univ. Prof. Dr. Bernhard Gramlich

Wien, 24. September 2009 ________________ ________________
(Unterschrift Verfasser) (Unterschrift Betreuer)

Infinitary Rewriting
–
Theory and Applications

Patrick Bahr

September 24, 2009

Abstract

Infinitary rewriting generalises usual finitary rewriting by providing infinite reduction sequences with a notion of convergence. The idea of – at least conceptually – assigning a meaning to infinite derivations is well-known, for example, from lazy functional programming or from process calculi. Infinitary rewriting makes it possible to apply rewriting in order to obtain a formal model for such infinite derivations. The goal of this thesis is to comprehensively survey the field of infinitary term rewriting, to point out its shortcomings, and to try to overcome some of these shortcomings. The most significant problem that arises in infinitary rewriting is the inherent difficulty to finitely represent and, hence, to implement it. To this end, we consider term graph rewriting, which is able to finitely represent restricted forms of infinitary term rewriting. Moreover, we study different models that are used to formalise infinite reduction sequences: The well-established metric approach as well as an alternative approach using partial orders. Both methods together with the consequent infinitary versions of confluence and termination properties are analysed on an abstract level. Based on this, we argue that the partial order model has more advantageous properties and represents the intuition of convergence in a more natural way. This assessment is also backed up by the results that we obtain for infinitary term rewriting: Unlike the metric approach, the partial order approach admits to generalise some results known from finitary orthogonal term rewriting – most importantly, confluence. It is also shown that so-called Böhm trees, usually constructed rather intricately, naturally arise as normal forms in the partial order model. Finally, we devise a complete ultrametric and a complete semilattice on term graphs both of which are used to introduce infinitary term graph rewriting. This is supposed to serve as a tool in order to investigate the limitations of term graph rewriting for implementing infinitary term rewriting.

Zusammenfassung

Infinitäre Ersetzungssysteme erweitern finitäre Ersetzungssysteme durch das Einführen eines Konvergenzbegriffs für unendliche Reduktionsfolgen. Die Idee, auch unendlichen Ableitungen zumindest konzeptionell eine Bedeutung zuzuordnen, ist zum Beispiel von der nicht-strikten funktionalen Programmierung oder auch von Prozesskalkülen her bekannt. Infinitäre Ersetzung ermöglicht es, diese unendlichen Ableitungen durch Ersetzungssysteme zu formalisieren. Die vorliegende Arbeit verfolgt das Ziel, einen umfassenden Überblick über das Gebiet der infinitären Termersetzungssysteme darzulegen, dessen Defizite aufzuzeigen und zu versuchen einige dieser Defizite zu beheben. Das bedeutendste Problem, welches sich bei infinitären Systemen zeigt, ist die inhärente Schwierigkeit diese endlich darzustellen, um somit eine Implementierung zu ermöglichen. Zu diesem Zweck werden Termgraphersetzungssysteme betrachtet, welche es erlauben eingeschränkte Formen infinitärer Termersetzung endlich darzustellen. Darüber hinaus untersuchen wir verschiedene Modelle der infinitären Ersetzung: die etablierte metrische Methode sowie eine alternative Methode unter Zuhilfenahme von Halbordnungen. Beide Ansätze werden – zusammen mit den daraus resultierenden infinitären Varianten von Konfluenz- und Terminierungseigenschaften – auf abstrakter Ebene analysiert. Darauf aufbauend erörtern wir, dass das Halbordnungsmodell vorteilhaftere Eigenschaften besitzt und die Intuition von Konvergenz natürlicher wiedergeben kann. Diese Einschätzung wird ebenso von unseren Erkenntnissen belegt, die wir für infinitäre Termersetzungssysteme erhalten: Im Gegensatz zum metrischen Ansatz erlaubt es der Halbordnungsansatz einige Resultate zu verallgemeinern, die von finitären orthogonalen Systemen bekannt sind – insbesondere Konfluenz. Es wird außerdem gezeigt, dass sogenannte Böhm-Bäume, welche üblicherweise relativ aufwendig konstruiert werden müssen, als Normalformen im Halbordnungsmodell entstehen. Schließlich werden eine vollständige Ultrametrik und ein vollständiger Halbverband auf Termgraphen entwickelt, welche beide genutzt werden um infinitäre Termgraphersetzung einzuführen. Dies soll als Instrument dienen, um die Grenzen der Möglichkeiten der Termgraphersetzung für die Implementierung der infinitären Termersetzung zu untersuchen.

Contents

Contents i

List of Figures iii

1 Introduction 1
- 1.1 Reduction Systems . . . 1
- 1.2 Motivation . . . 2
- 1.3 Structure of the Thesis . . . 4
- 1.4 Main Contributions . . . 5

2 Preliminaries 7
- 2.1 Set Theory . . . 7
  - 2.1.1 Partial Orders . . . 8
  - 2.1.2 Ordinal Numbers and Sequences . . . 9
  - 2.1.3 Limits in Partial Orders . . . 12
- 2.2 Topology . . . 13
  - 2.2.1 Topological Spaces . . . 14
  - 2.2.2 Metric Spaces . . . 15
- 2.3 Reduction Systems . . . 18
  - 2.3.1 Abstract Reduction Systems . . . 18
  - 2.3.2 Terms . . . 20
  - 2.3.3 Term Rewriting Systems . . . 23

3 Transfinite Reductions 27
- 3.1 Metric Reduction Systems . . . 28
- 3.2 Partial Reduction Systems . . . 36
- 3.3 Transfinite Abstract Reduction Systems . . . 42
  - 3.3.1 Properties of Transfinite Reductions . . . 42
  - 3.3.2 Relating MRSs to PRSs . . . 45
- 3.4 Alternative Models of Transfinite Reductions . . . 47

4 Term Graphs 49
- 4.1 Graphs and Term Graphs . . . 50
- 4.2 Homomorphisms . . . 53
- 4.3 Canonical Term Graphs . . . 57
- 4.4 A Partial Order on Term Graphs . . . 62
- 4.5 A Metric on Term Graphs . . . 75

5 Infinitary Term Rewriting 87
- 5.1 Finitary Properties on Infinite Terms . . . 88
  - 5.1.1 Termination Properties . . . 88
  - 5.1.2 Confluence Properties . . . 89
- 5.2 MRS vs. PRS Model of Infinitary Term Rewriting . . . 90

5.3 Weakly Convergent MRS Reductions . . . 91
5.3.1 Compression and Approximation . . . 92
5.3.2 Confluence . . . 93
5.3.3 Connection to Strongly Convergent Reductions . . . 95
5.4 Strongly Convergent MRS Reductions . . . 96
5.4.1 Compression and Approximation . . . 97
5.4.2 Complete Developments . . . 98
5.4.3 Tiling Diagrams and Projections . . . 102
5.4.4 Confluence . . . 105
5.4.5 Meaningless Terms and Böhm Trees . . . 107
5.4.6 Reduction Strategies . . . 110
5.4.7 Termination . . . 114
5.5 Strongly Convergent PRS Reductions . . . 114
5.5.1 Descendants . . . 117
5.5.2 Complete Developments . . . 123
5.5.3 Relation to Böhm Trees . . . 133

6 Term Graph Rewriting **139**
6.1 Term Graph Rewriting Systems . . . 139
6.2 Simulating Term Rewriting . . . 148

7 Conclusions **157**
7.1 Summary . . . 157
7.2 Future Work and Possible Applications . . . 158

Bibliography **161**

Index **169**

Symbols **176**

List of Figures

1.1 Reduction sequences generating infinite list of natural numbers. 3
1.2 Cyclic term graph rewriting. 4

3.1 Weakly convergent reduction sequence. 31
3.2 Strongly convergent reduction sequence. 32
3.3 Reduction sequence with stable context. 39
3.4 Limits of the PRS reduction sequence. 39
3.5 Alternative definition of the infinitary normal form property. 43

4.1 Example for a tree representation of a term; generalisation to graphs. 50
4.2 Example for a replacement of a node by a term graph. 53
4.3 Isomorphism between Δ-homomorphisms. 57
4.4 Example of a $\perp$-homomorphism injective in non-$\perp$-nodes. 63
4.5 Counterexample for candiate partial orders. 64
4.6 Least upper bound $g \sqcup_\perp h$ of compatible term graphs g and h. 69
4.7 Example for truncation. 77
4.8 $\perp$-depth in truncated term graphs. 79
4.9 Fringe nodes and strong $\perp$-homomorphisms. 81

5.1 Semi-ω-confluence of Theorem 5.3.9. 94
5.2 Descendants by a single step. 99
5.3 Diamond property of complete developments according to Corollary 5.4.19. . . . 102
5.4 Tiling diagram. 103
5.5 The Infinitary Strip Lemma. 104
5.6 Semi-infinitary confluence according to Corollary 5.4.24. 104
5.7 Infinitary confluence modulo hypercollapsing terms. 106
5.8 Binary trees. 112
5.9 A parallel-outermost reduction. 113
5.10 Reduction sequence with volatile positions. 116
5.11 A path. 125
5.12 Turning a PRS reduction into a Böhm reduction. 134

6.1 Example term graph and term graph rewrite rule. 141
6.2 Alternative definition of term graph rewriting. 145
6.3 Weak completeness of term graph rewriting. 153
6.4 Example for weak completeness of term graph rewriting. 154
6.5 Necessity of (infinitary) CR for weak completeness of term graph rewriting. . . . 155
6.6 Infinitary term graph rewriting vs. infinitary term rewriting. 155

Chapter 1

Introduction

In this introduction we first sketch the background of this thesis, viz. reduction systems and their basic idea. Afterwards, we motivate the benefits of considering infinitely long reductions and show how they can be used for certain applications. Finally, a summary of the contents of the thesis is provided and the main contributions are briefly explained.

1.1 Reduction Systems

Reduction systems and in particular term rewriting systems are a well-established theory in computer science. The first such systems, the λ-calculi introduced by Church [Chu32], were originally conceived to serve as a logical foundation for mathematics. Later on, the main motivation for considering reduction systems was their ability to provide a formal model for computations.

Reduction systems usually consist of a set of rules defined on a particular set of objects, which in most cases consist of a language of terms. The rules of the system determine how the objects of the system can be rewritten. Each such rewrite step is usually supposed to model some kind of elementary computation step that is performed on some object.

A simple example is a system that operates on arithmetic expressions such as $(5+3)\cdot(2+3)$ and that has rules to simplify sums and products of two numbers. Such a system would, for example, contain the rules $5+3 \to 8$, $8 \cdot 5 \to 40$ etc. which state how sums and products can be replaced by a single number. This allows the following reduction:

$$\underline{(5+3)} \cdot (2+3) \to 8 \cdot \underline{(2+3)} \to \underline{8 \cdot 5} \to 40$$

The above reduction consists of three reduction steps. The underlining indicates the subexpressions to which rewrite rules are applied. There is no rule that can be applied to the numeral 40. Such *irreducible* objects are also called *normal forms*. They are usually considered to be the result of a computation.

Reduction systems can, however, be non-terminating, i.e. yielding infinite reduction sequences. If the above system would also contain rules which allow to employ the commutativity of addition and multiplication, viz. the rules $x + y \to y + x$ and $x \cdot y \to y \cdot x$, then the system would admit an infinite reduction

$$5 + 3 \to 3 + 5 \to 5 + 3 \to \dots$$

Such a reduction has no “final result”. Hence, the system is not terminating.

One of the most important classes of reduction systems is the class of term rewriting systems. These systems work on a term language by replacing subterms according to their rules. The system illustrated above is an example of a term rewriting system. The language it operates on consists of two binary symbols $+$ and $\cdot$ as well as a constant symbol n for each natural number $n \in \mathbb{N}$.

The rules of a term rewriting system can be seen as the program of the system. The execution of this program consists of performing rewriting steps according to these rules. Functional programming languages are based on this idea. In most cases, for a program to be useful, we need it to be terminating, i.e. not allowing infinite computations. Hence, in this setting, infinite reduction sequences as the one shown above should be avoided. In some functional programming languages, usually referred to as *lazy functional programming languages* (cf. [Lau93, Jos89]), non-termination is not necessarily problematic but is, in fact, desirable as it allows a more declarative style of programming. These functional programming languages are based on the idea of *lazy evaluation*, an evaluation strategy that only evaluates an expression when it is “needed” by some function.

We will see some examples in the next section which illustrate that infinite reductions may have a meaningful result.

1.2 Motivation

A computation that runs infinitely long is usually not desirable. One notable exception is the class of *reactive systems* [HP85] which are designed to interact with their environment and, thus, need not have a run time that is bounded a priori. The methods of infinitary term rewriting were successfully applied in this setting (e.g. in [IN95]). This is, however, not the subject of the thesis. Even in a setting in which a computation is performed in order to compute an output value for a given input value, infinitely long computations might still be adequate. A computation might try to approximate some value without reaching it in finite time but improving the accuracy of the approximation with each step. The various iterative algorithms to compute the transcendental number π are an example for this.

For a simple example in terms of reduction systems, consider the term rewriting system containing the rule $a \to f(a)$ that rewrites a subterm a to a term $f(a)$. This systems allows to construct an infinite reduction

$$a \to f(a) \to f(f(a)) \to f(f(f(a))) \to \dots$$

of increasingly large terms. In fact, the constant a cannot be reduced to a normal form.

Although the above reduction sequence does not terminate, i.e. it does not reach a normal form after finitely many steps, intuitively the reduction approximates the infinite term $f^\omega = f(f(f(f(\dots))))$ that consists of infinitely many f symbols. A finite prefix of this reduction is able to reach a term that coincides with f^ω up to an arbitrary depth. To put it in other terms: The reduction sequence *converges* to the term f^ω.

As another example consider the term rewrite rule $h(x) \to c(x, h(s(x)))$ and the induced infinite reduction

$$h(0) \to c(0, h(s(0))) \to c(0, c(s(0), h(s^2(0)))) \to c(0, c(s(0), c(s^2(0), h(s^3(0))))) \to \dots$$

This might seem a bit obscure. However, if we interpret the binary function symbol c as the list constructor operator “:”, and terms of the form $s^n(0)$ as natural numbers n, then the above reduction sequence looks like this:

$$h(0) \to 0 : h(1) \to 0 : 1 : h(2) \to 0 : 1 : 2 : h(3) \to \dots$$

The reduction produces increasingly large lists of natural numbers. The intuitive limit of that sequence is the infinite list of all natural numbers $0 : 1 : 2 : \dots = [0, 1, 2, \dots]$. By depicting the tree representation of the terms as illustrated in Figure 1.1, this can be seen more directly. In each step, the part of the tree that keeps unchanged grows. This growing stable part more and more resembles the tree representation of the infinite list $[0, 1, 2, \dots]$.

Yet, it does not need to end here. That is, a reduction does not have to stop after ω steps. Reconsider the rewrite rule $a \to f(a)$ and the infinite reduction sequence that converges to

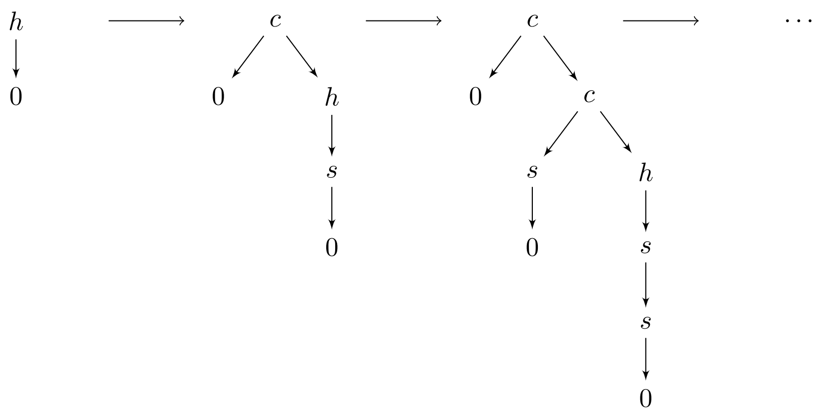

Figure 1.1: Reduction sequences generating infinite list of natural numbers.

the term f^ω. This term does not have to be a normal form. Suppose that there is an additional rule $f(x) \to g(x)$. Then we can continue the reduction like this:

$$a \to f(a) \to f(f(a)) \to \dots f^\omega \to g(f^\omega) \to g(g(f^\omega)) \to g(g(g(f^\omega))) \to \dots g^\omega$$

Let us return to the example of the system generating the infinite list $[0, 1, 2, \dots]$. Using functions that generate infinite data structures and in particular infinite lists is an ubiquitous technique in lazy functional programming. For example, in the lazy functional language Haskell, in order to write a function `number` that takes a list and numbers its elements by putting the i-th element e of the list in a pair (i, e), one can conveniently make use of infinite lists:

```
from n   = n : from (n + 1)
number l = zip (from 0) l
```

The function `from` is equivalent to the function h that generates infinite lists of natural numbers. The employed function `zip` takes two lists and returns a list whose i-th element is a pair consisting of the i-th element of the first list and the i-th element of the second list. Because `zip` stops as soon as the end of either argument lists is reached, we have to provide a list `[0,1,...,n-1]` as the first argument that has at least the same length as `l`. Since the argument list `l` can be of arbitrary size, we have to compute its length and then generate a list of appropriate length. Alternatively, we simply use an infinite list `[0,1,...]` which is generated by `from 0` in the same way $h(0)$ has generated such an infinite list in the example above. The fact that `from 0` does not terminate does not pose a problem in this setting. Provided the list `l` is finite, `zip` only "needs" a finite prefix of the infinite list `[0,1,...]`, and such a prefix can be obtained by only finitely often evaluating `from`.

The evaluation of the function `number` can also be seen in the sense of infinitary term rewriting: Before `zip` is applied to its two arguments they are evaluated to normal form. For `from 0`, this is a problem as it does not have a normal form in the finitary setting. However, it reaches a normal form after infinitely many steps, viz. the infinite list `[0,1,...]`. Then `zip` can be applied to these normal forms of its arguments. This yields the same result as the lazy evaluation. In this sense, infinitary term rewriting can be used as a natural model of lazy functional programming.

Rewriting can also be adopted for term graphs. Term graphs generalise terms by allowing to share common subterms. For example, the term $a(s(0), s(s(0)))$ can be represented by the term graph

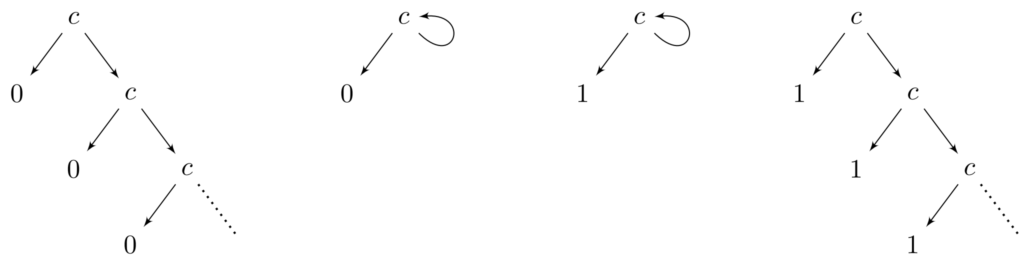

(a) Infinite list $[0, 0, \dots]$ (b) Term graph. (c) Reduct. (d) Infinite list $[1, 1, \dots]$

Figure 1.2: Cyclic term graph rewriting.

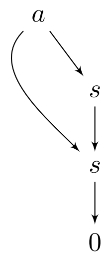

Term graph rewriting constitutes an important technique for efficiently implementing functional languages. In this setting, term graphs arise which contain cycles. Term graph rewriting is supposed to be an efficient implementation of term rewriting. However, as soon as cyclic graphs are present, a single term graph rewriting step might simulate an infinite number of term rewriting steps.

To see this, consider the infinite list $[0, 0, \dots]$ consisting of '0's. Again taking c as the list constructor, this list is represented by the term depicted in Figure 1.2a. This term can be represented by the term graph shown in Figure 1.2b. When the term graph is unravelled, we obtain the term $[0, 0, \dots]$. Now consider the rule $0 \to 1$. If this rule is applied to the term graph, then we obtain the term graph in Figure 1.2c which represents the infinite list $[1, 1, \dots]$ shown in Figure 1.2d. All '0's in the list were rewritten to '1's in one go. If we want to do this on the term side, i.e. if we want to rewrite the term $[0, 0, \dots]$ to the term $[1, 1, \dots]$ using the rule $0 \to 1$, we need infinitely many steps.

Hence, we need infinitary term rewriting to model this. But this can also be seen in the other direction: Term graph rewriting can serve as a tool to implement infinitary term rewriting since at least some infinite terms and some infinite reductions on them can be represented by finite term graphs and finite reductions on them.

1.3 Structure of the Thesis

Before we begin dealing with the topic of this thesis we provide, in Chapter 2, a brief summary of the fundamental mathematical tools that are needed for our presentation of infinitary rewriting. This includes basic set theory, order theory, topology and, most importantly, reduction systems. Topological spaces and partial orders will be of great importance as they are used to formalise transfinite reductions and their limits.

Subsequently, in Chapter 3, several models for formalising transfinite reductions are presented. The analysis of these models is performed on an abstract level. That is, we provide extensions to the notion of abstract reduction systems which allow to form limits of infinite reduction sequences. Two different approaches are investigated: The well-known method of endowing the system with a metric and a new method which uses a partial order.

Both approaches distinguish between a weak and a strong variant of transfinite reductions. We compare the resulting four different notions of transfinite rewriting, give criteria for which some of them coincide, and investigate their common properties. Moreover, transfinite versions of well-known finitary properties, such as termination and confluence properties, are introduced, and their interrelations are analysed.

In Chapter 4, term graphs are introduced together with several different notions of homomorphisms on them. The thus obtained theory is then used to introduce a metric and a partial order on term graphs which extend the usual metric resp. the usual partial order on terms. This provides the foundation for the two subsequent chapters, in particular Chapter 6.

Chapter 5 forms the core of this thesis. It provides an overview of the field of infinitary term rewriting. At first we discuss how the consideration of infinite terms changes the behaviour of (finitary) termination and confluence properties. Then it is shown that the partial order model provides a conservative extension to the well-known metric model of infinitary term rewriting. Subsequently, we summarise the most important results on infinitary term rewriting in the literature. This includes both the weak and the strong variant of transfinite reductions in the metric model. Finally, we investigate the properties of transfinite reductions in the partial order model and analyse in which cases the metric and the partial order approach yield equivalent reductions.

In Chapter 6, term graphs are reconsidered. We will present a notion of term graph rewriting and discuss some of its properties known from the literature. Moreover, we use the results of Chapter 4 in order to extend term graph rewriting such that it admits transfinite reductions. The focus is set on the ability of term graph rewriting to simulate term rewriting and, most importantly, infinitary term rewriting. This also includes an informal treatment of how infinitary term rewriting can be practically implemented by transforming a term rewriting system into a term graph rewriting system.

Finally, in Chapter 7, we conclude this thesis with a summary of its results and an outlook on possible future work.

1.4 Main Contributions

A large part of this thesis consists of presenting the current state of research in the field of infinitary rewriting. Nevertheless, we were also able to develop some new results:

- Kennaway [Ken92] developed an extension to abstract reduction systems which models transfinite reductions via metric spaces. In this thesis, a similar extension using *partial orders* is developed. We show that both approaches yield transfinite reduction sequences exhibiting a behaviour similar to that of finite reduction sequences, in particular in the way they can be composed. Additionally, we are able to show that the infinitary versions of confluence and termination properties are related to each other in the same way as they are in the finitary case. Finally, we provide a criterion which ensures that reductions described by the metric method are also reductions w.r.t. the partial order method. Both infinitary term rewriting systems and infinitary term graph rewriting systems are shown to meet this criterion.

- We investigate the newly introduced notion of strongly convergent reductions of term rewriting systems formalised by the partial order approach. The analysis is chiefly restricted to orthogonal systems. Our findings include:

  - orthogonal systems are *infinitary confluent*
  - orthogonal systems are *infinitary normalising*
  - orthogonal systems allow arbitrary *complete developments*
  - orthogonal systems have the *compression property*, i.e. reductions can always be performed in at most ω steps

We argue that this shows that strongly convergent reductions in the partial order approach are *more well-behaved* than corresponding reductions in the metric approach. Additionally, it is shown that reduction sequences do *always* weakly and strongly converge and that the metric model is *subsumed* by the partial order model. Furthermore, we show that partial order reductions are equivalent to metric reductions in the corresponding *Böhm reduction* system. Böhm reduction systems are term rewriting systems that are augmented by additional rules which assure that certain terms are identified by the system. This offers new insights both into Böhm reductions, including the entailed notion of Böhm trees, and into the partial order approach to infinitary term rewriting.

- We introduce a partial order and a metric on term graphs which extend the partial order and the metric on terms, respectively. Moreover, we show that the partial order forms a *complete semilattice* on term graphs and that the metric is, in fact, a *complete ultrametric*. Both results extend corresponding properties of the partial order and the metric on terms, respectively. This allows us to define *infinitary term graph rewriting* both in the partial order and the metric approach.

The positive results for partial order infinitary term rewriting indicate that it has more advantageous properties than infinitary term rewriting in the metric model. Since all reduction sequences do converge in the partial order model and the metric model is subsumed, we obtain a model of infinitary term rewriting which is superior to the well-known metric model. It also provides a more *fine-grained* view of convergence. Instead of the two possibilities – convergence or divergence – which the metric model provides, the partial order model allows to identify several *intermediary states* between full convergence and full divergence.

The equivalence of partial order infinitary term rewriting and metric infinitary term rewriting in the Böhm reduction illustrates the *fundamental difference* between the two models of convergence quite concisely. Moreover, it shows that the rather intricate construction of *Böhm trees* can also be performed quite naturally.

The application of the models of infinitary rewriting to the setting of term graph rewriting allows to compare term rewriting and term graph rewriting also in terms of their transfinite reductions. In the past, this was conducted in an informal way only. In addition, we hope that infinitary term graph rewriting can be employed as a tool for investigating which class of infinitary term rewriting can be finitely implemented using term graph rewriting. We also think that the partial order on term graphs can be used to generalise Böhm trees of term rewriting systems to *"Böhm graphs"* of term graph rewriting systems.

Chapter 2

Preliminaries

This chapter introduces the basic concepts, notations and theorems that are necessary for the investigations made in this thesis. This comprises – with few exceptions – well-known facts and notions of the respective branches of mathematics. Most of the material presented in this chapter can be found in the corresponding standard textbooks. At the beginning of each section we mention some of the relevant textbooks that can be consulted for a more detailed presentation. For some theorems that are mentioned here, the author was not able to find them in a textbook. In this case, an argument for the claimed property is given.

2.1 Set Theory

This section covers the basic concepts of set theory including partial orders and ordinal numbers. A detailed treatment of these subjects can be found in [Lev79] and [Sie65].

We use $\mathbb{N}^+$ to denote the set of positive integers $\{1, 2, \ldots\}$, $\mathbb{N}$ for the set of non-negative integers $\mathbb{N}^+ \cup \{0\}$, $\mathbb{R}$ for the set of all real numbers, $\mathbb{R}^+$ for the set of positive real numbers $\{r \in \mathbb{R} \,|\, r > 0\}$, and $\mathbb{R}^+_0$ for the set of non-negative real numbers $\mathbb{R}^+ \cup \{0\}$.

Let $f\colon A \to B$, $g\colon B \to C$ be two *functions*. We use $g \circ f$ to denote the *composition* of f and g, i.e. the function $(g \circ f)\colon A \to C$ with $(g \circ f)(a) = g(f(a))$ for all $a \in A$. Moreover, we write $\mathsf{range}(f)$ to denote the *range* of f, i.e. the set $\{f(a) \,|\, a \in A\}$. For the *identity function* on the set A, we write id_A, i.e. we have $\mathsf{id}_A(a) = a$ for all $a \in A$. A function $h\colon B \to A$ is called the *inverse* of f iff $h \circ f = \mathsf{id}_A$ and $f \circ h = \mathsf{id}_B$. There is at most one such inverse for each function. If it exists for a function f, then we use f^{-1} to denote it.

Let $f\colon A \rightharpoonup B$, $g\colon B \rightharpoonup C$ be two *partial functions*. The *domain* of f, written $\mathsf{dom}(f)$, is the set $\{a \in A \,|\, f(a) \text{ is defined}\}$. Accordingly, the *range* of f, denoted $\mathsf{range}(f)$, is the set $\{f(a) \,|\, a \in \mathsf{dom}(f)\}$. The *composition* of f and g, written as $g \circ f\colon A \rightharpoonup C$, is defined as $(g \circ f)(a) = g(f(a))$ for all $a \in \mathsf{dom}(f)$ with $f(a) \in \mathsf{dom}(g)$. For all other $a \in A$, the value $(g \circ f)(a)$ is not defined. Hence, $\mathsf{dom}(g \circ f) = \{a \in \mathsf{dom}(f) \,|\, f(a) \in \mathsf{dom}(g)\}$. Note that (partial) function composition is associative. A partial function $h\colon B \rightharpoonup A$ is called the *inverse* of f if we have $h(f(a)) = a$ for each $a \in \mathsf{dom}(f)$ and $f(h(b)) = b$ for each $b \in \mathsf{dom}(h)$. Also for partial functions, such an inverse, if existent, is unique and we use f^{-1} to denote it.

Whenever we have a function $f\colon A \to B$ and sets $A' \subseteq A$, $B' \subseteq B$, then we write $f(A')$ for the set $\{f(a) \,|\, a \in A'\}$ and $f^{-1}(B')$ for the set $\{a \in A \,|\, f(a) \in B'\}$.

Let $R, S \subseteq A \times A$ be two binary relations over a set A. We write aRb as a shorthand for $(a, b) \in R$. Moreover, we use $R \circ S$ to denote the composition of R and S, i.e. the binary relation over A with $R \circ S = \{(a, b) \,|\, \exists c \in A.\ aRc, cSb\}$. Just as for functions, the composition $\circ$ is associative. Additionally, we use the following definitions:

$$R^0 = \Delta_A = \{(a, a) \,|\, a \in A\}, \qquad R^+ = \bigcup_{n>0} R^n$$

$$R^{n+1} = R^n \circ R \qquad R^* = \bigcup_{n\geq 0} R^n$$

R^+ is known to be the transitive closure of R, i.e. the least transitive relation containing R; and R^* is known to be the reflexive, transitive closure of R, i.e. the least reflexive, transitive relation containing R. Also note that $R^1 = R$ and $R^{n+1} = R \circ R^n$.

For an *equivalence relationR* on a set A, i.e. a transitive, reflexive and symmetric relation, we use A/R to denote the *quotient* of A by R. That is, $A/R = \{[a]_R \mid a \in A\}$, where $[a]_R$ denotes the equivalence class $\{b \in A \mid aRb\}$ for each $a \in A$.

2.1.1 Partial Orders

In this section we give an overview over partial orders.

Definition 2.1.1 (ordered class)

(i) Let A be a class. A binary relation $< \subseteq A \times A$ is called a *partial order* (or simply *order*) on A if it is irreflexive and transitive.

(ii) A partial order $<$ on a class A is called *well-founded* if any non-empty subclass B of A has a minimal element $a \in B$ w.r.t. $<$, i.e. there is no $b \in B$ with $b < a$.

(iii) A partial order $<$ on a class A is called a *well-order* if it is *left-narrow*, i.e. $\{a \in A \mid a < b\}$ is a set for all $b \in A$, and every non-empty subclass B of A has a least element $a \in B$ w.r.t. $<$, i.e. $a < b$ for all $b \in B$.

(iv) A pair $(A, <)$ consisting of a class (resp. set) A and a partial order $<$ on it is called a *partially ordered class* (resp. *partially ordered set*). If $<$ is a well-order, then $(A, <)$ is called a *well-ordered class* (resp. *well-ordered set*). Usually, if $<$ is clear from the context, we simply refer to A as the partially ordered class.

Notation 2.1.2. Usually, we denote a partial order as $<$ (possibly with some index). In some cases it is convenient to refer to the reflexive closure of $<$ which we will denote by $\le$. This refers to the relation $\{(a, b) \in A \times A \mid a < b \text{ or } a = b\}$, which is then reflexive, antisymmetric and transitive.

Definition 2.1.3 (monotone functions)
Let $(A, <_A)$ and $(B, <_B)$ be two partially ordered classes. A function $f\colon A \to B$ is called *monotone*, if, for all $a, a' \in A$ with $a <_A a'$, also $f(a) <_B f(a')$ holds. If, additionally, f is surjective and, for all $a, a' \in A$ with $f(a) <_B f(a')$, also $a <_A a'$ holds, then f is called an *order isomorphism*. If such an order isomorphism exists, then $(A, <_A)$ and $(B, <_B)$ are called *order isomorphic* (or simply *isomorphic*), written $(A, <_A) \cong (B, <_B)$

Definition 2.1.4 (bounds, directed set)
Let $(A, <)$ be a partially ordered class and $B \subseteq A$ a subset.

(i) An *upper bound* resp. a *lower bound* of B is an element $a \in A$ with $b \le a$ resp. $a \le b$ for all $b \in B$.

(ii) An upper bound resp. a lower bound b of B is called *greatest element* resp. *least element* in B if $b \in B$.

(iii) Two elements $a, a' \in A$ are called *compatible* if the set $\{a, a'\}$ has an upper bound.

(iv) The *least upper bound* (*lub*) resp. *greatest lower bound* (*glb*) of B, denoted $\bigsqcup B$ resp. $\bigsqcap B$, is the least element of $\{a \in A \mid a \text{ upper bound of } B\}$ resp. the greatest element of $\{a \in A \mid a \text{ lower bound of } B\}$.

(v) B is called *directed* if it is non-empty and each pair of elements in B has an upper bound in B, i.e.

$$\forall a, b \in B \exists c \in B \quad a, b \le c.$$

Notation 2.1.5. We make use of some abbreviations for lubs and glbs: Instead of $\bigsqcup\{x, y\}$ we also write $x \sqcup y$, and instead of $\bigsqcup\{x_i \,|\, P(i)\}$, for some condition P on i, we also write $\bigsqcup_{P(i)} x_i$. Accordingly, we use the notation $x \sqcap y$ and $\bigsqcap_{P(i)} x_i$ for glbs.

Definition 2.1.6 (complete partial order, bounded complete partial order, complete semilattice)
Let $(<, A)$ be a partially ordered class.

(i) $<$ is called a *complete partial order* (*cpo*) if it has a least element and every directed subset of A has a least upper bound.

(ii) $<$ is called a *bounded complete partial order* (*bcpo*) if every subset of A that has an upper bound has a least upper bound.

(iii) $<$ is called a *complete semilattice* if it is a bounded complete cpo.

Note that, although the terminology may suggest it, cpos and bcpos are incomparable concepts. That is, neither is a cpo necessarily a bcpo nor is a bcpo necessarily a cpo. For instance, $\mathbb{N}$ endowed with the natural order $<$ constitutes a bcpo. Yet, it is not a cpo as the set $\mathbb{N}$ itself is directed but does not have a lub. On the other hand, the partial order given by the Hasse diagram

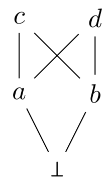

is a cpo but not a bcpo. It fails to be a bcpo since the set $\{a, b\}$ has an upper bound, viz. both c and d, however, it does not have a lub.

Proposition 2.1.7 (cpos and complete semilattices, [KP93])
A complete partial order is a complete semilattice iff every pair of compatible elements has a least upper bound.

Proposition 2.1.8 (bcpos admit glbs, [KP93])
In a bounded complete partial order, any non-empty set has a greatest lower bound.

2.1.2 Ordinal Numbers and Sequences

In this section we summarise the relevant notions for ordinal numbers and sequences.

Definition 2.1.9 (ordinal number)

(i) A class A is called *transitive* if every member of a member of A is a member of A, i.e. if $a \in A$ implies $a \subseteq A$.

(ii) A set α is called an *ordinal* if it is transitive and well-ordered by the relation $< = \{(a, b) \in \alpha \times \alpha \,|\, a \in b\}$, i.e. $a < b$ iff $a \in b$. We use On to denote the class of all ordinals.

Proposition 2.1.10 (ordinals are sets of ordinals)
Every member of an ordinal is an ordinal. That is, for each ordinal α, also every element $\beta \in \alpha$ is an ordinal.

Proposition 2.1.11 (On is a well-ordered proper class)

(i) $< = \{(a, b) \in \mathsf{On} \times \mathsf{On} \,|\, a \in b\}$ *is a well-order on* On.

(ii) On *is a proper class.*

Proposition 2.1.12 (construction of ordinals)

(i) If A is a set of ordinals, then $\bigcup A$ is an ordinal. $\bigcup A$ is the least upper bound of A. Thus every set of ordinals is bounded, and an unbounded class of ordinals is a proper class.

(ii) For every ordinal α, also $\alpha \cup \{\alpha\}$ is an ordinal and there is no β such that $\alpha < \beta < \alpha \cup \{\alpha\}$.

Notation 2.1.13. If α is an ordinal, we write $\mathsf{S}(\alpha)$ to denote its successor, the ordinal $\alpha \cup \{\alpha\}$. Moreover, we use 0 to denote the least ordinal $\varnothing$, 1 to denote its successor $\{0\}$, 2 for $\{0,1\}$ etc. ω is used to denote the least infinite ordinal $\{0,1,2,\dots\}$. Additionally, we use ω_1 to denote the least uncountable ordinal $\{\alpha \in \mathsf{On} \,|\, \alpha \text{ is countable}\}$.

Definition 2.1.14 (successor ordinal, limit ordinal)

(i) α is called a *successor ordinal* if there is an ordinal β such that $\alpha = \mathsf{S}(\beta)$.

(ii) α is called a *limit ordinal* if it is neither 0 nor a successor ordinal.

Note that $\bigcup \alpha$ is α' if $\alpha = \mathsf{S}(\alpha')$, and α if α is not a successor ordinal.

Lemma 2.1.15 (limit ordinals)
If α is a limit ordinal and $\beta < \alpha$, then there is an ordinal γ such that $\beta < \gamma < \alpha$; in particular $\beta < \mathsf{S}(\beta) < \alpha$.

Proposition 2.1.16 (lub of countably many countable ordinals is countable)
Let M be a countable subset of ω_1. Then $\bigcup M < \omega_1$. That is, the lub of a countable set of countable ordinals is again a countable ordinal.

Proposition 2.1.17 (principle of well-founded induction/recursion)
Let A be a class partially ordered by a left-narrow well-founded order $<$.

(i) Let B be a subclass of A. If, for each $b \in A$, it holds that $\{a \in A \,|\, a < b\} \subseteq B$ implies $b \in B$, then $A = B$.

(ii) Let G be a binary function on V, the class of all sets. Then there is a unique function F such that $F(a) = G(a, F|_{\{b \in A \,|\, b < a\}})$.

It is immediate from the definition that a well-order is also a left-narrow well-founded partial order. Since On is well-ordered, well-founded induction and well-founded recursion can also be applied to On or any ordinal. In this case we rather refer to these principles as *transfinite induction* resp. *transfinite recursion*. An instance of the transfinite recursion principle is the following definition of addition and multiplication on ordinals.

Definition 2.1.18 (ordinal arithmetic)

(i) Addition $+$ on On:

$$\begin{aligned} \alpha + 0 &= \alpha \\ \alpha + \mathsf{S}(\beta) &= \mathsf{S}(\alpha + \beta) \\ \alpha + \lambda &= \bigcup \{\alpha + \beta \,|\, \beta < \lambda\} \qquad \text{if } \lambda \text{ is a limit ordinal} \end{aligned}$$

(ii) Multiplication $\cdot$ on On:

$$\begin{aligned} \alpha \cdot 0 &= 0 \\ \alpha \cdot \mathsf{S}(\beta) &= (\alpha \cdot \beta) + \alpha \\ \alpha \cdot \lambda &= \bigcup \{\alpha \cdot \beta \,|\, \beta < \lambda\} \qquad \text{if } \lambda \text{ is a limit ordinal} \end{aligned}$$

Lemma 2.1.19 (properties of ordinal addition)

(i) $+$ *is associative, i.e.* $\alpha + (\beta + \gamma) = (\alpha + \beta) + \gamma$.

(ii) If $\alpha \leq \gamma$*, then there is a unique ordinal* β *such that* $\alpha + \beta = \gamma$*. Moreover,* β *is a limit ordinal iff* γ *is.*

Definition 2.1.20 (sequence)
Let α be an ordinal. A *sequence* S of *length* α (also called an α-*sequence*) in a set A, written $(a_\iota)_{\iota<\alpha}$, is a function from α to A with $\iota \mapsto a_\iota$ for all $\iota \in \alpha$. If α is a limit ordinal, S is called *open*. Otherwise, it is called *closed*. If α is a finite ordinal, S is called *finite*. Otherwise, it is called *transfinite* (or also *infinite*).

Notation 2.1.21. Let M be a set. We use the notations M^α, $M^{\leq\alpha}$ and $M^{<\alpha}$ to denote the set of all sequences in M of length α, $\leq \alpha$ and $< \alpha$, respectively. Instead of $M^{<\omega}$ we use the more common notation M^*. We use ε to denote the empty sequence. In order to refer to the length of a sequence S, we write $|S|$.

Definition 2.1.22 (concatenation of sequences)
Let $S = (a_\iota)_{\iota<\alpha}$, $T = (b_\iota)_{\iota<\beta}$ be two sequences in some set A. Then the sequence $S \cdot T$, called the *concatenation* of S and T, is the sequence $(a'_\iota)_{\iota<\delta}$, where $\delta = \alpha + \beta$ and $a'_\iota = a_\iota$ for $\iota < \alpha$ and $a'_{\alpha+\iota} = b_\iota$ for $\iota < \beta$.

From the definition, it is immediate that the concatenation of sequences is associative. Hence, we take the convenience of omitting brackets when concatenating several sequences. That is, we write $R \cdot S \cdot T$ for the sequence $(R \cdot S) \cdot T = R \cdot (S \cdot T)$.

Definition 2.1.23 (subsequence, segment, prefix, suffix)
Let $S = (a_\iota)_{\iota<\alpha}$ be an α-sequence. A β-sequence $T = (b_\iota)_{\iota<\beta}$ is called a *subsequence* of S if there is a monotone function $f\colon \beta \to \alpha$ such that $b_\iota = a_{f(\iota)}$. To indicate this, we write S/f for the subsequence T. If $f(\iota) = f(0) + \iota$ for all $\iota < \alpha$, then S/f is called a *segment* of S. Furthermore, such a segment S/f is called a *prefix* resp. a *suffix* of S if $f(0) = 0$ resp. if $f(0) + \beta = \alpha$. A subsequence S/f is called *proper* if f is not a surjection, i.e. S/f has a smaller length than S.

Two distinct sequences S, T are called *disjoint* if neither S is a prefix of T nor vice versa. Moreover, if S is a prefix of T, then T is called an *extension* of S.

Notation 2.1.24. Let $S = (a_\iota)_{\iota<\alpha}$ be an α-sequence and $\beta \leq \gamma < \alpha$. Then $S|_{[\beta,\gamma)}$ or equivalently $(a_\iota)_{\beta\leq\iota<\gamma}$ denotes the segment S/f, where $f\colon \alpha' \to \alpha$ is the mapping defined by $f(\iota) = \beta + \iota$ for all $\iota < \alpha'$, and α' is the unique ordinal with $\alpha = \beta + \alpha'$ (cf. Lemma 2.1.19).

When dealing with segments, the definition above is a bit cumbersome to work with. The following lemma provides more insight into the definition and makes it easier to reason about segments. It follows immediately from the definition of concatenation and the different forms of segments.

Lemma 2.1.25 (characterisation of segments)
Let S *and* S' *be two sequences.*

(i) S' *is a (proper) segment of* S *iff there are two sequences* R, T *such that* $S = R \cdot S' \cdot T$ *(and not both* R *and* T *are empty).*

(ii) S' *is a (proper) prefix of* S *iff there is a (non-empty) sequence* T *such that* $S = S' \cdot T$.

(iii) S' *is a (proper) suffix of* S *iff there is a (non-empty) sequence* R *such that* $S = R \cdot S'$.

Notation 2.1.26. For two sequences S and T, we write $S \leq T$ if S is a prefix of T, and $S < T$ if S is a proper prefix of T. When used in this way, we refer to $\leq$ and $<$ as the (strict) prefix order. One can easily see that the prefix order $\leq$ is indeed a partial order. Proposition 2.1.28 below shows that it even a complete semilattice.

With this knowledge of sequences we can generalise addition of ordinal numbers to arbitrary sequences of ordinal numbers.

Definition 2.1.27 (general addition of ordinals)
Let $(\beta_\iota)_{\iota<\alpha}$ be a sequence of ordinals. The sum of $(\beta_\iota)_{\iota<\alpha}$, written $\sum_{\iota<\alpha}\beta_\iota$, is defined as follows:

$$
\begin{aligned}
u\sum_{\iota<\alpha}\beta_\iota &= 0 && \text{if } \alpha = 0\\
\sum_{\iota<\alpha}\beta_\iota &= \sum_{\iota<\alpha'}\beta_\iota + \beta_{\alpha'} && \text{if } \alpha = \alpha' + 1\\
\sum_{\iota<\alpha}\beta_\iota &= \bigcup\left\{\sum_{\iota<\gamma}\beta_\iota \,\middle|\, \gamma < \alpha\right\} && \text{if } \alpha \text{ is a limit ordinal}
\end{aligned}
$$

To also generalise concatenation to arbitrary sequences of sequences, we need a means for defining limits on sequences. The following proposition shows that the prefix order does provide such a means.

Proposition 2.1.28 (prefix order on sequences is a complete semilattice)
The prefix order $\leq$ on sequences is a complete semilattice.

Proof. That $\leq$ is a partial order is obvious. In the following, we will show that $\leq$ is a cpo and that two compatible sequences have a lub. The proposition then follows from Proposition 2.1.7.

The empty sequence is the least element for $\leq$. Let D be a directed set of sequences (over the same set). We will show that D has a lub. Define $\alpha = \bigcup\{|S| \mid S \in D\}$. α is the length of the sequence $\overline{S}$ that we will construct. Define $\overline{S} = (a_\iota)_{\iota<\alpha}$ by choosing, for each $\gamma < \alpha$, some sequence $(b_\iota)_{\iota<\beta} \in D$ with $\beta > \gamma$ and setting $a_\gamma = b_\gamma$. One can easily check that $\overline{S}$ is well-defined and is indeed the lub of D. This shows that $\leq$ is a cpo.

Next consider two compatible sequences S_1 and S_2. That is, there is a sequence $\widehat{S}$ with $S_1, S_2 \leq \widehat{S}$. Hence, there are two sequences T_1 and T_2 with $S_1 \cdot T_1 = \widehat{S} = S_2 \cdot T_2$. One can easily see that then $S_1 \leq S_2$ or vice versa. In the former case, S_2 is the lub of $\{S_1, S_2\}$, in the latter case, it is S_1. □

Definition 2.1.29 (general concatenation of sequences)
Let $(S_\iota)_{\iota<\alpha}$ be a sequence of sequences over a common set. The concatenation of $(S_\iota)_{\iota<\alpha}$, written $\prod_{\iota<\alpha} S_\iota$, is defined as follows:

$$
\begin{aligned}
\prod_{\iota<\alpha} S_\iota &= \varepsilon && \text{if } \alpha = 0\\
\prod_{\iota<\alpha} S_\iota &= \prod_{\iota<\alpha'} S_\iota \cdot S_{\alpha'} && \text{if } \alpha = \alpha' + 1\\
\prod_{\iota<\alpha} S_\iota &= \bigsqcup\left\{\prod_{\iota<\gamma} S_\iota \,\middle|\, \gamma < \alpha\right\} && \text{if } \alpha \text{ is a limit ordinal}
\end{aligned}
$$

Remark 2.1.30. The general concatenation of sequences of sequences is well-defined in the limit ordinal case since the set $\left\{\prod_{\iota<\gamma} S_\iota \,\middle|\, \gamma < \alpha\right\}$ is directed as one can easily see. Proposition 2.1.28 then yields that the lub of this set exists.

2.1.3 Limits in Partial Orders

In this section we provide a means to define limits for sequences using a partial order.

Definition 2.1.31 (limit inferior)
Let $(A, \leq)$ be a partial order and $(a_\iota)_{\iota<\alpha}$ a non-empty α-sequence in A. The *limit inferior* of $(a_\iota)_{\iota<\alpha}$, written $\liminf_{\iota\to\alpha} a_\iota$, is defined as:

$$\liminf_{\iota\to\alpha} a_\iota = \bigsqcup_{\beta<\alpha}\left(\bigsqcap_{\beta\leq\iota<\alpha} a_\iota\right)$$

It is easy to see that the limit inferior of closed sequences is simply the last element of the sequence. This is, however, only a special case of the following more general proposition:

Proposition 2.1.32 (invariance of limit inferior)
Let $(a_\iota)_{\iota<\alpha}$ be a sequence in a partially ordered set and $(b_\iota)_{\iota<\beta}$ a non-empty suffix of $(a_\iota)_{\iota<\alpha}$. Then $\liminf_{\iota\to\alpha} a_\iota = \liminf_{\iota\to\beta} b_\iota$.

Proof. Since $\liminf_{\iota\to\alpha} a_\iota = \bigsqcup_{\gamma<\alpha}\bigsqcap_{\gamma\leq\iota<\alpha} a_\iota$, we have to show that

$$\overline{a} = \bigsqcup_{\gamma<\alpha}\bigsqcap_{\gamma\leq\iota<\alpha} a_\iota = \bigsqcup_{\beta\leq\gamma<\alpha}\bigsqcap_{\gamma\leq\iota<\alpha} a_\iota = \overline{a}' \quad \text{holds for each } \beta<\alpha.$$

Let $b_\gamma = \bigsqcap_{\gamma\leq\iota<\alpha} a_\iota$ for each $\gamma < \alpha$, $A = \{b_\gamma \,|\, \gamma < \alpha\}$ and $A' = \{b_\gamma \,|\, \beta \leq \gamma < \alpha\}$. Note that $\overline{a} = \bigsqcup A$ and $\overline{a}' = \bigsqcup A'$. Since $A' \subseteq A$, we have that $\overline{a}' \leq \overline{a}$. On the other hand, since $b_\gamma \leq b_{\gamma'}$ for $\gamma \leq \gamma'$, we find, for each $b_\gamma \in A$, some $b_{\gamma'} \in A'$ with $b_\gamma \leq b_{\gamma'}$. Hence, $\overline{a} \leq \overline{a}'$. Therefore, due to the antisymmetry of $\leq$, we can conclude that $\overline{a} = \overline{a}'$. □

For sequences of limit ordinal length, the limit inferior might not exist. The following proposition provides a sufficient condition for the partial order that guarantees the existence of the limit inferior.

Proposition 2.1.33 (limit inferior for complete semilattices)
The limit inferior is defined for any non-empty sequence in a complete semilattice.

Proof. Let $(A, \leq)$ be complete semilattice and $(a_\iota)_{\iota<\alpha}$ a non-empty sequence in A. As $\leq$ is also a cpo, it suffices to show that $D = \left\{\bigsqcap_{\beta\leq\iota<\alpha} a_\iota \,\middle|\, \beta < \alpha\right\}$ is a directed set. Let $a, b \in D$, i.e. there are $\beta', \beta'' < \alpha$ such that $a = \bigsqcap M_{\beta'}$ and $b = \bigsqcap M_{\beta''}$ for $M_\beta = \{a_\iota \,|\, \beta \leq \iota < \alpha\}$. Assume w.l.o.g. that $\beta' \leq \beta''$. Thus, $M_{\beta''} \subseteq M_{\beta'}$ and, consequently, $a \leq b$. A fortiori, a and b have an upper bound in D. □

Proposition 2.1.34 (limit inferior of open sequences)
Let $(A, \leq)$ be a complete semilattice and $(a_\iota)_{\iota<\lambda}$ an open sequence in A. Then it holds that $\liminf_{\iota<\lambda} a_\iota = \liminf_{\iota<\lambda}(a_\iota \sqcap a_{\iota+1})$.

Proof. Let $\overline{a} = \liminf_{\iota<\lambda} a_\iota$ and $\widehat{a} = \liminf_{\iota<\lambda}(a_\iota \sqcap a_{\iota+1})$. Since $a_\iota \sqcap a_{\iota+1} \leq a_\iota$ for each $\iota < \lambda$, we have $\widehat{a} \leq \overline{a}$. On the other hand, consider the sets $\overline{A}_\alpha = \{a_\iota \,|\, \alpha \leq \iota < \lambda\}$ and $\widehat{A}_\alpha = \{a_\iota \sqcap a_{\iota+1} \,|\, \alpha \leq \iota < \lambda\}$ for each $\alpha < \lambda$. Of course, we then have $\bigsqcap \overline{A}_\alpha \leq a_\iota$ for all $\alpha \leq \iota < \lambda$, and thus also $\bigsqcap \overline{A}_\alpha \leq a_\iota \sqcap a_{\iota+1}$ for all $\alpha \leq \iota < \lambda$. Hence, $\bigsqcap \overline{A}_\alpha$ is a lower bound of $\widehat{A}_\alpha$ which implies that $\bigsqcap \overline{A}_\alpha \leq \bigsqcap \widehat{A}_\alpha$. Consequently, $\overline{a} \leq \widehat{a}$ and, due to the antisymmetry of $\leq$, we can conclude that $\overline{a} = \widehat{a}$. □

2.2 Topology

In this section the basic notions of topological and in particular metric spaces are recapitulated. A recent textbook on this matter is [Mun00].

2.2.1 Topological Spaces

At first we give an overview over general topological spaces.

Definition 2.2.1 (topological space)

(i) A *topology* on a set X is a class $\mathcal{T}$ of subsets of X having the following properties:

(1) $\varnothing, X \in \mathcal{T}$

(2) If $A \subseteq \mathcal{T}$, then $\bigcup A \in \mathcal{T}$

(3) If $U, V \in \mathcal{T}$, then $U \cap V \in \mathcal{T}$.

(ii) A set X together with a topology $\mathcal{T}$ on it is called a *topological space*.

Definition 2.2.2 (open/closed set, neighbourhood, Hausdorff space)
Let X be a topological space with topology $\mathcal{T}$.

(i) A subset $U \subseteq X$ is called *open* resp. *closed* if $U \in \mathcal{T}$ resp. $U \smallsetminus X \in \mathcal{T}$. A set $V \subseteq X$ is called a *neighbourhood* of a point $p \in X$ if it contains an open set that contains p, i.e. $p \in U \subseteq V$ for some $U \in \mathcal{T}$.

(ii) X is called a *Hausdorff space* if, for each pair x_1, x_2 of distinct points in X, there exist neighbourhoods U_1 and U_2 of x_1 and x_2, respectively, that are disjoint.

The notion of a basis for a topology helps to define a topology on a set.

Definition 2.2.3 (basis for a topology)

(i) Let X be a set. A *basis* for a topology on X is a class $\mathcal{B}$ of subsets of X, called *basis elements*, such that

(1) For each $x \in X$, there is at least one basis element $B \in \mathcal{B}$ containing x.

(2) If x belongs to the intersection of two basis elements $B_1, B_2 \in \mathcal{B}$, then there is a basis element $B \in \mathcal{B}$ containing x such that $B \subseteq B_1 \cap B_2$.

(ii) If $\mathcal{B}$ is a basis for a topology on X, the topology $\mathcal{T}$ generated by $\mathcal{B}$ is defined as $\{\bigcup B \mid B \subset \mathcal{B}\}$.

Example 2.2.4 (ordinal spaces)
The concept of a basis can be used to conveniently define the standard topology on ordinals. Every ordinal α can be endowed with a topology $\mathcal{T}_\alpha$ that is generated by the basis $\mathcal{B}_\alpha = \{[\beta,\gamma) \mid \beta,\gamma \leq \alpha; \beta \text{ not a limit ordinal}\}$, where $[\beta,\gamma) = \{\delta \mid \beta \leq \delta < \gamma\}$. The resulting topological space is called an *ordinal space*.

Definition 2.2.5 (continuity)
Let X and Y be topological spaces. A function $f\colon X \to Y$ is called *continuous at point* $x \in X$ if, for each neighbourhood V of $f(x)$, there is a neighbourhood U of x such that $f(U) \subseteq V$. A function $f\colon X \to Y$ is called *continuous* if it is continuous at each point $x \in X$. In particular, if f is a sequence, i.e. X is an ordinal space, we call the sequence f a *continuous sequence*.

It is easy to see from this definition, that a sequence is continuous at any non-limit ordinal point. So, for checking continuity of sequences, only limit ordinal points have to be considered. This is summarised in the following lemma:

Lemma 2.2.6 (continuity of sequences)
Let $(x_\iota)_{\iota<\alpha}$ be a sequence in a topological space X. $(x_\iota)_{\iota<\alpha}$ is continuous iff, for each limit ordinal $\gamma < \alpha$ and each neighbourhood V of x_γ, there is an ordinal $\beta_V < \gamma$ such that $x_\iota \in V$ for each $\beta_V < \iota < \gamma$.

Proof. For the "if" direction, let $\gamma < \alpha$ and V be a neighbourhood of x_γ. To show that $(x_\iota)_{\iota<\alpha}$ is continuous, we have to provide a neighbourhood U of γ with $x_\iota \in V$ for all $\iota \in U$. Suppose γ is not a limit ordinal. Then $\{\gamma\}$ is an open set and we choose $U = \{\gamma\}$. If γ is a limit ordinal, we use the hypothesis and choose $U = \{\iota \in \alpha \,|\, \beta_V < \iota < \gamma\}$. For the "only if" direction, let $(x_\iota)_{\iota<\alpha}$ be a continuous sequence, i.e., for each $\gamma < \alpha$ and each neighbourhood V of x_γ, there is a neighbourhood U of γ such that $x_\iota \in V$ for all $\iota \in U$. We consider the particular case where γ is a limit ordinal. One can easily see that in this case the neighbourhood V of x_γ has to contain an open set $\{\iota \in \alpha \,|\, \beta_V < \iota < \gamma\}$ for some $\beta_V < \gamma$. □

Definition 2.2.7 (convergence)
Let X be a topological space and $(x_\iota)_{\iota<\alpha}$ a non-empty sequence in X. The sequence $(x_\iota)_{\iota<\alpha}$ *converges* to a point $x \in X$ if, for each neighbourhood U of x, there is a $\gamma < \alpha$ such that $x_\iota \in U$ for all $\gamma < \iota < \alpha$. Then x is called a *limit* of $(x_\iota)_{\iota<\alpha}$.

Note that $\mathsf{S}(\alpha)$-sequences $(a_\iota)_{\iota<\mathsf{S}(\alpha)}$ uniquely converge to the point a_α. The notion of convergence is only non-trivial for open sequences.

Proposition 2.2.8 (Hausdorff spaces admit at most one limit point)
If X is a Hausdorff space, then any sequence in X converges to at most one point.

Notation 2.2.9. If $(x_\iota)_{\iota<\alpha}$ is a sequence in a Hausdorff space, we write $\lim_{\iota\to\alpha} x_\iota = x$ whenever $(x_\iota)_{\iota<\alpha}$ converges to x.

In particular, every ordinal space is a Hausdorff space and, thus, admits at most one limit point.

Lemma 2.2.10 (concatenation of sequences)
Let X be a topological space and $S = (a_\iota)_{\iota<\alpha}$, $T = (b_\iota)_{\iota<\beta}$ two non-empty sequences in X. Then the following holds:

(i) $S \cdot T$ is continuous iff S and T are continuous, and S is closed or converges to b_0.

(ii) $S \cdot T$ converges to a iff T converges to a.

Proof. This follows immediately from Lemma 2.2.6 and Lemma 2.1.19. □

2.2.2 Metric Spaces

In this section we consider a particular class of topological spaces – metric spaces.

Definition 2.2.11 ((ultra-)metric space)

(i) Let M be a set. A function $\mathbf{d}\colon M \times M \to \mathbb{R}_0^+$ is called a *metric* on M if the following properties hold for all $x, y, z \in M$:

(1) $\mathbf{d}(x,y) = 0$ iff $x = y$ (identity)
(2) $\mathbf{d}(x,y) = \mathbf{d}(y,x)$ (symmetry)
(3) $\mathbf{d}(x,z) \le \mathbf{d}(x,y) + \mathbf{d}(y,z)$ (triangle inequality)

If $\mathbf{d}$ instead of (3) satisfies the stronger property

(3') $\mathbf{d}(x,z) \le \max\{\mathbf{d}(x,y), \mathbf{d}(y,z)\}$, (strong triangle)

it is called an *ultrametric*.

(ii) A *metric space* is a pair $(M, \mathbf{d})$, where M is a set and $\mathbf{d}$ is a metric on M. If $\mathbf{d}$ is even an ultrametric, $(M, \mathbf{d})$ is called an *ultrametric space*.

An example for a complete metric space that is not an ultrametric space is the set $\mathbb{R}$ of real numbers together with the usual metric $\mathbf{d}(x,y) = |x - y|$. Later we will see an example for an ultrametric space.

Definition 2.2.12 (metric topology)
Let $(M, \mathbf{d})$ be a metric space.

(i) Let $m \in M$ and $\varepsilon \in \mathbb{R}^+$. The *open ε-ball centred at m*, denoted $B_{\mathbf{d}}(m, \varepsilon)$ is defined as

$$B_{\mathbf{d}}(m, \varepsilon) = \{n \mid \mathbf{d}(m, n) < \varepsilon\}.$$

(ii) The class $\mathcal{B} = \{B_{\mathbf{d}}(m, \varepsilon) \mid m \in M, \varepsilon \in \mathbb{R}^+\}$ of open ε-balls is a basis for a topology on M. The generated topology $\mathcal{T}$ on M is called *metric topology.*

Metric topologies are Hausdorff, i.e., in particular, limits of sequences are unique if they exist. With the above definition of metric topologies the criterion for the continuity and for the convergence of sequences in a metric space can be formulated as follows:

Lemma 2.2.13 (continuity and convergence in metric spaces)
Let $(M, \mathbf{d})$ be a metric space and $(m_\iota)_{\iota<\alpha}$ a sequence in M.

(i) The sequence $(m_\iota)_{\iota<\alpha}$ converges to a point $m \in M$ iff, for each $\varepsilon \in \mathbb{R}^+$, there is a $\beta < \alpha$ such that $\mathbf{d}(m, m_\iota) < \varepsilon$ for every $\beta < \iota < \alpha$.

(ii) The sequence $(m_\iota)_{\iota<\alpha}$ is continuous iff, for each limit ordinal $\lambda < \alpha$ and $\varepsilon \in \mathbb{R}^+$, there is a $\beta < \lambda$ such that $\mathbf{d}(m_\lambda, m_\iota) < \varepsilon$ for every $\beta < \iota < \lambda$; in other words, for each limit ordinal $\lambda < \alpha$, the prefix $(m_\iota)_{\iota<\lambda}$ converges to m_λ.

Definition 2.2.14 (Cauchy sequence)
Let $(M, \mathbf{d})$ be a metric space and $(m_\iota)_{\iota<\alpha}$ a sequence in M. $(m_\iota)_{\iota<\alpha}$ is called a *Cauchy sequence* if, for any $\varepsilon \in \mathbb{R}^+$, there is a $\beta < \alpha$ such that, for all $\beta < \iota < \iota' < \alpha$, we have that $\mathbf{d}(m_\delta, m'_\delta) < \varepsilon$.

It is clear that every closed sequence is a Cauchy sequence. If one deals with ultrametric spaces, there is a simpler characterisation of open Cauchy sequences:

Lemma 2.2.15 (Cauchy sequence in ultrametric spaces)
Let $(M, \mathbf{d})$ be an ultrametric space. An open continuous sequence $(m_\iota)_{\iota<\alpha}$ is Cauchy iff $\lim_{\iota\to\alpha} \mathbf{d}(m_\iota, m_{\iota+1}) = 0$.

Proof. The "only if" direction follows straightforwardly from the definition of Cauchy sequences. For the "if" direction, suppose $(m_\iota)_{\iota<\alpha}$ is an open continuous sequence with $\lim_{\iota\to\alpha} \mathbf{d}(m_\iota, m_{\iota+1}) = 0$. That is, according to Lemma 2.2.13, for each $\varepsilon \in \mathbb{R}^+$, there is a $\beta_\varepsilon < \alpha$ such that

$$\mathbf{d}(m_\iota, m_{\iota+1}) < \varepsilon \quad \text{for all } \beta_\varepsilon < \iota < \alpha. \tag{1}$$

We will show that this implies $\mathbf{d}(m_\delta, m_\gamma) < \varepsilon$ for all $\beta_\varepsilon < \delta < \gamma < \alpha$, which proves that $(m_\iota)_{\iota<\alpha}$ is Cauchy. We will proceed by transfinite induction on $\gamma > \delta$.

If $\gamma = \delta + 1$, then $\mathbf{d}(m_\delta, m_\gamma) < \varepsilon$ follows from (1). If $\gamma = \gamma' + 1 > \delta$, then we get $\mathbf{d}(m_\delta, m_{\gamma'}) < \varepsilon$ by applying the induction hypothesis. From (1), we obtain $\mathbf{d}(m_{\gamma'}, m_\gamma) < \varepsilon$. By combining these inequalities using the stronger triangle inequality of the ultrametric, we get

$$\mathbf{d}(m_\delta, m_\gamma) \le \max\{\mathbf{d}(m_\delta, m_{\gamma'}), \mathbf{d}(m_{\gamma'}, m_\gamma)\} < \varepsilon.$$

Let $\gamma > \delta$ be a limit ordinal. Since $(m_\iota)_{\iota<\alpha}$ is continuous, according to Lemma 2.2.13, there is a $\delta' < \gamma$ such that $\mathbf{d}(m_\gamma, m_\iota) < \varepsilon$ for each $\delta' < \iota < \gamma$. In particular, this holds for some $\iota > \delta$. Additionally, we can apply the induction hypothesis to get $\mathbf{d}(m_\delta, m_\iota) < \varepsilon$. These two inequalities can be combined by the triangle inequality and the symmetry of $\mathbf{d}$:

$$\mathbf{d}(m_\delta, m_\gamma) \le \max\{\mathbf{d}(m_\delta, m_\iota), \mathbf{d}(m_\iota, m_\gamma)\} = \{\mathbf{d}(m_\delta, m_\iota), \mathbf{d}(m_\gamma, m_\iota)\} < \varepsilon.$$

□

From the definition of Cauchy sequences, it is easy to see that the following holds:

Fact 2.2.16 (convergences implies Cauchy)
Every converging sequence in a metric space is Cauchy.

Then the following property of continuous sequences in metric spaces is obvious:

Corollary 2.2.17 (continuous sequences in metric spaces)
For every continuous sequence $(m_\iota)_{\iota<\alpha}$ *in a metric space,* $\lim_{\iota\to\lambda} \mathbf{d}(m_\iota, m_{\iota+1}) = 0$ *holds for each limit ordinal* $\lambda < \alpha$.

Proof. This follows immediately from Lemma 2.2.13 and Fact 2.2.16. □

Definition 2.2.18 (completeness of a metric space)
Let $(M, \mathbf{d})$ be a metric space.

(i) M is called α-complete if every Cauchy sequence of length α is converging.

(ii) M is called complete if it is α-complete for each ordinal $\alpha \neq 0$.

Note that in the literature mostly ω-sequences are considered and, therefore, the notion of ω-completeness is simply called completeness in these contexts. But this will cause no confusion since both notions actually coincide:

Lemma 2.2.19 (completeness)
Let $(M, \mathbf{d})$ *be a metric space.*

(i) M *is* α*-complete for every successor ordinal* α.

(ii) M *is* ω*-complete iff* M *is complete.*

Proof. (i) Trivial.

(ii) The "if" direction is trivial. For the "only if" direction, let $S = (a_\iota)_{\iota<\alpha}$ be a non-empty Cauchy sequence over an ω-complete metric space $(M, \mathbf{d})$. We have to show that S converges. If α is a successor ordinal, this is trivial. If, on the other hand, α is a limit ordinal, consider an ω-sequence $(\varepsilon_i)_{i<\omega}$ in $\mathbb{R}^+$ that converges to 0, e.g. $\varepsilon_i = \frac{1}{i}$. Since S is a Cauchy sequence, for each ε_i with $i < \omega$, there is an ordinal $\gamma_i < \alpha$ such that

$$\mathbf{d}(a_\delta, a_{\delta'}) < \varepsilon_i \quad \text{for all } \gamma_i < \delta, \delta' < \alpha. \tag{1}$$

Since α is a limit ordinal, Lemma 2.1.15 allows us to choose the γ_i's such that $f\colon \omega \to \alpha$ given by $i \mapsto \gamma_i$ is a monotone function. Consider the subsequence $S' = S/f$. In order to see that this is a Cauchy ω-sequence, let $\varepsilon \in \mathbb{R}^+$. Since $(\varepsilon_i)_{i<\omega}$ converges to 0, there is some $n < \omega$ with $\varepsilon_n \leq \varepsilon$. Because, according to (1), $\mathbf{d}(a_\delta, a_{\delta'}) < \varepsilon_n \leq \varepsilon$ for all $\gamma_n < \delta, \delta' < \alpha$ and because of the monotonicity of f, we have

$$\mathbf{d}(a'_i, a'_j) = \mathbf{d}(a_{f(i)}, a_{f(j)}) < \varepsilon_n \leq \varepsilon \quad \text{for all } n < i, j < \omega.$$

Hence, $S' = (a'_i)_{i<\omega}$ is a Cauchy ω-sequence and, thus, converges to some $a \in M$ by the ω-completeness of M. That is, for each $\varepsilon \in \mathbb{R}^+$, there is some $n_\varepsilon < \omega$ such that

$$\mathbf{d}(a, a'_i) < \varepsilon \quad \text{for all } n_\varepsilon < i < \omega. \tag{2}$$

We will show that also S converges to a. To this end, let $\varepsilon \in \mathbb{R}^+$. Then, by (2), we have

$$\mathbf{d}(a, a'_i) < \frac{\varepsilon}{2} \quad \text{for all } n_{\frac{\varepsilon}{2}} < i < \omega. \tag{3}$$

Since $(\varepsilon_i)_{i<\omega}$ converges to 0, there is some $n_{\frac{\varepsilon}{2}} < k < \omega$ such that $\varepsilon_k \leq \frac{\varepsilon}{2}$. Consequently, as an instance of (1) we have

$$\mathbf{d}(a_\delta, a_{\delta'}) < \varepsilon_k \leq \frac{\varepsilon}{2} \quad \text{for all } f(k) < \delta, \delta' < \alpha$$

In particular, we get

$$\mathbf{d}(a_{f(k+1)}, a_\delta) < \frac{\varepsilon}{2} \quad \text{for all } f(k) < \delta < \alpha \tag{4}$$

As $k + 1 > n_{\frac{\varepsilon}{2}}$, we obtain from (3) that

$$\mathbf{d}(a, a_{f(k+1)}) = \mathbf{d}(a, a'_{k+1}) < \frac{\varepsilon}{2}. \tag{5}$$

By combining (4) and (5) using the triangle inequality, we get

$$\mathbf{d}(a, a_\delta) \leq \mathbf{d}(a, a_{f(k+1)}) + \mathbf{d}(a_{f(k+1)}, a_\delta) < \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon.$$

Consequently, S converges to a as well. □

Definition 2.2.20 (dense set)
Let $(M, \mathbf{d})$ be a metric space. A subset $N \subseteq M$ is called *dense* if

$$M = \left\{ \lim_{i \to \omega} m_i \,\middle|\, m_i \in N \text{ for all } i < \omega \right\}.$$

That is, M is the set of limits of all convergent ω-sequences in N.

For any metric space M, one can construct a complete metric space $\overline{M}$ that contains M as a dense subspace. $\overline{M}$ is unique up to isomorphism. It is called the *completion* of M.

2.3 Reduction Systems

In this section the most fundamental notions and results in the field of reduction systems, most importantly term rewriting systems, are collected in order to set the preparations for the investigations in this thesis. For more details, we recommend consulting [BN98] and [Ter03].

2.3.1 Abstract Reduction Systems

At first we consider abstract reduction systems.

Definition 2.3.1 (abstract reduction system)
An *abstract reduction system* (*ARS*) is a quadruple $\mathcal{A} = (A, \Phi, \mathsf{src}, \mathsf{tgt})$ consisting of a set of objects A, a set of steps Φ, and the *source* and *target functions* $\mathsf{src}\colon \Phi \to A$ and $\mathsf{tgt}\colon \Phi \to A$, respectively. We write $\varphi\colon a \to_{\mathcal{A}} b$ whenever there are $\varphi \in \Phi$, $a, b \in A$ such that $\mathsf{src}(\varphi) = a$ and $\mathsf{tgt}(\varphi) = b$. Usually, the name of the involved step is not needed to be explicitly indicated in which case we simply write $a \to_{\mathcal{A}} b$. If the ARS is clear from the context, we drop the explicit reference to it and write $a \to b$ instead.

This definition is rather technical and, for most purposes, it is too cumbersome to define the components $\Phi, \mathsf{src}, \mathsf{tgt}$ explicitly. Therefore, we allow to "reverse" the notation that we have introduced above to use it also to define an ARS. We want to be able to define an ARS by giving a collection of statements of the form $a \to b$, where a, b are objects of the ARS. The set Φ is then defined to be the set of these statements $a \to b$, i.e. pairs of objects, and we define $\mathsf{src}(a \to b) = a$ and $\mathsf{tgt}(a \to b) = b$.

Remark 2.3.2. The usual definition of an ARS comprises a pair consisting of a set of objects and a binary relation on this set. The reason for choosing a finer structured system becomes clearer when we extend this concept to *transfinite abstract reduction systems* in Chapter 3 and even more so when we extend the semantics of (infinitary) term rewriting systems to *transfinite abstract reduction systems*. To this end, consult Remark 3.1.5 for a motivation of this definition.

Definition 2.3.3 (reduction sequence)
An α-sequence $S = (\varphi_\iota)_{\iota<\alpha}$ of reduction steps in an ARS $\mathcal{A}$ is called an *α-reduction sequence* (or simply *reduction sequence*) if there is a sequence of objects $(a_\iota)_{\iota<\alpha'}$ in the underlying set A, where $\alpha' = \alpha$ if S is open, $\alpha' = \alpha+1$ if S is closed, such that $\varphi_\iota\colon a_\iota \to a_{\iota+1}$ for all $\iota < \alpha$. For such a sequence, we also write $(\varphi\colon a_\iota \to a_{\iota+1})_{\iota<\alpha}$ or simply $(a_\iota \to a_{\iota+1})_{\iota<\alpha}$. The reduction sequence S is said to start from a whenever $a = a_0$. In some cases we are not interested in the actual sequence of steps of a reduction sequence, in which case we will simply call it a *reduction*.

Remark 2.3.4. For now, reduction sequences of length greater than ω are not meaningful. One can easily see that, for example, the ω-th step in such a sequence is not related to the preceding steps of the reduction sequence. This holds in general for all reduction steps indexed by a limit ordinal. For successor ordinals, this is not a problem as by the above definition the $(\beta+1)$-st step is required to start in the object that the β-th step ends in. Meaningful definitions for reduction sequences of length beyond ω have to include a notion of continuity which bridges the gaps caused by limit ordinals. A variety of different approaches to such a notion is given in Chapter 3.

Notation 2.3.5. Let $\mathcal{A}$ be an ARS, and $S = (a_i \to a_{i+1})_{i<n}$ a finite reduction sequence in $\mathcal{A}$ of length n. To indicate this fact, we write $S\colon a \to^*_{\mathcal{A}} a_n$. For the special case of $n = 1$ and $n > 0$, we write $S\colon a_0 \to_{\mathcal{A}} a_n$ resp. $S\colon a_0 \to^+_{\mathcal{A}} a_n$. On some occasions it is convenient to "reverse" this notation in which case we will use $S\colon a_n \leftarrow^*_{\mathcal{A}} a_0$, $S\colon a_n \leftarrow_{\mathcal{A}} a_0$, and $S\colon a_n \leftarrow^+_{\mathcal{A}} a_0$, respectively. If $\mathcal{A}$ is clear from the context, it is omitted from the notation. If the name of the sequence is not of relevance, it is omitted as well.

Whenever $a \to^*_{\mathcal{A}} b$ we say that a *reduces* to b in $\mathcal{A}$ or that b is a *reduct* of b in $\mathcal{A}$.

In some contexts the direction, in which reductions are performed, is irrelevant. In this case we will consider the symmetric closure of an ARS:

Definition 2.3.6 (symmetric closure of an ARS)
Let $\mathcal{A} = (A, \Phi, \mathsf{src}, \mathsf{tgt})$ be an ARS. The *symmetric closure* $\mathcal{A}^s = (A, \Phi^s, \mathsf{src}^s, \mathsf{tgt}^s)$ of $\mathcal{A}$ is constructed by adjoining, for each step φ in $\mathcal{A}$, a symmetric copy φ' of it. More specifically, $\Phi^s = \varphi \uplus \{\varphi' \mid \varphi \in \Phi\}$; and $\mathsf{src}^s(\varphi) = \mathsf{src}(\varphi)$, $\mathsf{tgt}^s(\varphi) = \mathsf{tgt}(\varphi)$, $\mathsf{src}^s(\varphi') = \mathsf{tgt}(\varphi)$, and $\mathsf{tgt}^s(\varphi') = \mathsf{src}(\varphi)$ for all $\varphi \in \Phi$. That is, whenever $\varphi\colon a \to_{\mathcal{A}} b$, then we have both $\varphi\colon a \to_{\mathcal{A}^s} b$ and $\varphi'\colon b \to_{\mathcal{A}^s} a$.

Notation 2.3.7. Usually, we will use the symmetric closure $\mathcal{A}^s$ of an ARS $\mathcal{A}$ only implicitly. A finite reduction sequence in $\mathcal{A}^s$ is called a *finite conversion sequence* in $\mathcal{A}$. And, referring to Notation 2.3.5, we will use the notations $S\colon a \leftrightarrow^*_{\mathcal{A}} b$, $S\colon a \leftrightarrow_{\mathcal{A}} b$, and $S\colon a \leftrightarrow^+_{\mathcal{A}} b$ instead of $S\colon a \to^*_{\mathcal{A}^s} b$, $S\colon a \to_{\mathcal{A}^s} b$, and $S\colon a \to^+_{\mathcal{A}^s} b$, respectively.

Notation 2.3.8. To make statements about one or more ARSs more concise, we will further shorten the notation by using $\to_{\mathcal{A}}$ (or simply $\to$) to refer to the relation

$$\{(a,b) \in A \times A \mid a \to_{\mathcal{A}} b\}.$$

The same is done for the other "arrow notations" such as $\to^*_{\mathcal{A}}$, $\leftrightarrow_{\mathcal{A}}$ etc. One can then easily verify that $\to^+_{\mathcal{A}}$ is the *transitive closure*, $\to^*_{\mathcal{A}}$ the *reflexive transitive closure*, and $\leftarrow_{\mathcal{A}}$ the *reverse* of $\to_{\mathcal{A}}$. Furthermore, $\leftrightarrow_{\mathcal{A}}$ is its *symmetric closure*, $\leftrightarrow^+_{\mathcal{A}}$ its *symmetric transitive closure*, and $\leftrightarrow^*_{\mathcal{A}}$ its *reflexive symmetric transitive closure*.

Definition 2.3.9 (properties of ARSs)
Let $\mathcal{A} = (A, \Phi, \mathsf{src}, \mathsf{tgt})$ be an ARS.

(i) Every element $a \in A$ is called a *normal form* of $\mathcal{A}$ if there is no $b \in A$ such that $a \to b$. We use $\mathsf{NF}_{\mathcal{A}}$ to denote the set of all normal forms of $\mathcal{A}$.

(ii) $\mathcal{A}$ is *confluent* (CR) if $\leftarrow^* \circ \to^* \subseteq \to^* \circ \leftarrow^*$.

(iii) $\mathcal{A}$ has the *diamond property* (DP) if $\leftarrow \circ \rightarrow \subseteq \rightarrow \circ \leftarrow$.

(iv) $\mathcal{A}$ is *terminating* (SN) if there is no ω-reduction sequence in $\mathcal{A}$.

(v) $\mathcal{A}$ is *normalising* (WN) if every element in A reduces to a normal form.

(vi) $\mathcal{A}$ is *complete* (COMP) if it is both confluent and terminating.

(vii) $\mathcal{A}$ has the *normal form property* (NF) if $a \leftrightarrow^* b$ implies $a \rightarrow^* b$ for all $a \in A$ and any normal form $b \in A$.

(viii) $\mathcal{A}$ has the *unique normal form property* (UN) if $a \leftrightarrow^* b$ implies $a = b$ for all normal forms $a, b \in A$.

(ix) $\mathcal{A}$ has the *unique normal form property w.r.t. reduction* ($\mathsf{UN}_{\rightarrow}$) if $a \leftarrow^* c \rightarrow^* b$ implies $a = b$ for all normal forms $a, b \in A$ and all $c \in A$.

Notation 2.3.10. For a property P of an ARS, we take the convenience of writing $P(\mathcal{A})$ if we want to indicate that the ARS has the property P. Moreover, if P is a universal property, i.e. a property that states that all elements should have a certain property P', we choose to write $P(\mathcal{A}, a)$ whenever the object a of $\mathcal{A}$ satisfies the property P'. Moreover, if we have two universal properties P_1, P_2, we write $P_1 \Longrightarrow P_2$ to indicate that, for any ARS $\mathcal{A}$ and any object a in $\mathcal{A}$, we have that $P_1(\mathcal{A}, a)$ implies $P_2(\mathcal{A}, a)$. Note, that in this case we also have that $P_1(\mathcal{A})$ implies $P_2(\mathcal{A})$.

2.3.2 Terms

In this section we provide an overview over finite and infinite terms.

Definition 2.3.11 (finite terms)

(i) A *signature* is a countable set Σ of function symbols each of which is associated with a natural number denoting its arity, i.e. the number of its arguments. $\Sigma^{(n)}$ denotes the set of function symbols in Σ of arity n. We use $\mathsf{ar}(f)$ as a shorthand to denote the arity of the symbol f in Σ, i.e. $\mathsf{ar}(f) = k$ iff $f \in \Sigma^{(k)}$.

(ii) The set $\mathcal{T}(\Sigma, \mathcal{V})$ of *finite terms* over a signature Σ and some (usually countably infinite) set $\mathcal{V}$ of variables with $\Sigma \cap \mathcal{V} = \varnothing$ is the smallest set T containing $\mathcal{V}$ such that $f(t_1, \ldots, t_n) \in T$ whenever $f \in \Sigma^{(n)}$ and $t_1, \ldots, t_n \in T$. $\mathcal{T}(\Sigma) = \mathcal{T}(\Sigma, \varnothing)$ denotes the set of *finite ground terms*, i.e. terms without variables.

In fact, to be able to have an appropriate concept of transfinite reductions over terms, we need to extend the notion of terms. The set of finite terms is defined inductively. If we turn this definition into a coinductive one, it turns out, as we will show later, that we get precisely the desired notion of infinite terms that we will need for our purposes.

Definition 2.3.12 (infinitary terms)
The set $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ of *infinitary terms* over a signature Σ and some set $\mathcal{V}$ of variables with $\Sigma \cap \mathcal{V} = \varnothing$ is the greatest set T such that, for each element $t \in T$, we either have $t \in \mathcal{V}$ or $t = f(t_1, \ldots, t_k)$, where $k \geq 0$, $f \in \Sigma^{(k)}$, and $t_1, \ldots, t_k$. $\mathcal{T}^\infty(\Sigma) = \mathcal{T}^\infty(\Sigma, \varnothing)$ denotes the set of *infinitary ground terms*, i.e. infinitary terms without variables. Usually, as we are mostly interested in the general setting of infinitary terms, we omit the adjective "infinitary" and simply say set of *terms* or set of *ground terms*, respectively.

Remark 2.3.13. The duality between finite and infinitary terms becomes more explicit if one considers the category-theoretic formulations of inductive and coinductive definitions.

The definition of the set of finite terms $\mathcal{T}(\Sigma, \mathcal{V})$ as given in Definition 2.3.11 is equivalent to the definition of the initial $F_{\Sigma,\mathcal{V}}$-algebra, where the functor $F_{\Sigma,\mathcal{V}}$: **Set** → **Set** is defined by

$$X \mapsto \coprod_{n \geq 0} (\Sigma^{(n)} \times X^n) + V$$

On the other hand, the set of infinitary terms $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ as defined in Definition 2.3.12 can equivalently be defined as the final $F_{\Sigma,\mathcal{V}}$-coalgebra. By abuse of notation, we will identify the final $F_{\Sigma,\mathcal{V}}$-coalgebra with its underlying set $\mathcal{T}^\infty(\Sigma, \mathcal{V})$.

Note that not all elements of the set of infinitary terms are infinite. In fact, it is a superset of the set of finite terms, i.e. $\mathcal{T}(\Sigma, \mathcal{V}) \subseteq \mathcal{T}^\infty(\Sigma, \mathcal{V})$. This inclusion is strict iff Σ contains at least one non-nullary symbol or $\mathcal{V}$ is not empty. Thus, subsequently, we will define several concepts for the set of infinitary terms $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ which will then also apply to the set of finite trees $\mathcal{T}(\Sigma, \mathcal{V})$. However, we have to be careful since the usual technique of defining concepts inductively on the structure of terms does not work as we are also dealing with infinite terms. This is explicitly taken care of in the following definition of positions in terms.

Definition 2.3.14 (position, subterm)
Let $t \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$ be a term.

(i) A *position* is a finite sequence over $\mathbb{N}$. The *set of positions* in the term t, denoted $\mathcal{P}(t)$, is defined by

$$\pi \in \mathcal{P}(t) \quad :\Longleftrightarrow \quad \begin{cases} \pi = \varepsilon & \text{or} \\ \pi = i \cdot \pi', t = f(t_0, \dots, t_{k-1}), k > 0, 0 \leq i < k, \pi' \in \mathcal{P}(t_i) \end{cases}$$

The set $\mathcal{P}(t)$ is sometimes also referred to as the *set of occurrences* in t. Accordingly, elements of $\mathcal{P}(t)$ are called *positions* or *occurrences*. The latter notion is preferred whenever the subterm at a particular position is referred to rather than the position itself.

(ii) The *set of subterms* of the term t, denoted $\mathcal{S}(t)$, is the least set S such that

(a) $t \in S$, and

(b) if $f(t_1, \dots, t_k) \in S$, then also $t_1, \dots, t_k \in S$.

A term s is called a *subterm* of t if $s \in \mathcal{S}(t)$. A subterm s of t is called *proper* if $s \neq t$.

Since positions are sequences, all notions that we have for sequences carry over to positions. The most important ones are the *prefix order* $\leq$ and the *disjointness* of two sequences.

Note that $\mathcal{P}(t)$ is well-defined for all terms t as, in the right-hand side occurrence of the $\dots \in \mathcal{P}(\dots)$ statement in the recursive definition, the referred position π' is smaller than the position π on the left-hand side. Subsequently, we will often use induction on (the length of) the positions of terms rather than induction on the structure of terms in order to be able make definitions and proofs work also on infinite terms. Alternatively, we will use coinductive definitions and arguments, i.e. we will make use of the finality of the $F_{\Sigma,\mathcal{V}}$-coalgebra $\mathcal{T}^\infty(\Sigma, \mathcal{V})$.

Beside these formal difficulties also practical problems arise. The foremost of these problems is that of how to represent infinite terms. We deal with a partial solution to this problem in Chapter 4. But also when presenting examples, we need a notation for denoting infinite terms. One abbreviation that we will use is the notation f^ω for the infinite term $f(f(f(\dots)))$.

Definition 2.3.15 (positions as pointers)
Let $s, t \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$ be terms and $\pi \in \mathcal{P}(t)$ be a position in t.

(i) The *subterm* of t at π, denoted $t|_\pi$, is inductively defined as

$$t|_\pi := \begin{cases} t & \text{if } \pi = \varepsilon \\ t_i|_{\pi'} & \text{if } \pi = i \cdot \pi', t = f(t_0, \dots, t_{k-1}) \text{ and } 0 \le i < k \end{cases}$$

(ii) The *symbol* of t at π, denoted $t(\pi)$, is defined as $t|_\pi$ if $t|_\pi \in \mathcal{V}$ and f if $t|_\pi = f(t_1, \dots, t_k)$, $k \ge 0$.

(iii) If $\pi \in \mathcal{P}(s)$ and $s(\pi) = t(\pi)$, then we say that s and t *coincide* in π.

(iv) The *replacement* of the subterm of t at π by s, denoted by $t[s]_\pi$, is inductively defined as

$$t[s]_\pi := \begin{cases} s & \text{if } \pi = \varepsilon \\ f(t_0, \dots, t_{i-1}, t_i[s]_{\pi'}, t_{i+1}, \dots, t_{k-1}) & \text{if } \pi = i \cdot \pi' \text{ and } t = f(t_0, \dots, t_{k-1}) \end{cases}$$

Notation 2.3.16. Sometimes we need to restrict the occurrences in a term t to the non-variable occurrences, i.e. the set $\{\pi \in \mathcal{P}(t) \mid t(\pi) \in \Sigma\}$. We denote this set of non-variable occurrences in a term t as $\mathcal{P}_\Sigma(t)$. Moreover, we use the notation $\mathsf{Var}(t)$ to denote the set $\{t(\pi) \mid t(\pi) \in \mathcal{V}\}$ of variables occurring in t.

Remark 2.3.17. Infinitary terms in $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ can also be specified by a mapping $\varphi \colon D \to \Sigma \cup \mathcal{V}$ where D is a subset of $\mathbb{N}^*$ such that, for each $\pi \cdot i \in D$, we have

$$\pi \in D, \quad \text{and} \qquad \text{(closure under prefixes)}$$
$$0 \le i < \mathsf{ar}(\varphi(\pi)). \qquad \text{(ranked)}$$

This mapping then uniquely determines a term $t \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$ with $\mathcal{P}(t) = D$ and $t(\pi) = \varphi(\pi)$ for all $\pi \in D$.

Definition 2.3.18 (context, substitution)

(i) A *context* is a "term with holes" which are represented by a distinguished variable $\Box$. We write $C[, \dots,]$ for a context with at least one occurrence of $\Box$, $C\langle, \dots, \rangle$ for a context with zero more occurrences of $\Box$, $C\{, \dots, \}$ for a context different from $\Box$ having zero or more occurrences of $\Box$, and $C[]$ for a context with exactly one occurrence of $\Box$. $C[t_1, \dots, t_n]$ denotes the result of replacing the occurrences of $\Box$ in C (from left to right) by $t_1, \dots, t_n$. $C\langle t_1, \dots, t_n\rangle$ and $C\{t_1, \dots, t_n\}$ are defined accordingly.

(ii) A *substitution* σ is a mapping from $\mathcal{V}$ to $\mathcal{T}^\infty(\Sigma, \mathcal{V})$. Its *domain*, denoted $\mathsf{dom}(\sigma)$, is the set $\{x \in \mathcal{V} \mid \sigma(x) \ne x\}$ of variables not mapped to itself by σ. If the range of σ is a subset of $\mathcal{T}(\Sigma, \mathcal{V})$, we call it a *finite substitution*. Substitutions are uniquely extended to morphisms from $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ to $\mathcal{T}^\infty(\Sigma, \mathcal{V})$, by the finality of the $F_{\Sigma,\mathcal{V}}$-coalgebra $\mathcal{T}^\infty(\Sigma, \mathcal{V})$, via $\sigma(f(t_1, \dots, t_n)) = f(\sigma(t_1), \dots, \sigma(t_n))$ for $f \in \Sigma^{(n)}$ and $t_1, \dots, t_n \in \mathcal{T}(\Sigma, \mathcal{V})$. Instead of $\sigma(s)$ we shall also write $s\sigma$.

On $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ we can define the following metric: Let $t, t' \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$. At first define the *similarity* of t and t', denoted $\mathsf{sim}(t, t')$, as the minimal depth where t and t' differ, i.e.

$$\mathsf{sim}(t, t') = \min \{|\pi| \mid \pi \in \mathbb{N}^*, t(\pi) \ne t'(\pi)\} \cup \{\infty\}$$

Note that $\mathsf{sim}(t, t') = \infty$ iff $t = t'$. Now we define the distance $\mathbf{d}$ as

$$\mathbf{d}(t, t') = 2^{-\mathsf{sim}(t, t')},$$

where we interpret $2^{-\infty}$ as 0.

If not stated otherwise, we will consider $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ and all subsets of it (in particular $\mathcal{T}(\Sigma, \mathcal{V})$) always with this metric. The following can be shown for the metric space $\mathcal{T}^\infty(\Sigma, \mathcal{V})$.

Proposition 2.3.19 (properties of the metric space $\mathcal{T}^\infty(\Sigma, \mathcal{V})$, [AN80])

(i) $(\mathcal{T}^\infty(\Sigma, \mathcal{V}), \mathbf{d})$ *is an ultrametric space.*

(ii) $(\mathcal{T}^\infty(\Sigma, \mathcal{V}), \mathbf{d})$ *is complete.*

(iii) $(\mathcal{T}(\Sigma, \mathcal{V}), \mathbf{d})$ *is dense in* $\mathcal{T}^\infty(\Sigma, \mathcal{V})$.

(iv) $(\mathcal{T}^\infty(\Sigma, \mathcal{V}), \mathbf{d})$ *is the metric completion of* $(\mathcal{T}(\Sigma, \mathcal{V}), \mathbf{d})$.

The set of terms can also be endowed with a partial order $\leq_\bot$, assuming a special constant symbol $\bot$ that is supposed to denote "undefinedness". We write $\Sigma_\bot$ to denote the signature $\Sigma \uplus \{\bot\}$. The terms in $\mathcal{T}^\infty(\Sigma_\bot, \mathcal{V})$ are also called *partial terms* over Σ. If it is necessary to make it explicit, we refer to the terms in $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ as *total terms*. $\leq_\bot$ is defined on $\mathcal{T}^\infty(\Sigma_\bot, \mathcal{V})$ as the least partial order with $\bot \leq t$ for all $t \in \mathcal{T}^\infty(\Sigma_\bot, \mathcal{V})$ that is *monotone*, where monotone means that, for all $k > 0, f \in \Sigma^{(k)}$ and $s, t_0, \ldots, t_{k-1} \in \mathcal{T}^\infty(\Sigma_\bot, \mathcal{V})$, we have that

$$t_j \leq_\bot s \text{ with } 0 \leq j < k \quad \text{ implies } \quad f(t_0, \ldots, t_j, \ldots, t_{k-1}) \leq_\bot f(t_0, \ldots, s, \ldots, t_{k-1})$$

Equivalently, the order can be characterised as follows: For two terms $s, t \in \mathcal{T}^\infty(\Sigma_\bot, \mathcal{V})$, we have $s \leq_\bot t$ iff there is a context $C\langle, \ldots, \rangle \in \mathcal{T}^\infty(\Sigma_\bot, \mathcal{V})$ and terms $t_1, \ldots, t_n \in \mathcal{T}^\infty(\Sigma_\bot, \mathcal{V})$ such that $s = C\langle\bot, \ldots, \bot\rangle$ and $t = C\langle t_1, \ldots, t_n\rangle$.

Proposition 2.3.20 ($\leq_\bot$ is a complete semilattice, [GTWW77])
$\leq_\bot$ *is a complete semilattice on* $\mathcal{T}^\infty(\Sigma_\bot, \mathcal{V})$.

Notation 2.3.21. Since the symbol $\bot$ in partial terms is supposed to denoted undefinedness, it is often necessary to neglect the $\bot$-occurrences in a term. To this end, we use the notation $\mathcal{P}_{\setminus\bot}(t)$ to denote the set of occurrences of subterms in t different from $\bot$, i.e. the set $\{\pi \in \mathcal{P}(t) \mid t(\pi) \neq \bot\}$. Analogously, we use the notation $\mathcal{P}_\bot(t)$ for the set $\{\pi \in \mathcal{P}(t) \mid t(\pi) = \bot\}$ of $\bot$-occurrences in t.

2.3.3 Term Rewriting Systems

In this section we introduce the basic theory of term rewriting.

Definition 2.3.22 ((infinitary) term rewriting system)
Let Σ be a signature.

(i) A *term rewrite rule* (or simply *rewrite rule*) over Σ is a pair (l, r) of finite terms in $\mathcal{T}(\Sigma, \mathcal{V})$ with $l \notin \mathcal{V}$ and $\mathsf{Var}(r) \subseteq \mathsf{Var}(l)$. Instead of (l, r) we usually write $l \to r$ and call l its *left-hand side* and r its *right-hand side*. If also infinitary terms in $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ are allowed for the right-hand side, we call $l \to r$ an *infinitary term rewrite rule*.

(ii) An *infinitary term rewriting system* (*ITRS*) over Σ is a pair $\mathcal{R} = (\Sigma, R)$ consisting of a signature Σ and a set R of infinitary term rewrite rules over Σ. If R is a set of term rewrite rules, then $\mathcal{R}$ is called a *term rewriting system* (*TRS*).

(iii) Let $\mathcal{R}$ be an ITRS over Σ. We say that t is a term in $\mathcal{R}$ if $t \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$. Analogously, we say that t is a partial term in $\mathcal{R}$ if $t \in \mathcal{T}^\infty(\Sigma_\bot, \mathcal{V})$.

Notation 2.3.23. Sometimes it is convenient to give a rewrite rule a name. In this case, we write $\rho\colon l \to r$ to indicate that the rule $l \to r$ has been given the name ρ.

In the following, we will define some notions for ITRSs. As before, since TRSs are a special case of ITRSs, these notions are also applicable to TRSs. This is particularly true for the definition of the semantics of ITRSs which associates an ARS to an ITRS.

Definition 2.3.24 (semantics of ITRSs)
Let $\mathcal{R} = (\Sigma, R)$ be an ITRS.

(i) A *prestep* of $\mathcal{R}$ is a triple (t, π, ρ) consisting of a term $t \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$, a position $\pi \in \mathcal{P}(t)$, and a rule $\rho \in R$.

(ii) A rule $\rho\colon l \to r \in R$ is *applicable* to a term $t \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$ at $\pi \in \mathcal{P}(t)$ if there is a substitution σ such that $l\sigma = t|_\pi$. The term $t[r\sigma]_\pi$ is called the *result* of the application of ρ to t at position π. Moreover, the subterm $l\sigma$ is called a *ρ-redex*, the position π is called a *ρ-redex occurrence*. The reference to the rule ρ may be omitted if it is irrelevant.

(iii) A prestep $\varphi = (t, \pi, \rho)$ is called a *step* if ρ is applicable to t at π.

(iv) The *induced ARS* of $\mathcal{R}$, denoted $\mathcal{A}_\mathcal{R}$, is given by the tuple $(A, \Phi, \mathsf{src}, \mathsf{tgt})$, where the set of objects A is the set of terms $\mathcal{T}^\infty(\Sigma, \mathcal{V})$, and the set of steps Φ is the set consisting of the steps of $\mathcal{R}$. Let $\varphi = (t, \pi, \rho)$ be a step of $\mathcal{R}$. We define $\mathsf{src}(\varphi) = t$ and $\mathsf{tgt}(\varphi) = t'$, where t' is the result of the application of ρ to t at π.

Notation 2.3.25. In virtue of the preceding definition, every ITRS $\mathcal{R} = (\Sigma, R)$ can be associated with its induced ARS $\mathcal{A}_\mathcal{R}$. That is why we identify $\mathcal{R}$ with $\mathcal{A}_\mathcal{R}$ and consider ITRSs as a special case of ARSs. In particular, we will write $s \to_\mathcal{R} t$ instead of $s \to_{\mathcal{A}_\mathcal{R}} t$. Since the steps of an ITRS additionally contain information about the rule that was applied and the position where it was applied, we sometimes want to make this explicit and use the notation $\varphi\colon s \to_{\pi,\rho} t$ for a step $\varphi = (s, \pi, \rho)$.

Definition 2.3.26 (defined symbols and constructors)
Let $\mathcal{R} = (R, \Sigma)$ be an ITRS and $f \in \Sigma$. f is called a *defined symbol* of $\mathcal{R}$ if there is a rule $l \to r$ with $l(\varepsilon) = f$. The set of all defined symbols is denoted by $\mathcal{D}_\mathcal{R}$. f is called a *constructor* if it is not a defined symbol. We use $\mathcal{C}_\mathcal{R}$ to denote the set of all constructors in $\mathcal{R}$. If it is clear which ITRS is meant, the subscript $\mathcal{R}$ may be dropped.

Definition 2.3.27 (pattern and arguments)
Let $\rho\colon l \to r$ be an infinitary term rewrite rule.

(i) The *pattern* of ρ is the context $l\sigma_\Box$, where $\sigma_\Box$ is the substitution $\{x \mapsto \Box \mid x \in \mathcal{V}\}$ that maps all variables to $\Box$.

(ii) Let t be a ρ-redex. Then the pattern P of ρ is also called a the *redex pattern* of t w.r.t. ρ. Let $l = P[x_1, \ldots, x_n]$ and σ the substitution such that $t = l\sigma = P[x_1\sigma, \ldots, x_n\sigma]$. The terms $x_1\sigma, \ldots, x_n\sigma$ are called the *arguments* of the redex t.

(iii) When referring to the positions (or occurrences) in a pattern, occurrences of the symbol $\Box$ are neglected.

Definition 2.3.28 (critical pair)
Let $\mathcal{R}$ be an ITRS, $\rho_1\colon l_1 \to r_1$, $\rho_2\colon l_2 \to r_2$ two variable-renamed rules of $\mathcal{R}$ not sharing variables, i.e. $\mathsf{Var}(l_1)$ and $\mathsf{Var}(l_2)$ are disjoint, and π a non-variable position in l_1 such that $l_1|_\pi$ and l_2 are unifiable with most general unifier σ. Then $(l_1[r_2]_\pi\sigma, r_1\sigma)$ is called a *critical pair* of ρ_1 and ρ_2 at position π unless $\pi = \varepsilon$ and ρ_1 and ρ_2 are two renamed versions of the same rule. If such a critical pair exists, then ρ_2 is said to *overlap* ρ_1 at position π. If $\pi = \varepsilon$, then the critical pair is called an *overlay*. The critical pair is called *trivial* if $l_1[r_2]_\pi\sigma = r_1\sigma$. The critical pair is called *joinable* if $l_1[r_2]_\pi\sigma$ and $r_1\sigma$ have a common reduct $r_1\sigma$, i.e. there is a term t with $l_1[r_2]_\pi\sigma \to^\star t \leftarrow^\star r_1\sigma$.

A related notion is that of conflicting redex occurrences:

Definition 2.3.29 (conflicting redex occurrences)
Let $\mathcal{R}$ be an ITRS, t a term in $\mathcal{R}$, and u, v two distinct redex occurrences in t. u and v are called *conflicting* if there is a position π such that $v = u \cdot \pi$ and π is a pattern position of the redex at u, or, vice versa, $u = v \cdot \pi$ and π is a pattern position of the redex at v. If this is not the case, then u and v are called *non-conflicting*.

Conflicting redex occurrences correspond to critical pairs which are not overlays. More precisely, whenever there is a pair of conflicting redex occurrences, say u and $v = u \cdot \pi$, then there is a critical pair $(l_1[r_2]_\pi \sigma, r_1\sigma)$ of $\rho_1 \colon l_1 \to r_1$ and $\rho_2 \colon l_2 \to r_2$ at position π, with ρ_1 a rule for the redex occurrence u and ρ_2 a rule for the redex occurrence v. Consequently, a system without critical pairs does not admit conflicting redex occurrence. This is one of the motivations for the following definition:

Definition 2.3.30 (linearity, orthogonality)
Let $\mathcal{R}$ be an ITRS.

(i) A term t is called *linear* if each variable has at most one occurrence in t.

(ii) A rule $l \to r$ of $\mathcal{R}$ is called *left-linear* if l is linear.

(iii) $\mathcal{R}$ is called *left-linear* if each rule in $\mathcal{R}$ is linear.

(iv) $\mathcal{R}$ is called *orthogonal* if it is left-linear and has no critical pairs.

(v) $\mathcal{R}$ is called *almost orthogonal* if it is left-linear and all its critical pairs are trivial overlays.

(vi) $\mathcal{R}$ is called *weakly orthogonal* if it is left-linear and all its critical pairs are trivial.

(vii) $\mathcal{R}$ is called an *overlay system* if all its critical pairs are overlays.

It is easy to see that orthogonal ITRSs are also almost orthogonal, and in turn almost orthogonal ITRSs are also weakly orthogonal. Moreover, it is clear that overlay systems and, a fortiori, almost orthogonal systems do not admit conflicting redex occurrences.

Theorem 2.3.31 (confluence of weakly orthogonal ITRSs)
Each weakly orthogonal ITRS is confluent.

Definition 2.3.32 (collapsing rules)
Let $\mathcal{R}$ be an orthogonal ITRS.

(i) A rule $l \to r$ in $\mathcal{R}$ is called *collapsing* if r is a variable. The unique position of the variable r in l is called the *collapsing position* of the rule.

(ii) A ρ-redex is called *collapsing* if ρ is a collapsing rule.

(iii) A *collapsing tower* is a non-empty sequence $(u_i)_{i<\alpha}$ of collapsing redex occurrences in a term t such that $u_{i+1} = u_i \cdot \pi_i$ for each $i < \alpha$, where π_i is the collapsing position of the redex at u_i.

(iv) A collapsing tower $(u_i)_{i<\alpha}$ in a term t is said to be *maximal* if it is infinite or if it is finite and $u_{\alpha-1} \cdot \pi_{\alpha-1}$ is not a collapsing redex occurrence, where $\pi_{\alpha-1}$ is the collapsing position of the redex $t|_{u_{\alpha-1}}$.

Definition 2.3.33 (top-termination)
An ITRS $\mathcal{R}$ is called *top-terminating* if there are no ω-reduction sequences with infinitely many steps occurring at the root position.

Chapter 3

Transfinite Reductions

The purpose of this chapter is to present models for extending abstract reduction systems in order to define reduction sequences of transfinite length in a meaningful way. To this end, a theoretical tool is needed to formalise the intuition of a limit of a sequence of elements. The most general way of defining such limits is by considering topological spaces (cf. Section 2.2). In particular, we want to investigate metric spaces as a basis for transfinite reductions. This results in the theory of *metric reduction systems* which is presented in Section 3.1.

Another variant of limit constructions that can be explored is that of the limit inferior and limit superior which is more powerful as it might permit convergence even if a limit does not exist. This can again be formalised in terms of topological spaces. However, we restrict our analysis of that matter to the particular case of partial orders and their notion of limit inferior. We introduce reduction systems utilising this framework in Section 3.2. The thus obtained reduction systems are called *partial reduction systems*.

Both approaches to transfinite reductions offer two kinds of reduction sequences: a weak and a strong variant. The weak variant only takes into account the objects in a reduction sequence and how they differ from each other. On the other hand, the strong variant additionally considers how the objects of the reduction sequence have been created, i.e. it also takes into account the reduction steps "between" the objects. For example, in the case of term rewriting the position where each reduction rule in the reduction sequence is applied becomes significant in the strong variant of transfinite reductions. We will see in more detail what this means concretely for metric reduction systems and partial reduction systems, respectively.

As the models that we are going to present in this chapter are rather abstract, we will instantiate them for the case of term rewriting at a very early stage of the discussion in order to provide an intuition for them and also to create appreciation for their merit. However, an in-depth exploration of the consequences of the different models for transfinite term rewriting sequences is deferred until Chapter 5.

In Section 3.3, the different notions of transfinite reduction sequences are compared. This comprises a discussion of how to lift properties known from finite reductions to the realms of transfinite reductions and how theses lifted properties behave in comparison to their original finitary variants. Of course, in general this depends on which variant of transfinite reduction we consider. However, interestingly enough this choice is in most cases insignificant as we will see.

Moreover, we present a set of abstract criteria which will guarantee that the reduction sequences of MRSs and PRSs are the same up to some well-defined subclass of them. This is going to be crucial for the subsequent chapters, as this allows us to transfer results between the MRS and the PRS world.

At the end of this chapter, in Section 3.4, we mention other possibilities of defining transfinite reductions.

3.1 Metric Reduction Systems

In this section we want to formalise the extension of abstract reduction systems by a metric space which enables us to define reduction sequences of transfinite length. Historically, the idea of employing a metric space was the first one that was proposed for this purpose (cf. [DKP89, DK89, DKP91]). Later Kennaway et al. [KKSdV91] proposed a restriction to this definition of transfinite rewriting which required the depth of the reduction steps to tend to infinity in order to obtain a limit for reductions of limit ordinal length. This resulted in the distinction between weak and strong variants of transfinite reductions.

In [Ken92] Kennaway generalised this approach to transfinite reductions to abstract reduction systems. This section is largely based on this work. The definition below follows the idea of the notion of metric abstract reduction systems that Kennaway introduced in [Ken92].

Definition 3.1.1 (metric reduction system [Ken92])
A *metric reduction system* (*MRS*) is a tuple $\mathcal{M} = (A, \Phi, \mathsf{src}, \mathsf{tgt}, \mathbf{d}, \mathsf{hgt})$, such that

(i) $\mathcal{A} = (A, \Phi, \mathsf{src}, \mathsf{tgt})$ is an ARS, called the *underlying ARS* of $\mathcal{M}$,

(ii) $\mathbf{d}\colon A \times A \to \mathbb{R}^+_0$ is a function such that $(A, \mathbf{d})$ is a metric space,

(iii) $\mathsf{hgt}\colon \Phi \to \mathbb{R}^+$ is a function, called the *height function*, and

(iv) if $\varphi\colon a \to_{\mathcal{A}} b$, then $\mathbf{d}(a, b) \leq \mathsf{hgt}(\varphi)$.

If the metric of an MRS $\mathcal{M}$ is even an ultrametric, then $\mathcal{M}$ is called an *ultrametric reduction system* (*URS*). Furthermore, an MRS is referred to as *complete* if the underlying metric space is complete.

The heart of every MRS is the underlying ARS. The other two parts, viz. the metric and the height, are needed to define the limit behaviour of transfinite reduction sequences, viz. continuity and convergence, and to distinguish weak and strong variants thereof, respectively.

Notation 3.1.2. We consider an MRS $\mathcal{M}$ as an instance of the underlying ARS $\mathcal{A}$. This allows us to conveniently use the same notions and notations introduced for ARSs. In particular, we write $\varphi\colon a \to_{\mathcal{M}} b$ for a reduction step $\varphi\colon a \to_{\mathcal{A}} b$ in the underlying ARS. If we want to explicitly indicate the height of a reduction step in $\mathcal{M}$, we write $\varphi\colon a \to_h b$ whenever $\mathsf{hgt}(\varphi) = h$.

We also want to extend the convention of defining ARSs to the concept of MRSs. Instead of giving the elements $\Phi, \mathsf{src}, \mathsf{tgt}, \mathsf{hgt}$ explicitly we choose to define them implicitly where appropriate. To this end, we provide a collection of statements of the form $a \to_h b$, where a, b are objects of the MRS and $h \in \mathbb{R}^+$. This defines Φ, src, tgt and hgt as follows: Φ is the set of the given statements of the form $a \to_h b$, $\mathsf{src}(a \to_h b) = a$, $\mathsf{tgt}(a \to_h b) = b$, and $\mathsf{hgt}(a \to_h b) = h$.

Since the definition of MRSs is indeed rather abstract, we choose to already extend the semantics of ITRSs to the model of MRSs such that its intuition becomes clear.

Definition 3.1.3 (MRS semantics of ITRSs)
Let $\mathcal{R} = (\Sigma, R)$ be an ITRS. The MRS *induced* by $\mathcal{R}$, denoted $\mathcal{M}_{\mathcal{R}}$, is given by the tuple

$$(\mathcal{T}^\infty(\Sigma, \mathcal{V}), \Phi, \mathsf{src}, \mathsf{tgt}, \mathbf{d}, \mathsf{hgt}),$$

where $(\mathcal{T}^\infty(\Sigma, \mathcal{V}), \Phi, \mathsf{src}, \mathsf{tgt})$ is the ARS $\mathcal{A}_{\mathcal{R}}$ induced by $\mathcal{R}$, $\mathbf{d}$ is the usual ultrametric on $\mathcal{T}^\infty(\Sigma, \mathcal{V})$, and hgt is defined as

$$\mathsf{hgt}(\varphi) = 2^{-|\pi|}, \text{ where } \varphi\colon t \to_{\pi,\rho} t'.$$

The choice of the definition of the height of a reduction step is made in accordance with the definition of the metric on terms: Whereas the distance $\mathbf{d}(s,t)$ of two terms s and t is defined as 2^{-d} with d the minimal depth of a discrepancy between s and t (or ∞ if none exists), the height of a rewrite step $s \to t$ is defined as $2^{-d'}$ with d' the depth of the redex that was contracted. This illustrates that the height of a rewriting step is an overapproximation of the actual distance $\mathbf{d}(s,t)$ of the involved terms in the sense of clause (iv) of Definition 3.1.1. In fact, we have the following:

Proposition 3.1.4 (MRS semantics yields a complete URS)
Each ITRS $\mathcal{R}$ induces a complete URS $\mathcal{M}_\mathcal{R}$.

Proof. As shown in Proposition 2.3.19, $\mathcal{T}^\infty(\Sigma,\mathcal{V})$ forms a complete ultrametric space. Hence, it remains to be shown that $\mathbf{d}(s,t) \le \mathsf{hgt}(\varphi)$ for each reduction step $\varphi\colon s \to_{\pi,\rho} t$. This is immediate from the definition of $\mathbf{d}$ and hgt: Since the reduction step takes place at π, we have $\mathsf{sim}(s,t) \ge |\pi|$ and, therefore, $\mathbf{d}(s,t) = 2^{-\mathsf{sim}(s,t)} \le 2^{-|\pi|} = \mathsf{hgt}(\varphi)$. □

Remark 3.1.5. The MRS semantics of ITRSs as given above illustrates the need for reification of reduction steps in the definition of ARSs, i.e. making reduction steps distinguishable objects themselves. For example, consider the TRS with the single rule $\rho\colon f(x) \to x$ and the term $f(f(a))$. There are two different ways of applying the rule to this term. Either at position $\pi_1 = \varepsilon$, to the redex $f(f(a))$, or at position $\pi_2 = 1$, to the redex $f(a)$. In both cases we have a rewrite step of the form $f(f(a)) \to f(a)$. But to have a well-defined semantics of ITRSs in terms of MRSs, we need to distinguish these two steps $\varphi_1\colon f(f(a)) \to_{\pi_1,\rho} f(a)$ and $\varphi_2\colon f(f(a)) \to_{\pi_2,\rho} f(a)$ because $\mathsf{hgt}(\varphi_1) = 1$ whereas $\mathsf{hgt}(\varphi_2) = \frac{1}{2}$. The phenomenon described above, i.e. that the position and/or the applied rule is not necessarily derivable from the start and the end term of a reduction step, is also known as *syntactic accident* [Lév78].

The reification of reduction steps is the major difference in the definition of MRSs in this thesis and its definition in [Ken92].

Next, we want to revisit the notion of a reduction sequence. Section 2.3.1 already provided a definition for reduction sequences, even allowing sequences of arbitrary length. Yet, as mentioned in Remark 2.3.4, reduction sequences of length greater than ω are not meaningful. For example, considering the term rewriting rules $a \to f(a)$ and $b \to g(b)$, the following constitutes a valid reduction sequence of length 2ω according to Definition 2.3.3:

$$a \to f(a) \to f(f(a)) \to f(f(f(a))) \to \;\dots\; b \to g(b) \to g(g(b)) \to g(g(g(b))) \to \;\dots$$

The problem that occurs here is that the second half of the reduction sequence, the one starting with b, is completely arbitrary. The culprit of this phenomenon is that the reduction step $b \to g(b)$ has no immediate predecessor. It is the ω-th step in the reduction sequence. In general the definition of a reduction sequence, say $(a_\iota \to a_{\iota+1})_{\iota<\alpha}$, does not stipulate any relation between the object a_λ at a limit ordinal $\lambda < \alpha$ and the elements a_ι, $\iota < \lambda$, that preceded it. Therefore, a notion of continuity is needed.

Definition 3.1.6 (continuity, convergence)
Let $S = (a_\iota \to_{h_\iota} a_{\iota+1})_{\iota<\alpha}$ be a reduction sequence in an MRS $\mathcal{M}$.

(i) S is called *weakly continuous* if the underlying sequence $(a_\iota)_{\iota<\alpha'}$ of elements is continuous in the metric space of $\mathcal{M}$. If, additionally, $\lim_{\iota\to\lambda} h_\iota = 0$ for each limit ordinal $\lambda < \alpha'$, then the sequence is called *strongly continuous*.

(ii) A weakly continuous reduction sequence $S = (a_\iota \to_{h_\iota} a_{\iota+1})_{\iota<\alpha}$ is called *weakly convergent* if the sequence $(a_\iota)_{\iota<\alpha}$ converges, say to some element a_α. A strongly continuous reduction sequence $S = (a_\iota \to_{h_\iota} a_{\iota+1})_{\iota<\alpha}$ is called *strongly convergent* if it is weakly convergent and $\lim_{\iota\to\alpha} h_\iota = 0$ in case α is a limit ordinal. In the case of weak (resp.

strong) convergence, we also say that the reduction sequence S weakly (resp. strongly) converges to a_α or, alternatively, that a_0 weakly (resp. strongly) converges to a_α in α steps. A weakly (resp. strongly) continuous reduction sequence is called *weakly* (resp. *strongly*) *divergent* if it is not weakly (resp. strongly) convergent.

(iii) $\mathcal{M}$ is called *weakly α-convergent* or *strongly α-convergent*, respectively, if any weakly continuous resp. strongly continuous α-reduction sequence in $\mathcal{M}$ is weakly convergent resp. strongly convergent.

Remark 3.1.7. Note that weak resp. strong convergence of a reduction sequence requires it also to be weakly resp. strongly continuous. This deviates from the usual topological notion of convergence which is independent from continuity. The reason for the choice of this definition is that reduction sequences, which are not weakly or strongly convergent, are in general not meaningful. The question whether such sequences converge to a limit is, therefore, also of no relevance.

Returning to the example we had considered before, viz. the reduction sequence

$$a \to f(a) \to f(f(a)) \to f(f(f(a))) \to \dots\ b \to g(b) \to g(g(b)) \to g(g(g(b))) \to \dots,$$

we can see that this is not a weakly continuous reduction sequence. The prefix

$$a \to f(a) \to f(f(a)) \to f(f(f(a))) \to \dots$$

weakly (and also strongly) converges to f^ω. As we will see later, in order to extend a reduction sequence by another one, they have to be compatible, i.e. the first sequences has to weakly/strongly converge to the element the second sequence starts with. In the example, this means that the extending sequence has to start in f^ω instead of b. Assuming there is a rule $f(x) \to g(x)$, the following is a weakly (and strongly) continuous reduction sequence of length 2ω:

$$a \to f(a) \to f(f(a)) \to f(f(f(a))) \to \dots\ f^\omega \to g(f^\omega) \to g(g(f^\omega)) \to \dots$$

Let us have a look at the difference between strong and weak continuity resp. convergence. If a reduction sequence $S = (a_\iota \to_{h_\iota} a_{\iota+1})_{\iota<\alpha}$ is weakly continuous, then the underlying sequence $(a_\iota)_{\iota<\alpha}$ is continuous in the metric space of the MRS. According to Corollary 2.2.17, this means that, for each limit ordinal $\lambda < \alpha$, it holds that $\lim_{\iota\to\lambda} \mathbf{d}(a_\iota, a_{\iota+1}) = 0$. Because of the definition of the metric on terms, this means that $\mathsf{sim}(a_\iota, a_{\iota+1})$ has to tend to infinity as ι approaches λ. In other words, the minimal depth of discrepancies between two consecutive terms tends to in infinity, i.e., intuitively, the differences between the terms becomes more and more insignificant. For strong convergence, it is additionally required that even the height of the reduction steps $\varphi\colon a_\iota \to a_{\iota+1}$, which by definition has to overapproximate the distance $\mathbf{d}(a_\iota, a_{\iota+1})$ between the involved objects, tends to zero for each limit ordinal, i.e. $\lim_{\iota\to\lambda} h_\iota = 0$. Due to the definition of the height of reduction steps in ITRSs, this means that the depth at which reductions are performed tends to infinity as ι approaches λ. The same intuition is also valid for the difference between weak and strong convergence. For closed reduction sequences, both notions coincide given strong continuity. For weak convergence of open reduction sequences, we need, according to Fact 2.2.16, that $\lim_{\iota\to\alpha} \mathbf{d}(a_\iota, a_{\iota+1}) = 0$. And again, for strong convergence, this is strengthened to $\lim_{\iota\to\alpha} h_\iota = 0$.

Note that the conditions for weak continuity and convergence that were mentioned above are only necessary, not sufficient. Only the condition given for weak convergence is sufficient provided the MRS under consideration is a complete URS (cf. Lemma 2.2.13).

Example 3.1.8
Consider the TRS with the single rule $f(g(x)) \to f(g(g(x)))$. The induced reduction sequence

$$S\colon f(g(c)) \to f(g^2(c)) \to f(g^3(c)) \to \dots$$

weakly converges to the term $f(g^\omega)$. Yet it does not strongly converge as each reduction step takes place at the root and, thus, at depth 0. That is, the depth of the contraction

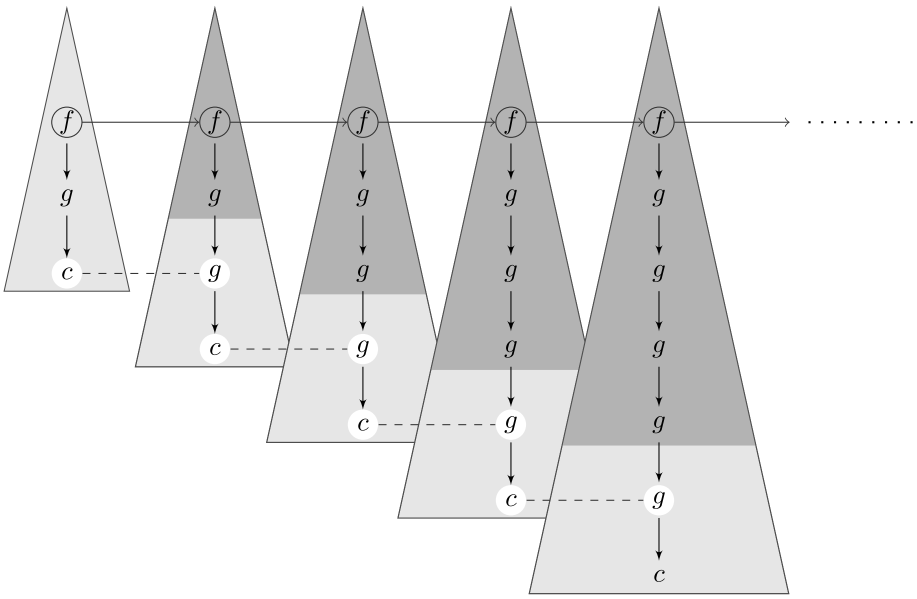

Figure 3.1: Weakly convergent reduction sequence.

site does not tend to infinity. The behaviour of the reduction sequence S is illustrated in Figure 3.1. The circles at the top of each term tree and the rewriting arrows indicate the position where the rewriting step is performed. When we follow the rewriting arrows we can see that they stay at the same depth. The white circles and the dashed lines indicate the shallowest position where two consecutive terms differ. The parts of the terms that remain unchanged are coloured in a darker shade of grey. One can see that these parts become bigger and bigger as the differing parts are pushed further down the term.

On the other hand, consider the TRS with the single rule $g(c) \to g(g(c))$. In this system the reduction sequence

$$T\colon f(g(c)) \to f(g^2(c)) \to f(g^3(c)) \to \ \ldots$$

does both weakly and strongly converge to the term $f(g^\omega)$. The terms of this reduction sequence are the same as in the reduction sequence S. In T, however, we have a different reduction rule which is now applied at different positions compared to S. One can see that the depth where the reduction takes place increases with each step. This can also be observed in Figure 3.2 which depicts the reduction sequence T in a similar fashion as Figure 3.1. One can see that the depth where the reduction is performed increases with each step. In particular, this means that it tends to infinity. Additionally, one can see that the context, where the reduction takes place and is, thus, kept untouched, grows with each step as indicated with the yet darker shade of grey.

Remark 3.1.9. One can easily see that the notion of a weakly or a strongly convergent reduction sequence over an MRS as presented above is a conservative extension of the notion of a finite reduction sequence over an ARSs. The additional structure that MRSs provide, as compared with ARSs, is irrelevant for finite reduction sequences. Similarly, the MRS semantics of TRSs in terms of MRSs is a conservative extension of the usual semantics in terms of ARSs. Therefore, we adopt all concepts and properties of ARSs for finite reduction sequences to the world of finite reduction sequences in MRSs. Furthermore, we can focus on the MRS semantics of TRSs; the usual semantics of ITRSs was only used to illustrate the MRS semantics appropriately.

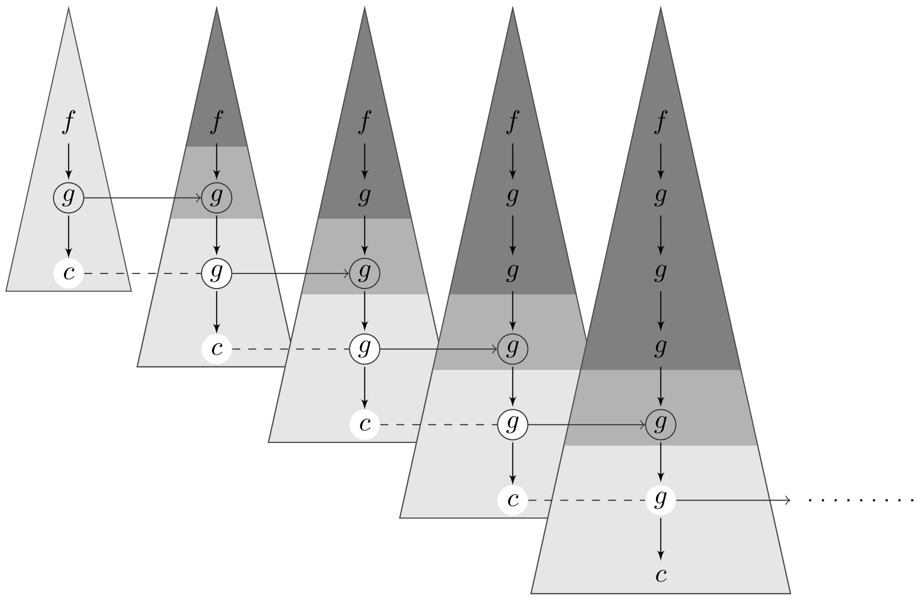

Figure 3.2: Strongly convergent reduction sequence.

In order to concisely state the continuity and convergence of reduction sequences or their lack thereof, we introduce some notation that generalises the notation $\to^*$ that we already have for finite reduction sequences.

Notation 3.1.10. Let $S = (a_\iota \to_{s_\iota} a_{\iota+1})_{\iota<\alpha}$ be a weakly continuous reduction sequence of length α in an MRS $\mathcal{A}$. We also denote this fact as $S\colon a_0 \hookrightarrow^{\alpha}_{\mathcal{A}} \ldots$. If S weakly converges to the limit a_α, we denote it as $S\colon a_0 \hookrightarrow^{\alpha}_{\mathcal{A}} a_\alpha$ (cf. [Luc01]) whereas if S is divergent, we denote it as $S\colon a_0 \nearrow^{\alpha}_{\mathcal{A}}$. If the name or the length of the sequence is not of importance, it is omitted, and we simply write $a_0 \hookrightarrow_{\mathcal{A}} \ldots$, $a_0 \hookrightarrow_{\mathcal{A}} a_\alpha$ or $a_0 \nearrow_\alpha$, respectively. Also, when the underlying MRS is clear from the context, it is dropped from the notation. To indicate a condition on the length of the reduction sequence, we may write this condition instead of the explicit length annotation. For example, we use $a \hookrightarrow^{\le\alpha} \ldots$, $a \hookrightarrow^{\le\alpha} a'$ or $a \nearrow^{\le\alpha}$ to indicate weakly continuous, convergent, resp. divergent reduction sequences of length at most α. If it should be indicated that a reduction sequence is strongly continuous (resp. convergent), we use $\twoheadrightarrow$ instead of $\hookrightarrow$ (cf. [Sim06]).

Remark 3.1.11. Note that the definition of a reduction sequence also permits a sequence of length 0. We consider such a reduction sequence as a reduction starting from any object and strongly converging to the same object. That is, we write $a \twoheadrightarrow^0 a$ (or also $a \hookrightarrow^0 a$) for such a reduction sequence, where a is an arbitrary object in the MRS under consideration.

Before we continue, let us review the definition of infinite reduction sequences in the light of the example of the MRS semantics of an ITRS. What does weak continuity of a reduction sequence $S = (t_\iota \to_{h_\iota} t_{\iota+1})_{\iota<\alpha}$ in an ITRS mean? According to Lemma 2.2.13, this means that, for each limit ordinal $\lambda < \alpha$, the distance $\mathbf{d}(t_\lambda, t_\iota)$ tends to 0 when ι approaches λ. Since $\mathbf{d}(t_\lambda, t_\iota)$ is defined as $2^{-\mathsf{sim}(t_\lambda, t_\iota)}$, this is equivalent to $\mathsf{sim}(t_\lambda, t_\iota)$ tending to infinity as ι approaches λ. That is, the minimal depth of discrepancies between t_λ and t_ι tends to infinity, or, informally speaking, the parts that t_λ and t_ι share increase as ι approaches λ. For strong continuity, it is additionally required that the height of the reduction steps tends to zero. For the definition of the MRS semantics, this means that the depth at which the reduction steps take places must tend to infinity. For weak and strong convergence, we can observe a similar relation: For closed sequences, weak and strong convergence trivially hold, provided we have strong continuity. For open sequences, there has to be term t_α such that

$\mathsf{sim}(t_\alpha, t_\iota)$ tends to infinity in order to obtain weak convergence. To get strong convergence, we additionally need that the depth at which the reduction steps take place tends to infinity.

By definition, strong continuity resp. strong convergence implies weak continuity resp. weak convergence:

Fact 3.1.12 (stong continuity/convergence implies weak continuity/convergence)
For every reduction sequence S in an MRS, it holds that

(i) $S\colon a \twoheadrightarrow \dots$ *implies* $S\colon a \hookrightarrow \dots$, *and that*

(ii) $S\colon a \twoheadrightarrow b$ *implies* $S\colon a \hookrightarrow b$.

It turns out that, for a strongly continuous reduction sequence $(\varphi_\iota)_{\iota<\alpha}$, it does not suffice to have $\lim_{\iota\to\alpha} \mathsf{hgt}(\varphi_\iota) = 0$ in order to be strongly converging. The reason for this can be twofold: Either the underlying sequence $(a_\iota \to_{s_\iota} a_{\iota+1})_{\iota<\alpha'}$ is not Cauchy, which might happen if the system is not ultrametric or even though the sequence is Cauchy, the limit might not exists due to the incompleteness of the metric space. The following example illustrates the former cause by giving a complete MRS which is not a complete URS.

Example 3.1.13
Let $\mathcal{M} = (\mathbb{R}, \Phi, \mathsf{src}, \mathsf{tgt}, d, \mathsf{hgt})$, where $\Phi = \{a_n \mid n > 0\}$, $\mathbf{d}$ is the usual metric on $\mathbb{R}$. Φ, src, tgt, and hgt are given by $a_n \to_{h_n} a_{n+1}$, $n \geq 0$, where $a_n = \sum_{k=1}^{n} \frac{1}{k}$ and $h_n = \frac{1}{n+1}$. Note that if $a_n \to_{h_n} a_{n+1}$, then $\mathbf{d}(a_n, a_{n+1}) = h_n$ and that $(\mathbb{R}, \mathbf{d})$ forms a complete metric space. Yet, it is not an ultrametric space. So let us consider the ω-reduction sequence $(a_\alpha \to_{h_i} a_{i+1})_{i<\omega}$. That is, we have the reduction sequence:

$$0 \to_1 1 \to_{\frac{1}{2}} \left(1 + \frac{1}{2}\right) \to_{\frac{1}{3}} \left(1 + \frac{1}{2} + \frac{1}{3}\right) \quad \dots$$

It is well-known that the series $\sum_{k=1}^{\infty} \frac{1}{k}$ does not converge. Hence, the above reduction sequence is not weakly convergent. However, it is vacuously strongly continuous, and the height of the reduction steps tends to 0, i.e. $\lim_{k\to\omega} h_k = \lim_{k\to\omega} \frac{1}{k} = 0$.

This odd behaviour cannot occur in a complete URS.

Proposition 3.1.14 (strong convergence in complete URSs, [Ken92])
Let $\mathcal{M}$ be a complete URS. Every strongly continuous reduction sequence $(\varphi_\iota)_{\iota<\alpha}$ in $\mathcal{M}$ is strongly convergent iff $\lim_{\iota\to\alpha} \mathsf{hgt}(\varphi_\iota) = 0$.

Proof. Let $S = (a_\iota \to_{h_\iota} a_{\iota+1})_{\iota<\alpha}$ be a strongly continuous sequence in a complete URS $\mathcal{M}$. The "only if" direction is immediate from the definition of strong convergence. For the "if" direction, assume that $\lim_{\iota\to\alpha} h_\iota = 0$. It needs to be shown that the underlying sequence $(a_\iota)_{\iota<\alpha'}$ is convergent. By Definition 3.1.1, it holds that $\mathbf{d}(a_\iota, a_{\iota+1}) \leq h_\iota$. Hence, $\lim_{\iota\to\alpha} h_\iota = 0$ implies $\lim_{\iota\to\alpha} \mathbf{d}(a_\iota, a_{\iota+1}) = 0$. Since the underlying metric space is an ultrametric and $(a_\iota)_{\iota<\alpha'}$ is continuous, this shows, by Lemma 2.2.15, that $(a_\iota)_{\iota<\alpha'}$ is Cauchy. This, under the given assumption that the underlying metric space is complete, shows that $(a_\iota)_{\iota<\alpha'}$ is convergent. □

For reduction sequences of finite length, it holds that the concatenation of two compatible reduction sequences, i.e. the latter sequence starts with the object the former ends with, again yields a reduction sequence. The following proposition shows that this also holds for the different notions of infinite reduction sequences.

Proposition 3.1.15 (concatenation of reduction sequences)
Let S and T be two non-empty reduction sequences over the same MRS $\mathcal{M}$, and a, b two objects in $\mathcal{M}$.

(i) $S \cdot T\colon a \hookrightarrow \dots$ *iff there is an object c in $\mathcal{M}$ with* $S\colon a \hookrightarrow c$ *and* $T\colon c \hookrightarrow \dots$.

(ii) $S \cdot T\colon a \hookrightarrow b$ *iff there is an object* c *in* $\mathcal{M}$ *with* $S\colon a \hookrightarrow c$ *and* $T\colon c \hookrightarrow b$.

(iii) $S \cdot T\colon a \twoheadrightarrow \dots$ *iff there is an object* c *in* $\mathcal{M}$ *with* $S\colon a \twoheadrightarrow c$ *and* $T\colon c \twoheadrightarrow \dots$.

(iv) $S \cdot T\colon a \twoheadrightarrow b$ *iff there is an object* c *in* $\mathcal{M}$ *with* $S\colon a \twoheadrightarrow c$ *and* $T\colon c \twoheadrightarrow b$.

Proof. (i) can be obtained by a straightforward argument using Lemma 2.2.10. Then (ii) follows from (i) and Lemma 2.2.10. (iii) and (iv) follow from (i) resp. (ii) using Definition 3.1.6. □

From this proposition and the definition of a segment of a sequence, we immediately get the following corollary that is important for induction proofs:

Corollary 3.1.16 (continuity/convergence of segments)
Let S *be a reduction sequence in an MRS. Then the following holds:*

(i) If S *is weakly (resp. strongly) continuous, then any segment of* S *is weakly (resp. strongly) continuous.*

(ii) If S *is weakly (resp. strongly) convergent, then any segment of* S *is weakly (resp. strongly) convergent.*

The following proposition reveals the intuition that continuity is simply the convergence of every proper prefix of limit ordinal length:

Proposition 3.1.17 (continuity and convergence)
Let $S = (a_\iota \to a_{\iota+1})_{\iota<\alpha}$ *be a reduction sequence in an MRS. Then the following are equivalent:*

(i) S *is weakly continuous.*

(ii) Each open proper prefix $S|_{[0,\lambda)}$ *of* S *weakly converges to* a_λ.

(iii) Each proper prefix $S|_{[0,\beta)}$ *of* S *weakly converges to* a_β.

The same holds for strong continuity/convergence.

Proof. The equivalence of (i) and (ii) follows from Lemma 2.2.13. The equivalence of (ii) and (iii) is obvious as weak convergence is trivial for closed weakly continuous reduction sequences.

The case of strong continuity/convergence follows easily from the case of weak continuity/convergence proved above. □

From this proposition we can easily obtain the following corollary that provides a characterisation of continuity of open reduction sequences.

Corollary 3.1.18 (continuity of open reduction sequences)
Let S *be an open reduction sequence in an MRS. Then the following holds:*

(i) S *is weakly continuous iff every proper prefix of* S *is weakly continuous.*

(ii) S *is strongly continuous iff every proper prefix of* S *is strongly continuous.*

Proof. (i) The “only if” direction follows from Proposition 3.1.17. For the converse direction, suppose that each proper prefix of $S = (a_\iota \to a_{\iota+1})_{\iota<\lambda}$ is weakly continuous. We will show that this implies that each proper prefix $S|_{[0,\alpha)}$ weakly converges to a_α. Applying Proposition 3.1.17 then yields that S is weakly continuous.

If $S|_{[0,\alpha)}$ is a proper prefix of S, then, according to Lemma 2.1.15, so is $S|_{[0,\alpha+1)}$. Hence, $S|_{[0,\alpha+1)}$ is weakly continuous. Since $S|_{[0,\alpha)}$ is a proper prefix of $S|_{[0,\alpha+1)}$, we can employ Proposition 3.1.17 to obtain that $S|_{[0,\alpha)}$ weakly converges to a_α.

(ii) Analogously. □

The following theorem provides further insight into the intuition of strong convergence.

Theorem 3.1.19 (characterisation of strong convergence)
Let S be a reduction sequence in an MRS $\mathcal{M}$.

(i) If S is strongly convergent, then, for any $h \in \mathbb{R}^+$, there are at most finitely many steps in S whose height is greater than h.

(ii) If S is weakly continuous and, for any $h \in \mathbb{R}^+$, there are at most finitely many steps in S whose height is greater than h, then S is strongly continuous. If, additionally, $\mathcal{M}$ is a complete URS, then S is even strongly convergent.

Proof. Let $S = (a_\iota \to_{h_\iota} a_{\iota+1})_{\iota<\alpha}$ be a reduction sequence in $\mathcal{M}$.

(i) (cf. [Ken92][1]) Suppose that S is strongly convergent and let $h \in \mathbb{R}^+$. We show by transfinite induction on α, the length of S, that the number of steps in S, which are at depth greater than h, is finite. The use of transfinite induction is justified by Corollary 3.1.16. The case $\alpha = 0$ is trivial. If α is a successor ordinal, the claim follows immediately from the induction hypothesis. Suppose that α is a limit ordinal. Since S is strongly convergent, we have $\lim_{\iota\to\alpha} h_\iota = 0$, i.e. there is some $\beta < \alpha$ such that $h_\iota < h$ for all $\beta < \iota < \alpha$. Hence, S has as many steps of depth greater than h as its proper prefix $S|_{[0,\beta)}$. By induction hypothesis, this number is finite.

(ii) Suppose that S is weakly continuous, and that the set $\{\iota \,|\, h_\iota > h\}$ is finite for any $h \in \mathbb{R}^+$. According to Definition 3.1.6, we have to show that $\lim_{\iota\to\lambda} h_\iota = 0$ for each limit ordinal $\lambda < \alpha$ in order to prove that S is strongly continuous. To this end, let $\varepsilon > 0$ be arbitrary. Then choose some h such that $0 < h < \varepsilon$. Since, by hypothesis, the set $\{\iota \,|\, h_\iota > h\}$ is finite, there is some ordinal $\beta < \lambda$ such that $h_\iota \le h < \varepsilon$ for all $\beta < \iota < \lambda$. Hence, $\lim_{\iota\to\lambda} h_\iota = 0$.

For the second part of the statement, we additionally assume $\mathcal{M}$ to be a complete URS. Moreover, we can assume that α is a limit ordinal since otherwise the statement is trivially true. The paragraph above proved strong continuity, i.e., according to Proposition 3.1.14, it remains to be shown, that $\lim_{\iota\to\alpha} h_\iota = 0$. Because α is a limit ordinal, we can employ the same argument as above to show this. ☐

From this theorem, the following corollary follows as shown in [Ken92].

Corollary 3.1.20 (strongly convergent reductions are of countable length, [Ken92])
A strongly convergent reduction sequence has countable length.

By employing an argument similar to the one used in [KdV05] for the particular case of infinitary term rewriting, we can generalise this to strongly *continuous* reduction sequences in complete URSs.

Proposition 3.1.21 (strongly continuous reductions are of countable length)
Every strongly continuous reduction sequences in a complete URS has countable length.

Proof. Suppose there is a strongly continuous reduction sequence $S = (m_\iota \to_{h_\iota} m_{\iota+1})_{\iota<\alpha}$ of uncountable length. By Corollary 3.1.20, we can assume that S is strongly divergent, and, by Corollary 3.1.16, we can assume that α is the least uncountable ordinal ω_1. As S is strongly divergent, we can apply Theorem 3.1.19 in order to obtain some $h \in \mathbb{R}^+$ such that the set $H = \{\gamma \in \omega_1 \,|\, h_\gamma > h\}$ is infinite. Hence, we can construct a strictly increasing ω-sequence $(\gamma_i)_{i<\omega}$ of elements in H. Let $\beta = \bigcup_{i<\omega} \gamma_\iota$. Due to Proposition 2.1.16, it holds that $\beta < \omega_1$. The proper prefix $S|_{[0,\beta)}$ of S then also contains infinitely many steps with height greater than h. Consequently, according to Theorem 3.1.19, also $S|_{[0,\beta)}$ is strongly divergent. By Proposition 3.1.17, this contradicts the fact that S is strongly continuous. ☐

[1] Note that the original result in [Ken92] refers to strongly *continuous* reduction sequences instead, which are defined in that paper differently. In our terminology, these sequences are actually strongly *convergent*. Since the argument of the original proof is rather sloppy, we present a modified argument.

The above theorem is not true for weakly convergent reduction sequences as the following example of an arbitrary long weakly convergent reduction sequence shows:

Example 3.1.22 ([Ken92])
Consider the MRS with a single object $\bullet$, equipped with the obvious metric, and a single step given by $\bullet \to_1 \bullet$. Then, for any ordinal α, the sequence $(\varphi_\iota\colon \bullet \to \bullet)_{\iota<\alpha}$ constitutes a weakly convergent reduction sequence.

3.2 Partial Reduction Systems

This section is along the lines of Section 3.1. We introduce an alternative method for defining meaningful reduction sequences of transfinite length. The key observation that illustrates the motivation of searching for an alternative is the existence of (weakly) divergent reduction sequences. Even in complete MRSs there are reduction sequences that do not (weakly) converge. This is due to the fact that completeness (of a metric space) only ensures that Cauchy sequences converge. The goal of considering partial reduction systems, which will be defined shortly, is to make the requirements for convergence less strict.

For example, suppose that we have a TRS consisting of the rules

$$f(x,a) \to f(s(x),b), \quad f(x,b) \to f(s(x),a).$$

Then we can construct the ω-reduction sequence

$$f(0,a) \to f(s(0),b) \to f(s(s(0)),a) \to f(s(s(s(0))),b) \to \dots$$

which is neither strongly nor weakly convergent in terms of its MRS semantics. The culprit is the second argument of the f symbol which constantly changes between a and b. However, excluding this "flickering", the reduction sequence seems to converge somehow. The investigation of *partial reduction systems* is aimed at formalising this relaxation of the notion of convergence. With this tool we will be able to identify $f(s^\omega,\bot)$ as the limit of the reduction sequence above.

To this end, a partially ordered set is employed rather than a metric space, and the limit construction is replaced by the limit inferior. The idea of considering the partial order on terms in order to model infinite reductions was first pursued by Corradini [Cor93, CG95], however, in a much more restricted setting. Blom [Blo04] investigated the use of the partial order on terms and its notion of limit inferior in the setting of the λ-calculus. Our investigation of partial reduction systems was mainly inspired by this work.

Definition 3.2.1 (partial reduction system)
A *partial reduction system* (*PRS*) is a tuple $\mathcal{P} = (A,\Phi,\mathsf{src},\mathsf{tgt},\leq,\mathsf{cxt})$ such that

(i) $\mathcal{A} = (A,\Phi,\mathsf{src},\mathsf{tgt})$ is an ARS, called the *underlying ARS* of $\mathcal{P}$,

(ii) $(A,\leq)$ is a partially ordered set,

(iii) $\mathsf{cxt}\colon \Phi \to A$ is a function, called the *context function*, and

(iv) if $\varphi\colon a \to_{\mathcal{A}} b$, then $\mathsf{cxt}(\varphi) \leq a,b$.

If the partial order $\leq$ is a complete semilattice, then $\mathcal{P}$ is called *complete*.

Notice the similarities of PRSs and MRSs: Both contain an ARS core, of course. The underlying set of the ARS is endowed with a metric in MRSs and with a partial order in PRSs. They are needed in order to define the limit behaviour of transfinite sequences. Additionally, each step in an MRS has a designated height which overapproximates the distance between the two objects involved in the reduction step. An analogous concept does exist in PRSs as well: To each step a context is assigned. The intuitive meaning of this context is that it represents some information or structure that is shared between the two objects involved in the reduction step. This is, in fact, how clause (iv) of Definition 3.2.1 is meant to be interpreted. Shortly, we will see what this means for the concrete example of term rewriting systems. When reduction sequences will be defined further below, we will also learn that the context function of PRSs, similarly to the height function of MRSs, is used to distinguish between weak and strong continuity resp. convergence. If the PRS under consideration is complete, e.g. in the case of ITRSs, the similarity of contexts in PRSs and heights in MRS becomes even more evident: Since the glb of arbitrary sets is always defined for complete semilattices, we can, in fact, rephrase clause (iv) of Definition 3.2.1 as follows:

$$\text{If } \varphi\colon a \to_{\mathcal{A}} b, \text{ then } \mathsf{cxt}(\varphi) \leq a \sqcap b.$$

Intuitively, $a \sqcap b$ represents the common structure/information of a and b. Therefore, $\mathsf{cxt}(\varphi)$ is an underapproximation of the shared structure/information of a and b. In other words: Similarly to the height for MRS, also the context possibly overestimates the difference between the two objects.

Notation 3.2.2. Similar to the case of MRSs we confuse PRSs with their underlying ARSs in order to use the notation that we already have for ARSs. In particular, we write $\varphi\colon a \to_{\mathcal{P}} b$ for a reduction step $\varphi\colon a \to_{\mathcal{A}} b$ in the underlying ARS $\mathcal{A}$. If we want to explicitly indicate the context of a reduction step in $\mathcal{P}$, we write $\varphi\colon a \to_c b$ whenever $\mathsf{hgt}(\varphi) = c$.

To define PRSs in a convenient, way we use an approach similar to that for MRSs: Instead of statements of the form $a \to_h b$, with $h \in \mathbb{R}^+$, we consider statements of the form $a \to_c b$ with $c \in A$. c is then interpreted as the context of the defined reduction step, i.e. $\mathsf{cxt}(a \to_c b) = c$.

As promised, we want to present the PRS semantics of ITRSs in order to give an intuition of the intention of the rather abstract definition of PRSs

Definition 3.2.3 (PRS semantics of ITRSs)
Let $\mathcal{R} = (\Sigma, R)$ be an ITRS. The PRS *induced* by $\mathcal{R}$, denoted $\mathcal{P}_{\mathcal{R}}$, is given by the tuple

$$(\mathcal{T}^\infty(\Sigma_\perp, \mathcal{V}), \Phi, \mathsf{src}, \mathsf{tgt}, \leq_\perp, \mathsf{cxt}),$$

where $(\mathcal{T}^\infty(\Sigma_\perp, \mathcal{V}), \Phi, \mathsf{src}, \mathsf{tgt})$ is the ARS $\mathcal{A}_{\mathcal{R}'}$ induced by the ITRS $\mathcal{R}' = (\Sigma_\perp, R)$, $\leq_\perp$ is the usual partial order on $\mathcal{T}^\infty(\Sigma_\perp, \mathcal{V})$, and cxt is defined as

$$\mathsf{cxt}(\varphi) = t[\perp]_\pi, \text{ where } \varphi\colon t \to_{\pi,\rho} t'.$$

The definition of the PRS semantics of ITRSs shows the intention of the context of a reduction step: It is, very literally, the context where the rewrite rule is applied and, thus, represents the structure that is necessarily shared between the original term and its contraction. The proof of the following proposition establishes this observation formally:

Proposition 3.2.4 (PRS semantics yields a complete PRS)
Each ITRS $\mathcal{R}$ induces a complete PRS $\mathcal{P}_{\mathcal{R}}$.

Proof. Since $\leq_\perp$ is a complete semilattice on $\mathcal{T}^\infty(\Sigma_\perp, \mathcal{V})$, it remains to be shown that for each rewrite step $\varphi\colon t \to_c t'$ it holds that $c \leq_\perp t, t'$. Let $l \to r$ be the rule applied in φ. Then there is a substitution σ and a context $C[]$ such that $t = C[l\sigma]$, $t' = C[r\sigma]$ and $c = C[\perp]$. As $\perp \leq_\perp l\sigma, r\sigma$ and $\leq_\perp$ is monotone, we can conclude that $C[\perp] \leq_\perp C[l\sigma], C[r\sigma]$. ☐

Next we define continuity and convergence for reduction sequences in PRSs. This definition follows the structure of Definition 3.1.6.

Definition 3.2.5 (continuity, convergence)
Let $S = (a_\iota \to_{c_\iota} a_{\iota+1})_{\iota<\alpha}$ be a reduction sequence in a PRS $\mathcal{P}$.

(i) S is called *weakly continuous* if $\liminf_{\iota \to \lambda} a_\iota = a_\lambda$ holds for each limit ordinal $\lambda < \alpha$. If, instead, $\liminf_{\iota \to \lambda} c_\iota = a_\lambda$ holds for each limit ordinal $\lambda < \alpha$, then the sequence is called *strongly continuous*.

(ii) A weakly continuous reduction sequence $S = (a_\iota \to_{c_\iota} a_{\iota+1})_{\iota<\alpha}$ is called *weakly convergent* if it is closed or, if it is open and $a_\alpha = \liminf_{\iota \to \alpha} a_\iota$ exists. A strongly continuous reduction sequence $S = (a_\iota \to_{c_\iota} a_{\iota+1})_{\iota<\alpha}$ is called *strongly convergent* if it is closed or, if it is open and $a_\alpha = \liminf_{\iota \to \alpha} c_\iota$ exists. In the case of weak (resp. strong) convergence, we also say that the reduction sequence S weakly (resp. strongly) converges to a_α or, alternatively, that a_0 weakly (resp. strongly) converges to a_α in α steps. A weakly (resp. strongly) continuous reduction sequence is called *weakly* (resp. *strongly*) *divergent* if it is not weakly (resp. strongly) convergent.

(iii) $\mathcal{P}$ is called *weakly α-convergent* or *strongly α-convergent*, respectively, if any weakly continuous resp. strongly continuous α-reduction sequence in $\mathcal{P}$ is weakly convergent resp. strongly convergent.

We now have a notion of reduction sequences of PRSs which is similar to that of MRSs. Therefore, we use the same notations for reduction sequences and their properties. In particular, the notations introduced in Notation 3.1.10 as well as the Remarks 3.1.7, 3.1.11 and 3.1.9 also apply to PRSs.

Before we continue our discussion of PRSs let us have a look at an example that illustrates the difference between weak and strong convergence for the particular case of TRSs. It will also show the difference to MRSs.

Example 3.2.6
Consider the TRS with the single rule $f(x,y) \to f(y,x)$. This rule induces the following reduction sequence:

$$S\colon f(a, f(g(a), g(b))) \to f(a, f(g(b), g(a))) \to f(a, f(g(a), g(b))) \to \dots$$

S simply alternates between the terms $f(a, f(g(a), g(b)))$ and $f(a, f(g(b), g(a)))$ by swapping the arguments of the inner f occurrence. The reduction sequence is depicted in Figure 3.3. The picture illustrates the parts of the terms that remain unchanged and those that remain completely untouched by the corresponding reduction step in the same way as in Figure 3.2, i.e. by using a lighter shade resp. a darker shade of grey. The unchanged part corresponds to the glb of the two terms of a reduction step, viz. for the first step

$$f(a, f(g(a), g(b))) \sqcap_\bot f(a, f(g(n), g(a))) = f(a, f(g(\bot), g(\bot)))$$

By symmetry, the glb of the terms of the second step is the same one. It is depicted in Figure 3.4a. Let $(t_i)_{i<\omega}$ be the sequence of terms of the reduction S. By definition, S weakly converges to $\liminf_{i<\omega} t_i$. According to Proposition 2.1.34, this is equal to $\liminf_{i<\omega}(t_i \sqcap_\bot t_{i+1})$. Since $t_i \sqcap_\bot t_{i+1}$ is constantly $f(a, f(g(\bot), g(\bot)))$, the reduction sequence weakly converges to $f(a, f(g(\bot), g(\bot)))$.

Similarly, the part of the term that remains untouched by the reduction step corresponds to the context. For the first step, this is $f(a, \bot)$. It is depicted in Figure 3.4b. By definition, S strongly converges to $\liminf_{i<\omega} c_i$ for $(c_i)_{i<\omega}$, the sequence of contexts of S. As one can see in Figure 3.3, the context constantly remains $f(a, \bot)$. Hence, S strongly converges to $f(a, \bot)$. The example sequence is a particularly simple one as both the glbs $t_i \sqcap_\bot t_{i+1}$ and the contexts c_i remain stable. In general, this is not necessary, of course.

An important difference to MRSs is that in complete PRSs every weakly or strongly continuous reduction sequence is also weakly resp. strongly convergent:

Fact 3.2.7 (continuity implies convergence)
In a complete PRS,

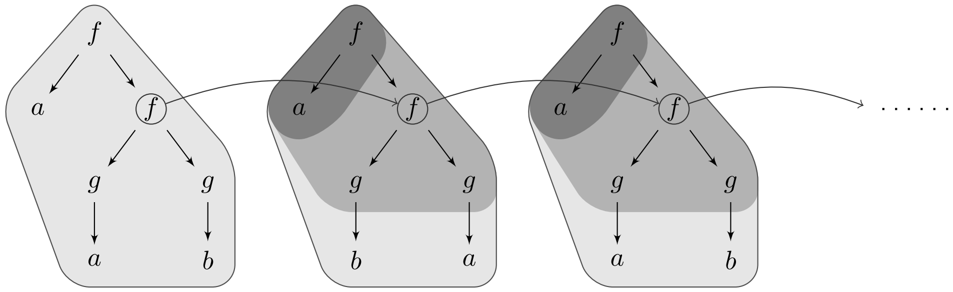

Figure 3.3: Reduction sequence with stable context.

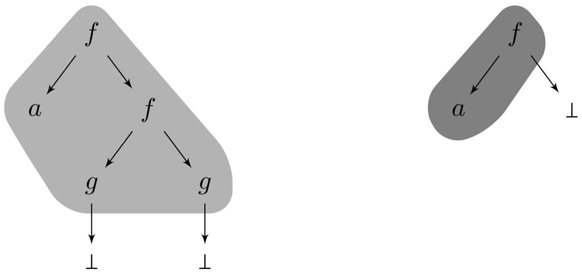

(a) Weak limit.

(b) Strong limit.

Figure 3.4: Limits of the PRS reduction sequence.

(i) every weakly continuous reduction sequence is weakly convergent, and

(ii) every strongly continuous reduction sequence is strongly convergent.

Proof. This follows immediately from Proposition 2.1.33. □

Another very important difference between the metric and the partial order model can be found in the nature of the distinction between weak and strong continuity/convergence. Strongly continuous/convergent reduction sequences in MRSs are, by definition, weakly continuous/convergent reduction sequences satisfying an additional property. That is, strong continuity/convergence is simply a more restrictive variant of weak continuity/convergence. The limit construction itself is the same. On the other hand, in PRSs the limit construction in weakly continuous/convergent reduction is different from the one used for strong continuity/convergence. The former considers the objects of the reduction sequence, the latter uses the contexts of the reduction steps. Note that clause (iv) of Definition 3.2.1 relates the objects and the context of a reduction sequence. However, this is not sufficient to ensure that strongly continuous/convergent reduction sequences are also weakly continuous/convergent in general as it is the case for MRSs.

For example, consider the PRS $\mathcal{P}$ given by the rules $a \to_\perp a$ and $\perp \to_\perp a$, where the underlying set A is $\{a, \perp\}$ with the partial order $\leq$ induced by $\perp \leq a$. Then the reduction sequence

$$a \to_\perp a \to_\perp a \to_\perp \ldots \perp \to_\perp a$$

is strongly continuous as the sequence of contexts $\perp, \perp, \ldots$ is constant and, thus, converges (via limit inferior) to $\perp$, the ω-th element in the sequence. On the other hand, this reduction sequence is not weakly continuous as the sequence of elements $a, a, \ldots$ converges to a and, therefore, not to the ω-th element of the sequence.

A similar counterexample exists for convergence: Consider the PRS $\mathcal{P}$ given by the rules $n \to_0 n+1$ for each $n \in \mathbb{N}$, where the underlying set is $\mathbb{N}$ with the natural partial order on them. Note that the 0 that is indicated for the reduction steps is supposed to be its context and not its height. The reduction sequence

$$0 \to_0 1 \to_0 2 \to_0 \dots$$

does strongly converge to 0 whereas it does not weakly converge. This is rather an artifact of the fact that the partial order on $\mathbb{N}$ is not a complete semilattice. But even if we extend the order to a complete semilattice (e.g. by adjoining a greatest element ∞), the reduction sequence is indeed weakly and strongly convergent. Yet, the respective limits are different: 0 for strong convergence and ∞ for weak convergence.

In order for PRSs to exhibit the same relation between strong and weak continuity resp. convergence that MRSs show, we have to consider total reductions:

Definition 3.2.8 (total reduction sequence)
Let $\mathcal{P}$ be a PRS and $S = (a_\iota \to a_{\iota+1})_{\iota<\alpha}$ a reduction sequence in $\mathcal{P}$. S is called *total* if each element a_ι is maximal w.r.t. the partial order of $\mathcal{P}$. If we write S as $S\colon a_0 \hookrightarrow a_\alpha$ or $S\colon a_0 \twoheadrightarrow a_\alpha$, i.e. the convergence of the reduction sequence is explicitly stated, we additionally require a_α to be maximal for S to be total.

For total reduction sequences, we indeed have that strong continuity/convergence implies weak continuity/convergence:

Proposition 3.2.9 (strong cont./conv. implies weak cont./conv.)
For every total reduction sequence S in a PRS, it holds that

(i) $S\colon a \twoheadrightarrow \dots$ *implies* $S\colon a \hookrightarrow \dots$*, and that*

(ii) $S\colon a \twoheadrightarrow b$ *implies* $S\colon a \hookrightarrow b$.[2]

Proof. Let $S = (a_\iota \to_{c_\iota} a_{\iota+1})_{\iota<\alpha}$

(i). If S is strongly continuous, then $\liminf_{\iota\to\lambda} c_\iota = a_\lambda$ for each limit ordinal $\lambda < \alpha$. We need to show that then also $\liminf_{\iota\to\lambda} a_\iota = a_\lambda$. By definition, we have $c_\iota \le a_\iota$ for each ι. Hence, also $a_\lambda = \liminf_{\iota\to\lambda} c_\iota \le \liminf_{\iota\to\lambda} a_\iota$ holds. Note that a_λ is a maximal element as S is total. Hence, we can conclude that $\liminf_{\iota\to\lambda} a_\iota = a_\lambda$.

(ii) Suppose that S strongly converges to b. Then S is also strongly continuous and, by (i), weakly continuous. If S is closed, then $b = a_\alpha$ and S also weakly converges to b. If S is open, then $b = \liminf_{\iota\to\alpha} c_\iota$. By the same argument used for (i), we obtain that $b = \liminf_{\iota\to\alpha} a_\iota$, i.e. S weakly converges to b. ☐

Analogously to Proposition 3.1.15 for MRSs, we can easily derive the following properties for concatenations of reduction sequences in PRSs:

Proposition 3.2.10 (concatenation of partial reduction sequences)
Let S and T be two non-empty reduction sequences over the same PRS $\mathcal{P}$, and a, b two objects in $\mathcal{P}$.

(i) $S \cdot T\colon a \hookrightarrow \dots$ *iff there is an object c in $\mathcal{P}$ with* $S\colon a \hookrightarrow c$ *and* $T\colon c \hookrightarrow \dots$

(ii) $S \cdot T\colon a \hookrightarrow b$ *iff there is an object c in $\mathcal{P}$ with* $S\colon a \hookrightarrow c$ *and* $T\colon c \hookrightarrow b$

(iii) $S \cdot T\colon a \twoheadrightarrow \dots$ *iff there is an object c in $\mathcal{P}$ with* $S\colon a \twoheadrightarrow c$ *and* $T\colon c \twoheadrightarrow \dots$

(iv) $S \cdot T\colon a \twoheadrightarrow b$ *iff there is an object c in $\mathcal{P}$ with* $S\colon a \twoheadrightarrow c$ *and* $T\colon c \twoheadrightarrow b$

[2]Here it is important that the totality of S also requires the final element b to be maximal.

Proof. This follows immediately from the definition of reduction sequences and continuity resp. convergence. □

Again this yields the following corollary – now for PRSs:

Corollary 3.2.11 (continuity/convergence of segments)
Let S be a reduction sequence over some PRS. Then the following holds:

(i) If S is weakly (resp. strongly) continuous, then any segment of S is weakly (resp. strongly) continuous.

(ii) If S is weakly (resp. strongly) convergent, then any segment of S is weakly (resp. strongly) convergent.

Also the relation between continuity and convergence in PRSs is the same as in MRSs:

Proposition 3.2.12 (continuity and convergence)
Let $S = (a_\iota \to a_{\iota+1})_{\iota<\alpha}$ be a reduction sequence over some PRS. Then the following are equivalent:

(i) S is weakly continuous.

(ii) Each open proper prefix $S|_{[0,\lambda)}$ of S weakly converges to a_λ.

(iii) Each proper prefix $S|_{[0,\beta)}$ of S weakly converges to a_β.

The same holds for strong continuity/convergence.

Proof. The implication from (i) to (ii) follows from Corollary 3.2.11, which asserts that the prefix $S|_{[0,\lambda)}$ is continuous as well, and the definition of weak continuity which requires that $\liminf_{\iota\to\lambda} a_\iota = a_\lambda$.

Consider the converse implication: If each open proper prefix of S is weakly convergent, then $\liminf_{\iota\to\lambda} a_\iota = a_\lambda$ holds for each limit ordinal $\lambda < \alpha$. Hence, S is continuous.

Items (ii) and (iii) are obviously equivalent as weak convergence is trivial for closed reduction sequences.

For the case of strong continuity/convergence, the argument is analogous. □

Corollary 3.2.13 (continuity of open reduction sequences)
Let S be an open reduction sequence in an PRS. Then the following holds:

(i) S is weakly continuous iff every proper prefix of S is weakly continuous.

(ii) S is strongly continuous iff every proper prefix of S is strongly continuous.

Proof. This can be proved in the same way as Corollary 3.1.18 using Proposition 3.2.12 instead of Proposition 3.1.17. □

Recall that strongly continuous reductions in MRSs have at most countable ordinal length. Such an upper bound for the length of reduction sequences does not exists for PRSs – neither for weakly nor for strongly convergent reductions. The following example illustrates this:

Example 3.2.14
Let $\mathcal{P}$ be a PRS on the singleton set $\{\bullet\}$ with the trivial partial order on it and let $\mathcal{P}$ have the reduction step $\varphi : \bullet \to_\bullet \bullet$. Then any constant sequence containing φ is a weakly and also strongly convergent reduction sequence.

3.3 Transfinite Abstract Reduction Systems

In this section we want to discuss both MRSs and PRSs. That is, we want to introduce notions that apply to both systems. To this end, we use the term *transfinite abstract reduction system* (*TARS*) in order to refer to an MRS or a PRS. In the second part of this section we also introduce some abstract criteria which allow us to relate MRSs to PRSs.

3.3.1 Properties of Transfinite Reductions

Now that we have several meaningful notions of transfinite reductions, we also want to investigate their properties. Therefore, the question arises which properties of transfinite reductions are of interest and might have a practical impact.

In principle, all properties known from the finitary setting can be lifted to the present setting of transfinite reductions (cf. [Ken92]): Simply replace $\to^{\star}$ by $\hookrightarrow$ resp. $\twoheadrightarrow$ in the definition of the property. Confluence (CR), for example, is defined as

$$\forall\ b \leftarrow^{\star} a \to^{\star} c \quad \Longrightarrow \quad \exists\ b \to^{\star} d \leftarrow^{\star} b$$

This can be lifted to *infinitary confluence* (CR^{∞}) by requiring

$$\forall\ b \twoheadleftarrow a \twoheadrightarrow c \quad \Longrightarrow \quad \exists\ b \twoheadrightarrow d \twoheadleftarrow b$$

when considering strong convergence, or

$$\forall\ b \hookleftarrow a \hookrightarrow c \quad \Longrightarrow \quad \exists\ b \hookrightarrow d \hookleftarrow b$$

when considering weak convergence. Usually, it is obvious from the context which variant, weak or strong, is meant. If not, then we additionally state whether we are considering weak or strong reductions. The illustrated lifting to infinitary properties can also be applied to the finitary properties WN and $\mathsf{UN}_{\to}$ which yields *infinitary normalisation* (WN^{∞}) and the *infinitary unique normal form property w.r.t. reduction* ($\mathsf{UN}^{\infty}_{\to}$), respectively.

Recall that, for the finitary properties NF and UN, the notion of convertibility is needed. Hence, we need to find an appropriate infinitary variant of convertibility in order to obtain infinitary versions of NF and UN. Devising a concept of infinitary convertibility is not straightforward and has to be done carefully as we will see.

The convertibility relation $\leftrightarrow^{\star}$ for ARSs is defined by taking finite reduction sequences of the symmetric closure of the original ARS (cf. Definition 2.3.6). The construction of the symmetric closure of an ARS can be straightforwardly generalised to TARSs. Infinitary convertibility can then be defined as the (possibly transfinite) convergent reduction sequences of this symmetric closure. That is, assuming we use $\leftrightarrow^{\infty}$ for infinitary convertibility, we define $a \leftrightarrow^{\infty}_{\mathcal{T}} b$ iff $a \hookrightarrow_{\mathcal{T}^s} b$ resp. $a \twoheadrightarrow_{\mathcal{T}^s} b$, for $\mathcal{T}^s$ the symmetric closure of $\mathcal{T}$. However, one can argue that this is not a reasonable choice for the notion of convertibility:

Example 3.3.1
Consider the TRS with the rules

$$f(0) \to 0, \quad f(1) \to 1$$

Let $\mathcal{T}$ be the induced MRS or PRS of the above system and $\mathcal{T}^s$ its symmetric closure. That is, $\mathcal{T}^s$, in particular, permits rewriting steps of the form $f^{n+1}(0) \leftarrow f^n(0)$ and $f^{n+1}(1) \leftarrow f^n(1)$. Hence, we have the following two reduction sequences in $\mathcal{T}^s$ which are both weakly and strongly convergent for both the MRS and the PRS semantics:

$$0 \to f(0) \to f(f(0)) \to f(f(f(0))) \to \ \dots\ f^{\omega}$$
$$1 \to f(1) \to f(f(1)) \to f(f(f(1))) \to \ \dots\ f^{\omega}$$

That is, we have $0 \leftrightarrow^{\infty} f^{\omega}$ as well as $1 \leftrightarrow^{\infty} f^{\omega}$. However, f^{ω} is a normal form in the symmetric system $\mathcal{T}^s$. Hence, there is no convergent sequence from f^{ω} to 0 or 1. Therefore, $\leftrightarrow^{\infty}$ is not symmetric.

$$\begin{tikzcd} & a \arrow[dl, twoheadrightarrow] \arrow[dr, twoheadrightarrow] & \\ b \arrow[rr, dashed, twoheadrightarrow] & & c \in \mathsf{NF} \end{tikzcd}$$

(a) NF^∞ according to [KKSdV95a].

$$\begin{tikzcd} b \arrow[r, leftrightarrow, bend left, "s"] \arrow[r, dashed, twoheadrightarrow, bend right] & c \in \mathsf{NF} \end{tikzcd}$$

(b) NF^∞ as defined here.

Figure 3.5: Alternative definition of the infinitary normal form property.

Even if we consider the symmetric closure of $\leftrightarrow^\infty$, the resulting relation would not be transitive: We have $0 \leftrightarrow^\infty f^\omega$ and $1 \leftrightarrow^\infty f^\omega$ but neither $1 \leftrightarrow^\infty 0$ nor $0 \leftrightarrow^\infty 1$.[3]

Because of the problem illustrated above, we, instead, define infinitary convertibility simply as the reflexive transitive symmetric closure of $\hookrightarrow$ resp. $\twoheadrightarrow$. This is a reasonable generalisation of the finitary case as $\leftrightarrow^\star$ is also the reflexive transitive symmetric closure of $\to^\star$.

Definition 3.3.2 (infinitary convertibility, [Ken92])
Let $\mathcal{T}$ be a TARS, and a, b objects in $\mathcal{T}$. a and b are called weakly resp. strongly convertible, written $a \leftrightarrow^w_\mathcal{T} b$ resp. $a \leftrightarrow^s_\mathcal{T} b$, whenever there is a finite sequence of objects $a_0, \dots, a_n$, $n \geq 0$, in $\mathcal{M}$ such that $a_0 = a$, $a_n = b$, and, for each $0 \leq i < n$, we have $a_i \hookrightarrow_\mathcal{T} a_{i+1}$ resp. $a_i \twoheadrightarrow_\mathcal{T} a_{i+1}$ or $a_i \hookleftarrow_\mathcal{T} a_{i+1}$ resp. $a_i \twoheadleftarrow_\mathcal{T} a_{i+1}$. The minimal n of such a sequence is called the length of $a \leftrightarrow^w_\mathcal{T} b$ resp. $a \leftrightarrow^s_\mathcal{T} b$.

With this notion of convertibility we can establish an alternative characterisation of CR^∞ which is analogous to its finitary version:

Proposition 3.3.3 (alternative characterisation of CR^∞)
Let $\mathcal{T}$ be a TARS. $\mathcal{T}$ is CR^∞ (for weak reductions) iff

$$\forall\ a \leftrightarrow^w b \quad \Longrightarrow \quad \exists\ a \hookrightarrow c \hookleftarrow b$$

The same also holds for strong reductions.

Proof. The argument is the same as for finitary reductions: The "if" direction is trivial, and the "only if" direction can be proved by and induction on the length of $a \leftrightarrow^w b$. □

Remark 3.3.4. The definition of NF^∞, the infinitary version of NF, that we are using is: If $b \leftrightarrow^s c$ for a normal form c, then $b \twoheadrightarrow c$; cf. Figure 3.5b. In [KKSdV95a] an alternative definition of NF^∞ is employed: If $b \twoheadleftarrow a \twoheadrightarrow c$ for a normal form c, then $b \twoheadrightarrow c$; cf. Figure 3.5a. However, both variants are equivalent. The implication from NF^∞ to the alternative variant is trivial. For the converse direction, assume that the system enjoys the alternative NF^∞ property and that $b \leftrightarrow^s c$ for a normal form c. A straightforward induction on the length of $b \leftrightarrow^s c$ shows that then $b \twoheadrightarrow c$ holds.
The same argument is, of course, also valid for weak reductions.

Proposition 3.3.5 (confluence properties)
For every TARS, we have the following implications (for both weak and strong reductions):

(i) $\mathsf{CR}^\infty \quad \Longrightarrow \quad \mathsf{NF}^\infty \quad \Longrightarrow \quad \mathsf{UN}^\infty \quad \Longrightarrow \quad \mathsf{UN}^\infty_\to$

(ii) $\mathsf{WN}^\infty\ \&\ \mathsf{UN}^\infty_\to \quad \Longrightarrow \quad \mathsf{CR}^\infty$

[3] In [Ken92] the very same example is presented for the same purpose. Its author, however, spuriously argues that it shows that 1, 0 and f^ω are pairwise infinitarily convertible and, thereby, questions the utility of this notion of infinitary convertibility.

Proof. The arguments are the same as for their finitary variants. □

Just as for the properties involving convertibility, we also have to be careful when defining the infinitary version of the termination property. In its finitary version, termination prohibits infinite reduction sequences. This is certainly not what we want for an infinitary termination property. As pointed out in [KdV05], another way of interpreting finitary termination is, that it guarantees that no matter how one extends an existing reduction sequence, in the end one always reaches a normal form. For transfinite reductions, we have to take into account that they can be extended by a limit construction. But such a construction can prohibit further extension of the thus obtained reduction sequence although it might not end in a normal form. This is precisely the case when we have a reduction sequence that is continuous but not convergent. Hence, infinitary termination has to require that any continuous reduction sequences is convergent as well. In a nutshell, the key difference to the finitary case is that there are not only two possibilities for the final object of a reduction sequence, viz. being a normal form or not being a normal form, but, in fact, there are tree: Either it is a normal form, it is not a normal form, or it is not existent due to divergence.

But there is another problem that might prohibit further extension of a reduction sequences despite the fact that no normal form was reached, yet: Consider the MRS from Example 3.1.22. It allows any reduction sequence to be extended to any arbitrary length. Yet, we can not extend these sequences by a limit construction, as we simply ran out of ordinal numbers! Hence, we have to additionally require for infinitary termination that there is an upper bound on the length of continuous reduction sequences. Intuitively speaking, just as non-terminating reduction sequences in the finitary setting leave the realms of ω, the finite ordinal numbers, non-terminating reduction sequences in the infinitary sense might leave the realms of On, the ordinal numbers.[4]

Definition 3.3.6 (infinitary termination property, [Rod98])
Let $\mathcal{T}$ be an TARS and a an object in $\mathcal{T}$. a is said to be *infinitarily terminating* (SN^∞) if

(a) every weakly (resp. strongly) continuous reduction sequence in $\mathcal{T}$ starting from a is weakly (resp. strongly) convergent, and

(b) the lengths of all weakly (resp. strongly) continuous reduction sequences starting from a are bounded from above, i.e. there is some ordinal β such that $a \hookrightarrow^\alpha \dots$ (resp. $a \twoheadrightarrow^\alpha \dots$) implies $\alpha < \beta$.

The TARS $\mathcal{T}$ itself is called *infinitarily terminating* (SN^∞) if every object in $\mathcal{T}$ is.

Remark 3.3.7. For strong reductions in MRSs, condition (b) of Definition 3.3.6 is always met due to Corollary 3.1.20. Therefore, the definition is equivalent to the one presented in [KdV05]. On the other hand, for complete PRSs, condition (a) of Definition 3.3.6 is always satisfied for both weak and strong reductions as Fact 3.2.7 shows.

The definition of the infinitary termination property is rather involved compared to its finitary version. The litmus test for its appropriateness is that it has to imply infinitary normalisation.

Proposition 3.3.8 (SN^∞ is stronger that WN^∞)
For every TARS, it holds that SN^∞ *implies* WN^∞ *for both weak and strong reductions.*

Proof. We prove the implication for weak reductions in MRSs. The other cases follow by the same argument.

In fact, we prove the contraposition of the implication. For this purpose, let $\mathcal{M}$ be an MRS and a some object in $\mathcal{M}$ that is not WN^∞. We show that then (a) or (b) of

[4] Strictly speaking, these "sequences" are, therefore, not sequences but rather a functions with domain On.

Definition 3.3.6 is violated. For this purpose, we assume (a) and show that then (b) does not hold. That is, we have to prove that there is no upper bound for the length of weakly continuous reductions in $\mathcal{M}$. We do this by defining a function f on the class On of ordinal numbers such that, for each $\alpha \in \mathsf{On}$, (1) $f(\alpha)$ is a weakly continuous reduction of length α starting in a and (2) $f(\alpha)$ is an extension of $f(\iota)$ for all $\iota < \alpha$. The construction of f is verified by the principle of transfinite recursion, and the properties (1) and (2) are established by transfinite induction (cf. Proposition 2.1.17).

For $\alpha = 0$, both (1) and (2) are trivial. Let α be a successor ordinal $\beta + 1$. By induction hypothesis, we have $f(\beta)\colon a \hookrightarrow^{\beta} \ldots$. By (a), there is some object b in $\mathcal{M}$ such that $f(\beta)\colon a \hookrightarrow^{\beta} b$. Since a is not WN^{∞}, b cannot be a normal form. Hence, there is a step $\varphi\colon b \to b'$ in $\mathcal{M}$. Define $f(\alpha) = f(\beta)\cdot\varphi$. That is, $f(\alpha)\colon a \hookrightarrow^{\alpha} b'$ which shows (1). (2) follows from the induction hypothesis since $f(\beta) \le f(\alpha)$.

Let α be a limit ordinal. According to the induction hypothesis, $f(\iota)\colon a \hookrightarrow^{\iota} \ldots$ holds for each $\iota < \alpha$. Moreover, since, by the induction hypothesis, also (2) holds for all $f(\iota)$, we have that $F = \{f(\iota) \mid \iota < \alpha\}$ is directed. Hence, $f(\alpha) = \bigsqcup F$ is well-defined according to Proposition 2.1.28. Consequently, all elements in F are prefixes of $f(\alpha)$. This shows (2) and, additionally, it shows that $f(\alpha)$ is a reduction sequence of length α starting in a. Moreover, we can employ Corollary 3.1.18 in order to conclude that $f(\alpha)$ is weakly continuous. □

The converse, however, is not true in general. As indicated before, Example 3.1.22 provides a counterexample for weak reductions in MRSs. The following example illustrates this for strong reductions:

Example 3.3.9
Consider the MRS $\mathcal{M}$ with object set $\{a, b\}$, equipped with an arbitrary metric (which must also be an ultrametric), and the steps given by $a \to_1 a$, and $a \to_h b$ with $h = \mathbf{d}(a, b)$. $\mathcal{M}$ is WN^{∞} as b is a normal form and, a reduces to b in a single step. Yet, a is not SN^{∞} as it has the strongly continuous reduction sequence $a \to_1 a \to_1 \ldots$ which is not strongly convergent.

Likewise, counterexamples can be constructed for PRSs as well.

3.3.2 Relating MRSs to PRSs

We have seen, especially in the preceding section, that reduction sequences in MRSs and PRSs share a variety of properties. In order to make a systematic analysis of both kinds of systems easier, it is desirable to relate MRSs to PRSs in such a way, that we are able to transfer certain results from MRSs to corresponding PRSs and vice versa. In this section we want to identify some sufficient criteria that ensure that the weakly/strongly continuous/convergent reduction sequences of an MRS are a well-defined subclass of weakly/strongly continuous/convergent reduction sequences in a corresponding PRS. The first set of criteria concerns the metric spaces and partially ordered sets which are the foundation of the limit construction of MRSs and PRSs, respectively. The second set of criteria relates the reduction steps of the MRS and the PRS:

Definition 3.3.10 (extension of metric spaces by partially ordered sets)

(i) Let $(A, \mathbf{d})$ be a metric space and let $(B, \le)$ be a partially ordered set. $(B, \le)$ is said to *extend* $(A, \mathbf{d})$ if

(1) $A = \{a \in B \mid a \text{ is maximal w.r.t. } \le\}$ (maximality)

(2) for all sequences $(a_\iota)_{\iota<\alpha}$ in A, it holds that (limits)

(a) $\lim_{\iota\to\alpha} a_\iota$ exists whenever $\liminf_{\iota\to\alpha} a_\iota \in A$, and

(b) $\liminf_{\iota\to\alpha} a_\iota = \lim_{\iota\to\alpha} a_\iota$ whenever $\lim_{\iota\to\alpha} a_\iota$ exists.

$(B, \le)$ is then called an *extension* of $(A, \mathbf{d})$.

(ii) Let $\mathcal{M} = (A, \Phi, \mathsf{src}, \mathsf{tgt}, \mathbf{d}, \mathsf{hgt})$ be an MRS and $\mathcal{P} = (B, \Phi', \mathsf{src}', \mathsf{tgt}', \leq, \mathsf{cxt})$ be a PRS. $\mathcal{P}$ is said to *extend* $\mathcal{M}$ if

(1) $(B, \leq)$ extends $(A, \mathbf{d})$,

(2) $\Phi \subseteq \Phi'$,

(3) $\mathsf{src}(\varphi) = \mathsf{src}'(\varphi)$ and $\mathsf{tgt}(\varphi) = \mathsf{tgt}'(\varphi)$ for all $\varphi \in \Phi$,

(4) $\mathsf{src}'(\varphi) \in B \setminus A$ for all $\varphi \in \Phi' \setminus \Phi$, and

(5) $S\colon a \twoheadrightarrow_{\mathcal{P}} b$ is total iff $S\colon a \twoheadrightarrow_{\mathcal{M}} b$ for any open reduction sequence S in $\mathcal{P}$.

$\mathcal{P}$ is then also called an *extension* of $\mathcal{M}$. If all conditions except (5) are met, then $\mathcal{P}$ is said to *weakly extend* $\mathcal{M}$; $\mathcal{P}$ is called a *weak extension* of $\mathcal{M}$. In order to emphasise this distinction, we sometimes also say *strong extension* instead of just extension.

Condition (1) of Definition 3.3.10(ii) guarantees that the metric space consists of the maximal elements of the ordered set and that the limit of a sequence of elements is equal to the limit inferior of the same sequence. Moreover, by the conditions (2) and (3), any reduction sequence in an MRS $\mathcal{M}$ is also a reduction sequence in any PRS $\mathcal{P}$ that extends $\mathcal{M}$. Condition (4) makes sure that the PRS $\mathcal{P}$ does not introduce new reduction steps unless they involve elements that do only occur in $\mathcal{P}$. And finally, (5) explicitly stipulates the relation between strongly convergent reduction sequences in both systems.

These condition are quite strong. However, as it will be shown in Section 5.2, the PRS semantics of an ITRS is always an extension of the MRSs semantics of the same ITRSs. The following proposition summarises the ramifications of the axiomatisation we have just introduced:

Proposition 3.3.11 (PRSs extending MRSs are conservative extensions) *Let $\mathcal{P}$ be a PRS extending an MRS $\mathcal{M}$. Then the following holds for all reduction sequences S in $\mathcal{P}$:*

(i) $S\colon a \hookrightarrow_{\mathcal{P}} \dots$ is total iff $S\colon a \hookrightarrow_{\mathcal{M}} \dots$.

(ii) $S\colon a \hookrightarrow_{\mathcal{P}} b$ is total iff $S\colon a \hookrightarrow_{\mathcal{M}} b$.

(iii) $S\colon a \twoheadrightarrow_{\mathcal{P}} \dots$ is total iff $S\colon a \twoheadrightarrow_{\mathcal{M}} \dots$.

(iv) $S\colon a \twoheadrightarrow_{\mathcal{P}} b$ is total iff $S\colon a \twoheadrightarrow_{\mathcal{M}} b$.

If $\mathcal{P}$ only weakly extends $\mathcal{M}$, then clauses (i) and (ii) hold.

Proof. Let $\mathcal{M} = (A, \Phi, \mathsf{src}, \mathsf{tgt}, \mathbf{d}, \mathsf{hgt})$ and $\mathcal{P} = (B, \Phi', \mathsf{src}', \mathsf{tgt}', \leq, \mathsf{cxt})$ and $S = (a_\iota \to a_{\iota+1})_{\iota<\alpha}$. In the following, we frequently use (1) – (5) to refer to the respective items of Definition 3.3.10(ii):

(i) For the "only if" direction, assume that $S\colon a \hookrightarrow_{\mathcal{P}} \dots$ is total. By the maximality condition for $(A, \mathbf{d})$ and $(B, \leq)$, the fact that S is total, and items (3) and (4), it follows that S is a reduction sequence in $\mathcal{M}$ as well. By Proposition 3.2.12, $S|_{[0,\lambda)}$ weakly converges to a_λ in $\mathcal{P}$ for each limit ordinal λ, i.e. $\liminf_{\iota\to\lambda} a_\iota = a_\lambda$. By the limits condition for $(A, \mathbf{d})$ and $(B, \leq)$, this implies $\lim_{\iota\to\lambda} a_\iota = a_\lambda$ for each limit ordinal $\lambda < \alpha$. According to Lemma 2.2.13, this means that S is weakly continuous in $\mathcal{M}$

For the converse direction, assume that $S\colon a \hookrightarrow_{\mathcal{M}} \dots$. By Lemma 2.2.13, we then have $\lim_{\iota\to\lambda} a_\iota = a_\lambda$, and, by the limits condition for $(A, \mathbf{d})$ and $(B, \leq)$, we have $\liminf_{\iota\to\lambda} a_\iota = a_\lambda$ for each limit ordinal $\lambda < \alpha$. Hence, $S\colon a \hookrightarrow_{\mathcal{P}} \dots$.

(ii) This immediately follows from the previous clause (i) if α is not a limit ordinal. If α is a limit ordinal, this follows from (i) and the maximality and limits condition for $(A, \mathbf{d})$ and $(B, \leq)$.

(iii) We can reason as follows:

$$
\begin{aligned}
& S\colon a_0 \twoheadrightarrow_{\mathcal{P}} \dots \text{ is total} \\
\overset{\text{(a)}}{\iff}\quad & \text{all } a_\iota \text{ are maximal and } S\colon a_0 \twoheadrightarrow_{\mathcal{P}} \dots \\
\overset{\text{(b)}}{\iff}\quad & \text{all } a_\iota \text{ are maximal and } S|_{[0,\lambda)}\colon a_0 \twoheadrightarrow_{\mathcal{P}} a_\lambda \text{ for all limit ordinals } \lambda < \alpha \\
\overset{\text{(c)}}{\iff}\quad & \text{all } a_\iota \text{ are maximal and } S|_{[0,\lambda)}\colon a_0 \twoheadrightarrow_{\mathcal{M}} a_\lambda \text{ for all limit ordinals } \lambda < \alpha \\
\overset{\text{(d)}}{\iff}\quad & S\colon a_0 \twoheadrightarrow_{\mathcal{M}} \dots
\end{aligned}
$$

Equivalence (a) follows from the definition of total reduction sequences, equivalence (b) follows from Proposition 3.2.12, equivalence (c) follows from (5), and equivalence (d) follows from Proposition 3.1.17 and items (3) and (4).

(iv) This immediately follows from (iii) if α is not a limit ordinal. If α is a limit ordinal, then this follows immediately from (5).

The proof of clauses (i) and (ii) is independent of condition (5). Hence, they are also valid for weak extensions. □

3.4 Alternative Models of Transfinite Reductions

We have limited our discussion of transfinite reductions chiefly to metric spaces and partially ordered sets. As indicated in the introduction to this chapter, there are also other models that are worthwhile considering.

The most important model is of course that of general topological spaces. Such a "bare" topological model was employed by Rodenburg [Rod98]. The main reason for this choice is, however, that a signature with arbitrary ordinal arities is studied. In this setting, the metric on terms that we use here is not suitable as the following example illustrates:

Example 3.4.1
Let f be a symbol of arity ω, and a and b be nullary. Consider the rewrite rule $a \to b$. and the following reduction sequence:

$$f(a,a,a,a,a\dots) \to f(b,a,a,a,a\dots) \to f(b,b,a,a,a,\dots) \to f(b,b,b,a,a,\dots)$$

Intuitively, this sequence should weakly and strongly converge to the term $f(b,b,b,\dots)$, as the reduction steps occur more and more far to the right. But, for the metric, the depth of the changes and the depth of the reduction steps is significant. Yet, both the depth of the changes in the terms and the depth of the reduction steps remain constant.

Therefore, [Rod98] uses a topology on terms in which the reduction illustrated in the above example indeed converges to $f(b,b,b,\dots)$. One can easily see, though, that the same result can be obtained with the PRS model: We did not consider signatures with transfinite arities. The partial order, however, can be extended to partial terms with transfinite arities in the obvious way. It is clear that the abovementioned reduction sequence weakly and strongly converges to $f(b,b,b,\dots)$ within the PRS semantics as well.

Another variant of transfinite reductions that one might consider is the one introduced by Corradini [Cor93]. There, also the partial order on terms is employed. But instead of defining transfinite reduction sequences, a theory of parallel reductions is established. Corradini shows that the parallel contraction of a set of mutually independent redexes of left-linear rules is well defined. Because an infinite term might have infinitely many redexes, such a set of mutually independent redexes can be infinite and, hence, a parallel contraction of these redexes corresponds to some sort of transfinite reduction sequence. However, for this result, the applicability and the application of a rewrite rule is defined quite differently by allowing a partial matching of left-hand sides. As we will see in Section 5.5.1, also the

PRS semantics of ITRSs allows the definition of a meaningful notion of parallel reductions – at least for orthogonal systems.

This thesis does not consider higher-order term rewriting or λ-calculus. As argued in [KKSdV97], when dealing with higher-order systems, it is desirable to consider alternative metrics which amounts to employ a context-sensitive definition of the notion of depth in a higher-order term. This is of course not prohibited by the definition of MRSs. Nevertheless, it is not entirely trivial how a corresponding PRS semantics should look like. Or to put it in terms of Section 3.3.2: How does a partial order have to be defined in order to extend such an alternative metric on terms? In his analysis of strongly convergent reductions of λ-calculus Blom [Blo04] solved this problem by using different definitions of the context function cxt depending on the metric under consideration. Yet, this approach, specifically designed for the needs of the λ-calculus, is unsatisfactory as it does not affect weakly converging reductions accordingly. In order to achieve this, changing the partial order is inevitable.

Chapter 4

Term Graphs

In this chapter we want to introduce term graphs as a generalisation of the concept of terms. Essentially, terms are ordered trees with labelled nodes. Term graphs generalise this concept by omitting the requirement of having a tree structure and, thus, allow a general graph structure instead. This idea is useful for implementing systems dealing with terms, e.g. functional programming languages or term rewriting systems. The advantage that term graphs provide is that they allow to represent terms in a compact way.

The benefits of using term graphs for this purpose is twofold: Firstly, it allows to enhance efficiency by representing several occurrences of the same subterm as a single subgraph. Secondly, and most importantly for our purposes, a certain class of infinite terms, so-called *rational terms*, can be represented by finite term graphs.

The former, also known as *sharing of subexpressions*, was first employed in the setting of reduction systems by Wadsworth [Wad71] focusing on λ-calculus and later by Staples [Sta80a, Sta80b, Sta80c] in the setting of term rewriting. Moreover, this idea was also used in implementations of functional programming languages, e.g. in [Tur79] and [PJ87], and it was even used as the basis for the formalisation of the semantics of functional programming languages (cf. [vESP97]) such as *Clean* [Pla95].

The latter application of term graphs, viz. cyclic term graphs, was initially used quite informally again by implementers of functional languages [Tur79, PJ87]. The idea behind cyclic term graphs is to represent infinitely repeating structures in a term by cycles in a term graph. The formal justification of these techniques was only given later by Farmer and Watro [FRW90, FW90]. Subsequently, shortly after the advent of infinitary rewriting, Kennaway et al. [KKSdV94] studied the connection of cyclic term graph rewriting and infinitary term rewriting. They were able to show that term graph rewriting can be used to finitely represent limited forms of infinitary term rewriting. This particular result provides the motivation for considering term graphs in this thesis. We shall come back to this topic in Chapter 6 in which we discuss term graph rewriting and its connection to (infinitary) term rewriting in more detail. Moreover, our aim is it to extend infinitary rewriting techniques known from term rewriting to the setting of term graphs.

Eventually, the goal is to be able to define transfinite reductions for term graph rewriting systems which are studied in Chapter 6. In Chapter 3, we have seen two approaches to define transfinite reductions: One using a metric space and another using a partially ordered set. We want to explore both alternatives. To this end, we will introduce a partial order and a metric on term graphs in Section 4.3 and Section 4.5, respectively. Both the metric and the partial order are designed such that they extend the corresponding concepts on terms. The main results of this chapter are that, just as in the setting of terms, the partial order forms a complete semilattice and the metric is a complete ultrametric. Additionally, it is shown that the partial order extends the metric in the sense of Definition 3.3.10.

Before we can start defining and investigating a partial order and a metric on term graphs we need to prepare the necessary notions and tools. The fundamental concepts of

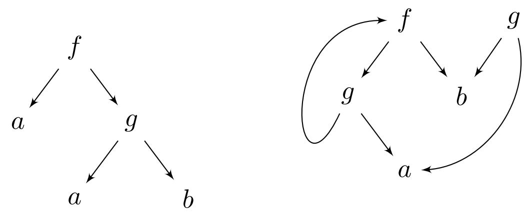

(a) $f(a, g(a, b))$. (b) A graph.

Figure 4.1: Example for a tree representation of a term; generalisation to graphs.

graphs and term graphs are given in Section 4.1. Subsequently, homomorphisms, by far the most important concept for our endeavour, are covered in Section 4.2. Homomorphisms give rise to the important equivalence relation of isomorphism. Technically, the partial order and the metric can only be meaningfully defined on the quotient of the term graphs by the isomorphism equivalence. This quotient construction is investigated in Section 4.3.

4.1 Graphs and Term Graphs

This section provides the basic notions for term graphs and more generally for graphs. Terms over a signature, say Σ, can be thought of as rooted trees whose nodes are labelled with symbols from Σ. Moreover, in these trees a node labelled with a k-ary symbol is restricted to have out-degree k and the outgoing edges are ordered. In this way the i-th successor of a node labelled with a symbol f is interpreted as the root node of the subtree that represents the i-th argument of f. For example, consider the term $f(a, g(a, b))$. The corresponding representation as a tree is shown in Figure 4.1a.

When turning to graphs, we simply remove the restriction of considering only trees but instead consider directed graphs. An example for a graph is depicted in Figure 4.1b. Note that this allows nodes to have more than one predecessor. Particularly, graphs also allow cycles as it can be seen in the example.

Definition 4.1.1 (graph)
Let Σ be a signature. A Σ-*graph* (or simply *graph*) is a tuple $g = (N, \mathsf{lab}, \mathsf{suc})$ consisting of

- a set N (of *nodes*),
- a *labelling function* $\mathsf{lab}\colon N \to \Sigma$, and
- a *successor function* $\mathsf{suc}\colon N \to N^*$.

The functions lab and suc are restricted to comply with the condition that $|\mathsf{suc}(n)| = \mathsf{ar}(\mathsf{lab}(n))$, i.e. a node labelled with a k-ary symbol has precisely k successors. Moreover, the graph g is called *finite* if the set N is finite.

Notation 4.1.2. Let $g = (N, \mathsf{lab}, \mathsf{suc})$ be a Σ-graph. If $\mathsf{suc}(n) = n_0 \cdot \ldots \cdot n_{k-1}$, then $\mathsf{suc}_i(n)$ denotes n_i, the i-th successor of n, for all $0 \le i < k$ and $\mathsf{ar}_g(n)$ denotes k, the arity of the label of n. Often we also use suc_i to define the function suc. More precisely, we then define a partial function $\mathsf{suc}_i\colon N \rightharpoonup N$ for each $i \in \mathbb{N}$ such that $\mathsf{suc}_i(n)$ is defined iff $i < \mathsf{ar}_g(n)$. The induced function suc is then defined as expected as

$$\mathsf{suc}(n) = \mathsf{suc}_0(n) \cdot \ldots \cdot \mathsf{suc}_{k-1}(n) \quad \text{for all } n \in N \text{ and } k = \mathsf{ar}_g(n).$$

In order to reduce notation, we sometimes refer to nodes labelled with a symbol, say σ, as σ-*nodes*, and nodes not labelled with σ as *non-σ-nodes*. Analogous notions are also used for sets of symbols Δ. That is, a node labelled with a symbol in Δ is referred to as Δ-*node* and a node labelled with a symbol not contained in Δ is called a *non-Δ-node*.

Example 4.1.3
Let $\Sigma = \{f/2, h/2, c/0\}$ be a signature and $g = (N, \mathsf{lab}, \mathsf{suc})$, where $N = \{n_1, n_2, n_3, n_4\}$, $\mathsf{lab}(n_1) = \mathsf{lab}(n_3) = f$, $\mathsf{lab}(n_2) = h$, $\mathsf{lab}(n_4) = c$ and $\mathsf{suc}(n_1) = n_2 \cdot n_3$, $\mathsf{suc}(n_2) = n_1 \cdot n_4$, $\mathsf{suc}(n_3) = n_2 \cdot n_4$, $\mathsf{suc}(n_4) = \varepsilon$. As suc respects the arities of the node labels, g is a Σ-graph. We often choose to depict a graph – well – graphically. For the example graph g, such a graphical representation would look like

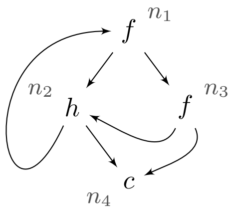

or simply

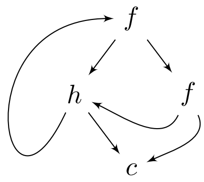

Usually, we are not interested in the names of the nodes, in which case we prefer the second representation which does not mention them. The edges in the graphical representation always emanate from the lower half of a node. The leftmost outgoing edge of a node n always points to $\mathsf{suc}_0(n)$, the edge to the right of it to $\mathsf{suc}_1(n)$ etc.

Since the outgoing edges of a node are ordered, a path in a graph can be described by giving the starting nodes and a sequence of numbers each of which defines which outgoing edge of the "current node" to take. This gives rise to the notion of a path. It is inspired by the notion of *access paths* introduced in [AK96].

Definition 4.1.4 (path, cyclicity)
Let $g = (N, \mathsf{lab}, \mathsf{suc})$ be a Σ-graph and $n, n' \in N$.

(i) A *path* in g from n to n' is a finite sequence $(p_i)_{i<l}$ in $\mathbb{N}$ such that either

- $n = n'$ and $(p_i)_{i<l}$ is empty, i.e. $l = 0$, or
- $0 \le p_0 < \mathsf{ar}_g(n)$ and the suffix $(p_i)_{1 \le i < l}$ is a path in g from $\mathsf{suc}_{p_0}(n)$ to n'.

(ii) If there exists a path from n to n' in g, we say that n' is *reachable* from n in g.

(iii) A path π in g from n to n' is called *cyclic* if there are two prefixes π_1, π_2 of π with $\pi_1 < \pi_2$ such that both π_1 and π_2 are paths from n to some common node n''. A path that is not cyclic is called *acyclic*.

(iv) g is called *cyclic* if it contains a cyclic path. Otherwise it is called *acyclic*.

Example 4.1.5
Consider the graph g from Example 4.1.3. The sequence $0 \cdot 1 \cdot 1$ is a path in g from n_2 to n_4. It is acyclic. The sequence $1 \cdot 0 \cdot 0 \cdot 1 \cdot 1$ is a path from n_1 to n_4. It is cyclic since it has the prefixes $\pi_1 = 1$ and $\pi_2 = 1 \cdot 0 \cdot 0 \cdot 1$ with $\pi_1 < \pi_2$ which are both paths from n_1 to n_3.

Now we can turn to term graphs which are simply graphs having a distinguished root node. Additionally, we require that every node in a term graph has to be reachable from the root node.

Definition 4.1.6 (term graph)
Let Σ be a signature. A *term graph* over Σ is a tuple $g = (N, \mathsf{lab}, \mathsf{suc}, r)$, where $h = (N, \mathsf{lab}, \mathsf{suc})$ is a Σ-graph, $r \in N$, and all nodes in N are reachable from r in the *underlying graph* h. r is called the *root node* of g. The term graph g is called *finite* if the underlying graph h is. The set of all term graphs over Σ is denoted as $\mathcal{G}^\infty(\Sigma)$, the set of all finite term graphs over Σ is denoted as $\mathcal{G}(\Sigma)$.

Notation 4.1.7. Let $g = (N, \mathsf{lab}, \mathsf{suc})$ be a graph and $r \in N$. For the sake of brevity, we will use (g, r) to denote the term graph $h = (N, \mathsf{lab}, \mathsf{suc}, r)$. Moreover, we use the notation N^h, lab^h, suc^h and r^h to refer to the respective components N,lab, suc and r of h.

Just as for graphs, we sometimes want to use a graphical notation for specifying term graphs. For term graphs, we use the same conventions that we have introduced for graphs. However, we additionally need a means to distinguish the root node of the term graph. To this end, we stipulate that the topmost node is considered to be the root node of the term graph. So the graphical representation given in Example 4.1.3 depicts the term graph (g, r), where r, the root node of the term graph, is the node n_1.

Definition 4.1.8 (occurrence, depth, tree)
Let $g = (N, \mathsf{lab}, \mathsf{suc}, r) \in \mathcal{G}^\infty(\Sigma)$ and $n, n' \in N$.

(i) A *path* in g from n to n' is a path in $(N, \mathsf{lab}, \mathsf{suc})$ from n to n'. An *occurrence* of n is a path in g from r to n. The set of all occurrences in g is denoted by $\mathcal{P}(g)$; the set of all occurrences of a node n in g is denoted $\mathcal{P}_g(n)$.

(ii) For each node $n \in N$, the *depth* of n, denoted $\mathsf{depth}_g(n)$, is the minimum of the lengths of the occurrences of n in g. The depth of g, denoted $\mathsf{depth}(g)$, is the maximum of the depths of the nodes in g if it exists and otherwise ∞. In sum,

$$\mathsf{depth}_g(n) = \min\{|\pi| \mid \pi \in \mathcal{P}_g(n)\}, \quad \mathsf{depth}(g) = \max\left\{\mathsf{depth}_g(n) \,\middle|\, n \in N\right\} \cup \{\infty\}$$

(iii) Let $\Delta \subseteq \Sigma$. The *depth of* Δ in g, denoted $\Delta\text{-}\mathsf{depth}(g)$, is the minimal depth of a Δ-node, i.e., a node labelled with a symbol in Δ, or ∞ if no such node exists in g. More precisely,

$$\Delta\text{-}\mathsf{depth}(g) = \min\left\{\mathsf{depth}_g(n) \,\middle|\, n \in N, g(n) \in \Delta\right\} \cup \{\infty\}$$

If Δ is a singleton set $\{\sigma\}$, we also write $\sigma\text{-}\mathsf{depth}(g)$ instead of $\{\sigma\}\text{-}\mathsf{depth}(g)$.

(iv) g is called *cyclic* if it contains a cyclic occurrence. Otherwise it is called *acyclic*.

(v) g is called a *term tree* if each node in g has exactly one occurrence.

Remark 4.1.9. Note that we use the same notation $\mathcal{P}(\cdot)$ for the set of node occurrences in a term graph as for positions in a term. The reason for doing so is that these two concepts are tightly related. In fact, they coincide on term trees. One can easily see that term trees essentially correspond to terms. The positions of a term are exactly the unique node occurrences in a corresponding term tree. Despite the fact that occurrences are a generalisation of the concept of positions we prefer the name "occurrence" as it conveys the idea that there might be multiple ones for a single node whereas positions in terms are unique. We should note here that there is one technical difference between terms and term tress: In term trees, nodes have names and, hence, can be distinguished from one another. That is, there are term trees which are “structurally equivalent” and, therefore, correspond to the same term, but which are, nonetheless, different from each other due to their nodes. This technical difference is overcome, however, when we consider isomorphism classes of term graphs.

Notation 4.1.10. Let $g = (N, \mathsf{lab}, \mathsf{suc})$ be a Σ-graph and $n \in N$. The *subgraph* of g restricted to n, denoted $g|n$, is the Σ-graph $(N', \mathsf{lab}', \mathsf{suc}')$, where N' is the set of nodes in g that are reachable from n in g, $\mathsf{lab}' = \mathsf{lab}|_{N'}$, and $\mathsf{suc}' = \mathsf{suc}|_{N'}$. For a term graph $h = (g, r)$, we also use this construction. That is, we write $h|n$ for the *sub-term graph* $(g|n, n)$. Additionally, for each $\pi \in \mathcal{P}(h)$, we use the notation $\mathsf{node}_h(\pi)$ to denote the unique node n that has an occurrence π in g and the notation $h(\pi)$ for $\mathsf{lab}^h(\mathsf{node}_h(\pi))$, i.e. the labelling of the node at π. Furthermore, by abuse of notation, we write $\mathsf{ar}_h(\pi)$ instead of $\mathsf{ar}_h(\mathsf{node}_h(\pi))$, and $h|\pi$ instead of $h|\mathsf{node}_h(\pi)$.

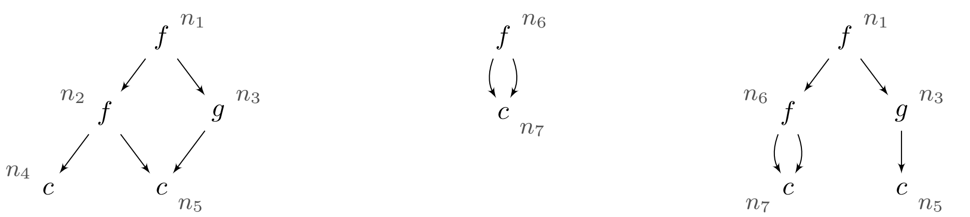

(a) Term graph g. (b) Term graph h. (c) Term graph $g[h]_{n_2}$.

Figure 4.2: Example for a replacement of a node by a term graph.

Definition 4.1.11 (replacement)
Let $g, h \in \mathcal{G}^\infty(\Sigma)$ and $n \in N^g$. Let M be the set of nodes only reachable through n in g (including n itself). That is,

$$M = \{m \in N^g \mid \forall \pi \in \mathcal{P}_g(m) \exists \pi' \in \mathcal{P}_g(n).\ \pi' \leq \pi\}.$$

The *replacement* of n in g by h, denoted $g[h]_n$, is the term graph $(N, \mathsf{lab}, \mathsf{suc}, r)$ with

$$N = (N^g \setminus M) \uplus N^h \qquad r = \begin{cases} r^g & \text{if } n \neq r^g \\ r^h & \text{if } n = r^g \end{cases}$$

$$\mathsf{lab}(m) = \begin{cases} \mathsf{lab}^g(m) & \text{if } m \in N^g \\ \mathsf{lab}^h(m) & \text{if } m \in N^h \end{cases}$$

$$\mathsf{suc}_i(m) = \begin{cases} \mathsf{suc}_i^g(m) & \text{if } m \in N^g, \mathsf{suc}_i^g(m) \neq n \\ r^h & \text{if } m \in N^g, \mathsf{suc}_i^g(m) = n \\ \mathsf{suc}_i^h(m) & \text{if } m \in N^h \end{cases}$$

Example 4.1.12
Consider the term graphs g and h illustrated in Figure 4.2a and Figure 4.2b, respectively. For constructing the replacement of n_2 in g by h, we have to identify the nodes in M, i.e. those nodes in g that are only reachable through n_2. Clearly, n_2 itself is in M. Additionally, also n_4 is in M since its only occurrence $0 \cdot 0$ is an extension of the n_2 occurrence 0. Therefore, these two nodes and all edges emanating from them are removed, and are replaced by the two nodes of h and their edges. The single edge that is pointing to n_2 in g is redirected to n_6, the root node of h. Since only those nodes are removed that are solely reachable through n_2, there are no other edges going from the remaining nodes to the now removed nodes. The resulting term graph $g[h]_n$ is depicted in Figure 4.2c.

For most properties of term graphs (just as for terms), the restriction to signatures with symbols of finite arity is not essential. Yet, for the following lemma that we will need later on, it is:

Lemma 4.1.13 (finitely many occurrences of bounded length)
Let $g \in \mathcal{G}^\infty(\Sigma)$ and $d \in \mathbb{N}$. Then there are only finitely many occurrences in g of depth at most d, i.e. the set $\{\pi \in \mathcal{P}(g) \mid |\pi| \leq d\}$ is finite.

Proof. Straightforward induction on d. □

4.2 Homomorphisms

This section is concerned with homomorphisms between graphs and most importantly between term graphs. Homomorphisms constitute the most important notion on term graphs.

They provide the foundation for the partial order and the metric on term graphs whose definitions are given in subsequent sections of this chapter. Furthermore, the definition of rewriting on term graphs that is investigated in Chapter 6 heavily depends on homomorphisms, too. This section also provides the fundamental properties of homomorphism which will become indispensable when studying the partial order, the metric and the rewrite relation on term graphs.

For term graph rewriting, there is a plethora of different approaches studied in the literature. In principle, we will follow the approach of Barendregt et al. [BvEG+87]. This approach uses graphs that might contain "empty" nodes, i.e. nodes without a label and without successors. These empty nodes represent variables of term graph rewrite rules. This is formalised by the definition of homomorphisms, which are not required to be homomorphic in empty nodes. Instead of considering graphs with empty nodes, we single out a set Δ of nullary symbols and regard nodes labelled with one of these symbols as "empty". As we have mentioned, homomorphisms in the approach of [BvEG+87] only have to be homomorphic in "non-empty" nodes. This motivates the notion of Δ-homomorphisms, where Δ is a set of nullary symbols. Δ-homomorphisms only have to be homomorphic in non-Δ-nodes. This provides much more flexibility which will be necessary as we will see. Δ-homomorphisms will be used to define the matching of term graph rewrite rules as well as a partial order on term graphs.

Definition 4.2.1 (Δ-homomorphism, Δ-isomorphism)
Let Σ be a signature, $\Delta \subseteq \Sigma^{(0)}$, and g, h Σ-graphs.

(i) Let $n \in N^g$ and $\varphi\colon N^g \to N^h$ a function. φ is called *homomorphic* in n if it satisfies the following two conditions:

$$\mathsf{lab}^g(n) = \mathsf{lab}^h(\varphi(n)) \qquad \text{(labelling)}$$
$$\varphi(\mathsf{suc}^g_i(n)) = \mathsf{suc}^h_i(\varphi(n)) \quad \text{for all } 0 \le i < \mathsf{ar}_g(n) \qquad \text{(successor)}$$

For a subset $N' \subseteq N$, we also say that φ is homomorphic in N' if φ is homomorphic in n for all $n \in N'$.

(ii) A Δ-*homomorphism* φ from g to h, denoted $\varphi\colon g \to_\Delta h$, is a function $\varphi\colon N^g \to N^h$ such that φ is homomorphic in n for all $n \in N^g$ with $\mathsf{lab}^g(n) \notin \Delta$.

(iii) A Δ-*homomorphism* φ from a term graph (g, r^g) to a term graph (h, r^h), denoted $\varphi\colon (g, r^g) \to_\Delta (h, r^h)$, is a Δ-homomorphism $\varphi\colon g \to_\Delta h$ with

$$\varphi(r^g) = r^h \qquad \text{(root)}$$

(iv) A Δ-homomorphism $\varphi\colon a \to_\Delta b$ (between two graphs or two term graphs) is called a Δ-*isomorphism*, written $\varphi\colon a \mathrel{\overset{\sim}{\to}}_\Delta b$, if there is a Δ-homomorphism $\varphi^{-1}\colon b \to_\Delta a$ that is the inverse of φ, i.e. $\varphi \circ \varphi^{-1}$ and $\varphi^{-1} \circ \varphi$ are identity functions. In this case, we also write $a \cong_\Delta b$ and say that a and b are Δ-*isomorphic*.

Notation 4.2.2. For a Δ-homomorphism φ between two objects a, b (both graphs or term graphs), we use the notation $\varphi\colon a \to_\sigma b$ and $\varphi\colon a \to b$ if $\Delta = \{\sigma\}$ or $\Delta = \emptyset$, respectively. Accordingly, we use the notions σ-homomorphism and homomorphism. The same convention applies to Δ-isomorphisms and Δ-isomorphic (term) graphs.

The first observation that can be made about Δ-homomorphisms is that they have a categorical structure:

Proposition 4.2.3 (categories of (term) graphs)
$\mathcal{G}^\infty(\Sigma)$ together with the Δ-homomorphisms on $\mathcal{G}^\infty(\Sigma)$ forms a category. Similarly, the set of Σ-graph and their Δ-homomorphisms form a category, too.

Proof. The identity Δ-homomorphism for a term graph (resp. graph) is the identity mapping on the nodes. Identity Δ-homomorphisms are easily seen to be the unit element w.r.t. function composition. Moreover, an easy equational reasoning reveals that the composition of two Δ-homomorphisms is again a Δ-homomorphism. Associativity of this composition is obvious as Δ-homomorphisms are functions. □

One can quite easily see that homomorphisms between term graphs are always surjective.

Lemma 4.2.4 (homomorphisms are surjective)
Let $g, h \in \mathcal{G}^\infty(\Sigma)$ and $\varphi\colon g \to h$. Then φ is surjective.

Proof. Follows from an easy induction on the depth of the nodes in h. □

Clearly, for the above lemma to hold, homomorphicness in each argument and the root condition are indispensable. Therefore, surjectivity is not necessary for homomorphisms between graphs or for general Δ-homomorphisms.

However, it turns out that, for each pair of term graphs g, h, there is at most one Δ-homomorphism $\varphi\colon g \to_\Delta h$.

Lemma 4.2.5 (at most one Δ-homomorphism)
Let $g, h \in \mathcal{G}^\infty(\Sigma)$. Then there is at most one Δ-homomorphism from g to h.

Proof. Suppose, there are two Δ-homomorphisms $\varphi_1, \varphi_2\colon g \to_\Delta h$. We prove that $\varphi_1 = \varphi_2$ by showing that $\varphi_1(n) = \varphi_2(n)$ for all $n \in N^g$ by induction on the depth of n.

Let $\mathsf{depth}_g(n) = 0$, i.e. $n = r^g$. By the root condition, we have that $\varphi_1(r^g) = r^h = \varphi_2(r^g)$. Let $\mathsf{depth}_g(n) = d > 0$. Then n has an occurrence $\pi \cdot i$ in g such that $\mathsf{depth}_g(n') < d$ for $n' = \mathsf{node}_g(\pi)$. Hence, we can employ the induction hypothesis for n' to obtain the following:

$$\begin{aligned}
\varphi_1(n) &= \mathsf{suc}_i^h(\varphi_1(n')) && \text{(successor condition for } \varphi_1\text{)}\\
&= \mathsf{suc}_i^h(\varphi_2(n')) && \text{(ind. hyp.)}\\
&= \varphi_2(n) && \text{(successor condition for } \varphi_2\text{)}
\end{aligned}$$

□

For the definition of the partial order and the metric on term graphs, we only need Δ-homomorphisms and Δ-isomorphisms where Δ is the empty set or a singleton set. For these two cases, the respective notions of Δ-isomorphism coincide.

Lemma 4.2.6 (isomorphism equals σ-isomorphism)
Let Σ be a signature, $\sigma \in \Sigma^{(0)}$, and g, h two Σ-graphs (or two term graphs over Σ). Then

$$g \cong h \quad \textit{iff} \quad g \cong_\sigma h.$$

Proof. The proof is the same for term graphs and for graphs: The "only if" direction is trivial. For the converse direction, assume that there is a σ-isomorphism $\varphi\colon g \cong_\sigma h$. By definition, its inverse $\varphi^{-1}\colon h \cong_\sigma g$ is also a σ-isomorphism. We have to show that φ and φ^{-1} are also homomorphic in σ-nodes. Since σ is nullary, the successor condition is vacuously satisfied for σ-nodes. Let $n \in N^g$ with $\mathsf{lab}^g(n) = \sigma$ and suppose that the labelling condition is not satisfied, viz. $\mathsf{lab}^h(\varphi(n)) \neq \sigma$. Hence, φ^{-1} is homomorphic in $\varphi(n)$ and we can derive a contradiction by the following equations:

$$\mathsf{lab}^g(n) = \mathsf{lab}^g(\varphi^{-1}(\varphi(n))) = \mathsf{lab}^h(\varphi(n)) \neq \sigma$$

By symmetry, we obtain that φ^{-1} is a homomorphism, too. Therefore, $\varphi\colon g \cong h$. □

Note that a bijective Δ-homomorphism is not necessarily a Δ-isomorphism. To realise this, consider two term graphs g, h, each with one node only. Let the node in g be labelled with a and the node in h with b then the only possible a-homomorphism from g to h is clearly a bijection but not an a-isomorphism. On the other hand, bijective homomorphisms are isomorphisms.

Lemma 4.2.7 (bijective homomorphisms are isomorphisms)
Let $g, h \in \mathcal{G}^\infty(\Sigma)$ and $\varphi\colon g \to h$. Then the following are equivalent

(a) φ is an isomorphism.

(b) φ is bijective.

(c) φ is injective.

Proof. The implication (a) $\Rightarrow$ (b) is trivial. The equivalence (b) $\Leftrightarrow$ (c) follows from Lemma 4.2.4. For the implication (b) $\Rightarrow$ (a), consider the inverse φ^{-1} of φ. We need to show that φ^{-1} is a homomorphism from h to g. The root condition follows immediately from the root condition for φ. Similarly, an easy equational reasoning reveals that the fact that φ is homomorphic in N^g implies that φ^{-1} is homomorphic in N^h □

From the proof we can see that the equivalence (a) $\Leftrightarrow$ (b) also holds true for graphs.

The following lemma provides another characterisation of Δ-isomorphism:

Lemma 4.2.8 (two Δ-homomorphisms = Δ-isomorphism)
Let $g, h \in \mathcal{G}^\infty(\Sigma)$. If there is a Δ-homomorphism from g to h and one from h to g, respectively, then $g \cong_\Delta h$.

Proof. Let $\varphi\colon g \to_\Delta h$ and $\psi\colon h \to_\Delta g$. According to Lemma 4.2.3, the composition of φ and ψ is a Δ-homomorphism $\psi \circ \varphi\colon g \to_\Delta g$. Since, by Lemma 4.2.5, this is the only Δ-homomorphism from g to g, it has to be the identity Δ-homomorphism for g. For the same reason, also $\varphi \circ \psi\colon h \to_\Delta h$ is the identity Δ-homomorphism for h. Hence, ψ is the inverse of φ which is, therefore, a Δ-isomorphism. □

Next we present three lemmas that state in which way Δ-homomorphisms preserve the depth of the nodes in the involved term graphs. These lemmas are of rather technical nature and are needed in Section 4.5 for analysing the connection between the partial order and the metric on term graphs.

Lemma 4.2.9 (weak depth preservation of Δ-homomorphisms)
Let $g, h \in \mathcal{G}^\infty(\Sigma)$ and $\varphi\colon g \to_\Delta h$. Then $\mathsf{depth}_g(n) \geq \mathsf{depth}_h(\varphi(n))$ for all $n \in N^g$.

Proof. We prove by induction on d that $\mathsf{depth}_h(\varphi(n)) \leq d$ for all $n \in N^g$ with $\mathsf{depth}_g(n) = d$. If $d = 0$, then $n = r^g$. By the root condition, we have $\varphi(r^g) = r^h$. Hence, $\mathsf{depth}_h(\varphi(r^g)) = 0$. If $d > 0$, then there is a node $m \in N^g$ with $\mathsf{depth}_g(m) = d - 1$ and $\mathsf{suc}^g_i(m) = n$ for some i. Applying the induction hypothesis yields $\mathsf{depth}_h(\varphi(m)) \leq d - 1$. From the successor condition, we can obtain $\varphi(n) = \mathsf{suc}^h_i(\varphi(m))$. Hence, $\mathsf{depth}_h(\varphi(n)) \leq \mathsf{depth}_h(\varphi(m)) + 1 \leq d$. □

Lemma 4.2.10 (reverse weak depth preservation of Δ-homomorphisms)
Let $g, h \in \mathcal{G}^\infty(\Sigma)$, $\varphi\colon g \to_\Delta h$, $d \in \mathbb{N}$ and Δ-$\mathsf{depth}(g) \geq d$. Then, for all $n \in N^h$ with $\mathsf{depth}_h(n) \leq d$, there is a node $m \in \varphi^{-1}(n)$ with $\mathsf{depth}_g(m) \leq \mathsf{depth}_h(n)$.

Proof. We prove the equivalent statement

$$\forall e \leq d\ \forall n \in N^h.\ (\mathsf{depth}_h(n) = e \implies \exists m \in N^g.(\mathsf{depth}_g(m) \leq e \wedge \varphi(m) = n))$$

by induction on e. If $e = 0$, then $n = r^h$. Take $m = r^g$. Then we have $\mathsf{depth}_g(m) = 0$ and, therefore, $\varphi(m) = n$ according to the root condition. If $e > 0$, then there is some $n' \in N^h$ with

$$\begin{tikzcd} g \arrow[r, "\Delta"'] \arrow[d, "\wr"] & h \\ g' \arrow[r, "\Delta"'] & h' \arrow[u, "\wr"'] \end{tikzcd}$$

Figure 4.3: Isomorphism between Δ-homomorphisms.

$\mathsf{suc}_i^h(n') = n$ and $\mathsf{depth}_h(n') = e-1$. Hence, we can employ the induction hypothesis to obtain some $m' \in N^g$ with $\mathsf{depth}_g(m') \le e-1$ and $\varphi(m') = n'$. Since $\Delta\text{-}\mathsf{depth}(g) \ge d \ge e > \mathsf{depth}_g(m')$, we have $\mathsf{lab}^g(m') \notin \Delta$. Hence, φ is homomorphic in m'. Let $m = \mathsf{suc}_i^g(m')$. We can reason as follows:

$$\varphi(m) = \varphi(\mathsf{suc}_i^g(m')) = \mathsf{suc}_i^h(\varphi(m')) = \mathsf{suc}_i^h(n') = n, \quad \text{and}$$
$$\mathsf{depth}_g(m) \le \mathsf{depth}_g(m') + 1 \le e$$

□

Lemma 4.2.11 (Δ-depth preservation)
Let $g, h \in \mathcal{G}^\infty(\Sigma)$ and $\varphi\colon g \to_\Delta h$, then $\Delta\text{-}\mathsf{depth}(g) \le \Delta\text{-}\mathsf{depth}(h)$.

Proof. Let $d = \Delta\text{-}\mathsf{depth}(g)$. If $d = \infty$, then $g \in \mathcal{G}^\infty(\Sigma \setminus \Delta)$. Hence, φ is a homomorphism which is, according to Lemma 4.2.4, surjective. Consequently, due to the labelling condition, $h \in \mathcal{G}^\infty(\Sigma \setminus \Delta)$, too, which implies that $\Delta\text{-}\mathsf{depth}(h) = \infty$. If $d = 0$, then $d \le \Delta\text{-}\mathsf{depth}(h)$ is trivially true. If $0 < d < \infty$, then, by Lemma 4.2.10, for each node n at depth $< d$ in h, there is a node m at depth $< d$ in g with $\varphi(m) = n$. Since then $\mathsf{lab}^g(m) \notin \Delta$, we also have $\mathsf{lab}^h(n) \notin \Delta$ by the labelling condition. Hence, $d \le \Delta\text{-}\mathsf{depth}(h)$. □

4.3 Canonical Term Graphs

We are not interested in distinguishing isomorphic term graphs. The fact that there are distinct term graphs that are isomorphic is rather an unintentional artifact of the definition of term graphs. Therefore, we want to consider the quotient $\mathcal{G}^\infty(\Sigma)/\cong$ by the isomorphism relation $\cong$. Accordingly, we also consider the quotient of the set of all Δ-homomorphisms on $\mathcal{G}^\infty(\Sigma)$ by $\cong$ where, for two Δ-homomorphisms $\varphi\colon g \to_\Delta h$, $\varphi'\colon g' \to_\Delta h'$, we have $\varphi \cong \varphi'$ iff there are isomorphisms $\psi_1\colon g \cong g'$ and $\psi_2\colon h' \cong h$ such that $\varphi = \psi_2 \circ \varphi' \circ \psi_1$; in other words, the diagram in Figure 4.3 commutes. We refer to the equivalence classes of this equivalence relation as Δ-homomorphisms (on $\mathcal{G}^\infty(\Sigma)/\cong$). Moreover, for a property P, we say that an equivalence class (of term graphs or homomorphisms) enjoys P iff all elements of the equivalence class enjoy P.

In order to work with this quotient construction properly, it is convenient to choose a canonical representation. That is, we need an appropriate subset C of $\mathcal{G}^\infty(\Sigma)$ such that each element $[g]_\cong \in \mathcal{G}^\infty(\Sigma)/\cong$ is represented by exactly one element $c \in C$. The idea to obtain such a set is to provide a mechanism to give nodes a unique name. This can be achieved by requiring that each node is the set of its occurrences (cf. [Plu99]):

Definition 4.3.1 (canonical term graph)
Let Σ be a signature. A term graph $g = (N, \mathsf{lab}, \mathsf{suc}, r)$ over Σ is called *canonical* if $n = \mathcal{P}_g(n)$ holds for each $n \in N$. That is, each node is the set of its occurrences in the term graph. The set of all canonical term graphs is denoted by $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$; the set of all finite canonical term graphs is denoted by $\mathcal{G}_\mathcal{C}(\Sigma)$.

Before we formally prove that canonical term graphs indeed canonically represent isomorphism classes of term graphs we want to introduce some alternative characterisations of Δ-homomorphisms which will simplify the reasoning over them.

Lemma 4.3.2 (characterisation of Δ-homomorphisms)
Let $g, h \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$ and $\varphi\colon N^g \to N^h$. Then $\varphi\colon g \to_\Delta h$ iff, for all $n \in N^g$, the following holds:

(a) $n \subseteq \varphi(n)$, *and*

(b) $\mathsf{lab}^g(n) = \mathsf{lab}^h(\varphi(n))$ *whenever* $\mathsf{lab}^g(n) \notin \Delta$.

Proof. For the "only if" direction, let $\varphi\colon g \to_\Delta h$. (b) is the labelling condition which has to be satisfied by φ. To establish (a), we show the equivalent statement

$$\forall \pi \in \mathcal{P}(g).\ \forall n \in N^g.\ \pi \in n \implies \pi \in \varphi(n)$$

We do so by induction on the length of π: If $\pi = \varepsilon$, then $\pi \in n$ implies $n = r^g$. By the root condition, we have $\varphi(r^g) = r^h$ and, therefore, $\pi = \varepsilon \in r^h$. If $\pi = \pi' \cdot i$, then let $n' = \mathsf{node}_g(\pi')$. Consequently, $\pi' \in n'$ and, by induction hypothesis, $\pi' \in \varphi(n')$. Since $\pi = \pi' \cdot i$, we have $\mathsf{suc}^g_i(n') = n$. By the successor condition we can conclude $\varphi(n) = \mathsf{suc}^h_i(\varphi(n'))$. This and $\pi' \in \varphi(n')$ yields that $\pi' \cdot i \in \varphi(n)$.

For the "if" direction, we assume (a) and (b). The labelling condition follows immediately from (b). For the root condition, observe that since $\varepsilon \in r^g$, we also have $\varepsilon \in \varphi(r^g)$. Hence, $\varphi(r^g) = r^h$. In order to show the successor condition, let $n, n' \in N^g$ and $0 \le i < \mathsf{ar}_g(n)$ such that $\mathsf{suc}^g_i(n) = n'$. Then there is an occurrence $\pi \in n$ with $\pi \cdot i \in n'$. By (a), we can conclude that $\pi \in \varphi(n)$ and $\pi \cdot i \in \varphi(n')$ which implies that $\mathsf{suc}^h_i(\varphi(n)) = \varphi(n')$. □

Example 4.3.3
In the above lemma, (a) states that the term graph h "has more sharing" than g. The simplest example is the following homomorphism φ from g to h:

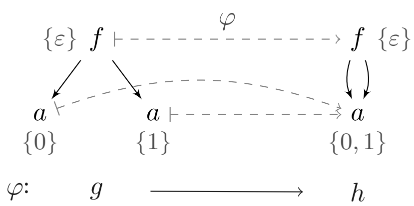

In this example, the nodes $\{0\}$ and $\{1\}$ are strictly less shared than their common image $\{0,1\}$.

Remark 4.3.4. Note that the lemma above also applies to non-canonical term graphs. It simply has to be rephrased such that instead referring to a node n its set of occurrences $\mathcal{P}_g(n)$ has to be referred to whenever the "inner structure" of n is used. We also use this the other way around: To avoid clutter in lengthy proofs, we often write n instead of $\mathcal{P}_g(n)$, i.e., we identify a node with its set of occurrences. We indicate this by writing that w.l.o.g. we assume the term graphs we deal with to be canonical.

From Lemma 4.2.5, we know that there is at most one Δ-homomorphism between two term graphs. The lemma above allows us to give this Δ-homomorphism right away if we know it exists. If there is a Δ-homomorphism from g to h, it is defined by $\varphi(n) = n'$, where n' is the unique node $n' \in N^h$ with $n \subseteq n'$.

Apart from that, it is clear that the set of nodes in a canonical partial term graph forms a partition of the set of occurrences. Hence, it defines an equivalence relation on the set of occurrences. For a canonical partial term graph g, we write $\sim_g$ for this equivalence relation. That is, $\pi_1 \sim_g \pi_2$ iff there is some $n \in N^g$ such that $\pi_1, \pi_2 \in n$. Using the convention mentioned in Remark 4.3.4 we can extend this to arbitrary term graphs. With this we are able to reformulate characterisation of Δ-homomorphisms:

Lemma 4.3.5 (characterisation of Δ-homomorphisms)
Let $g, h \in \mathcal{G}^\infty(\Sigma)$. Then there is a Δ-homomorphism from g to h iff, for all $\pi, \pi' \in \mathcal{P}(g)$, we have

(a) $\pi \sim_g \pi' \quad\Longrightarrow\quad \pi \sim_h \pi'$

(b) $g(\pi) = h(\pi) \quad$ *whenever* $\quad g(\pi) \notin \Delta$

Proof. W.l.o.g. we assume g and h to be canonical. For the "only if" direction, assume that φ is a Δ-homomorphism from g to h. Then we can use the properties (a) and (b) of Lemma 4.3.2, which we will refer to as (a') and (b') to avoid confusion. In order to show (a), assume $\pi \sim_g \pi'$. Then there is some node $n \in N^g$ with $\pi, \pi' \in n$. (a') yields $\pi, \pi' \in \varphi(n)$ and, therefore, $\pi \sim_g \pi'$. To show (b), we assume some $\pi \in \mathcal{P}(g)$ with $g(\pi) \notin \Delta$. Then we can reason as follows:

$$g(\pi) = \mathsf{lab}^g(\mathsf{node}_g(\pi)) \overset{\text{(b')}}{=} \mathsf{lab}^h(\varphi(\mathsf{node}_g(\pi))) \overset{\text{(a')}}{=} \mathsf{lab}^h(\mathsf{node}_h(\pi)) = h(\pi)$$

For the converse direction, assume that both (a) and (b) hold. Define the function $\varphi\colon N^g \to N^h$ by $\varphi(n) = n'$ iff $n \subseteq n'$ for all $n \in N^g$ and $n' \in N^h$. To see that this is well-defined, we show at first that, for each $n \in N^g$, there is at most one $n' \in N^h$ with $n \subseteq n'$. Suppose there is another node $n'' \in N^h$ with $n \subseteq n''$. Since $n \neq \emptyset$, this implies $n' \cap n'' \neq \emptyset$. Hence, $n' = n''$. Secondly, we show that there is at least one such node n'. Choose some $\pi^* \in n$. Since then $\pi^* \sim_g \pi^*$ and, by (a), also $\pi^* \sim_h \pi^*$ holds, there is some $n' \in N^h$ with $\pi^* \in n'$. For each $\pi \in n$, we have $\pi^* \sim_g \pi$ and, therefore, $\pi^* \sim_h \pi$ by (a). Hence, $\pi \in n'$. So we know that φ is well-defined. By construction, φ satisfies (a'). Moreover, because of (b), it is also easily seen to satisfy (b'). Hence, φ is a homomorphism from g to h. □

We can apply Lemma 4.2.8 to both of the above characterisations of Δ-homomorphisms to get another characterisation of Δ-isomorphisms:

Corollary 4.3.6 (characterisation of Δ-isomorphisms)
Let $g, h \in \mathcal{G}^\infty(\Sigma)$. Then the following holds:

(i) $\varphi\colon N^g \to N^h$ *is a Δ-isomorphism iff*

  (a) $\mathcal{P}_g(\varphi(n)) = \mathcal{P}_g(n)$, *and*

  (b) $\mathsf{lab}^g(n) = \mathsf{lab}^h(\varphi(n)) \quad$ *or* $\quad \mathsf{lab}^g(n), \mathsf{lab}^h(\varphi(n)) \in \Delta$.

(ii) $g \cong h$ *iff*

  (a) $\sim_g\; = \;\sim_h$, *and*

  (b) $g(\pi) = h(\pi) \quad$ *or* $\quad g(\pi), h(\pi) \in \Delta$.

Proof. Immediate consequence of Lemma 4.3.2 resp. Lemma 4.3.5 using Lemma 4.2.8. □

Now we can revisit the notion of canonical term graphs using the above characterisation of Δ-isomorphisms. We will define a function $\mathcal{C}(\cdot)\colon \mathcal{G}^\infty(\Sigma) \to \mathcal{G}^\infty_{\mathcal{C}}(\Sigma)$ that maps a term graph to its canonical representation. To this end, let $g = (N, \mathsf{lab}, \mathsf{suc}, r)$ be a term graph. Then define $\mathcal{C}(g) = (N', \mathsf{lab}', \mathsf{suc}', r')$ as follows.

$$N' = \{\mathcal{P}_g(n) \mid n \in N\} \qquad r' = \mathcal{P}_g(r)$$
$$\mathsf{lab}'(\mathcal{P}_g(n)) = \mathsf{lab}(n) \qquad \mathsf{suc}'_i(\mathcal{P}_g(n)) = \mathcal{P}_g(\mathsf{suc}_i(n)) \qquad \text{for all } n \in N, 0 \le i < \mathsf{ar}_g(n)$$

$\mathcal{C}(g)$ is easily seen to be a well-defined canonical term graph. With this definition we indeed capture the idea of a canonical representation of isomorphism classes as the following proposition confirms:

Proposition 4.3.7 (canonical partial term graphs are a canonical representation)
Let $g \in \mathcal{G}^\infty(\Sigma)$. $\mathcal{C}(g)$ canonically represents the equivalence class $[g]_\cong$. More precisely, it holds that

(i) $[g]_\cong = [\mathcal{C}(g)]_\cong$, *and*

(ii) $[g]_\cong = [h]_\cong \quad$ *iff* $\quad \mathcal{C}(g) = \mathcal{C}(h)$.

In particular, we have, for all canonical term graphs g, h, that $g = h$ iff $g \cong h$.

Proof. Straightforward consequence of Corollary 4.3.6. □

This means there is a one-to-one correspondence between canonical term graphs and their Δ-homomorphisms on the one hand, and equivalence classes of term graphs and their Δ-homomorphisms on the other hand. More precisely, the respective categories are isomorphic:

Proposition 4.3.8 (categories of term graphs)

(i) $\mathcal{G}^\infty(\Sigma)/\cong$ *together with the Δ-homomorphisms on* $\mathcal{G}^\infty(\Sigma)/\cong$ *forms a category.*

(ii) $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$ *together with the Δ-homomorphisms on* $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$ *forms a category.*

(iii) The categories described above are isomorphic.

Proof. (i) In order to obtain a category, we need to define compositions of Δ-homomorphisms. Δ-homomorphisms on isomorphism classes are sets of Δ-homomorphisms on term graphs. The composition is defined by forming the set of all compositions of Δ-homomorphisms in the corresponding sets. That is, for two Δ-homomorphisms $\Phi\colon [g_1]_\cong \to_\Delta [g_2]_\cong$ and $\Psi\colon [g_2]_\cong \to_\Delta [g_3]_\cong$, their composition $\Phi \circ \Psi\colon [g_1]_\cong \to_\Delta [g_3]_\cong$ is defined as $\Phi \circ \Psi = \{\varphi \circ \psi \mid \varphi \in \Phi, \psi \in \Psi\}$. One can easily check that this is indeed well-defined and gives rise to a category.

(ii) is obvious.

For (iii), we need a bijective functor $F\colon \mathcal{G}^\infty_\mathcal{C}(\Sigma) \to \mathcal{G}^\infty(\Sigma)/\cong$ between the categories. F maps each canonical term graph g to its isomorphism class $[g]_\cong$ and each Δ-homomorphism $\varphi\colon g \to_\Delta h$ on canonical term graphs to its isomorphism class $[\varphi]_\cong\colon [g]_\cong \to_\Delta [h]_\cong$. Functoriality of F is immediate and bijectivity follows from Proposition 4.3.7. □

Remark 4.3.9. In the following, we will make use of this fact. Instead of dealing with $\mathcal{G}^\infty(\Sigma)/\cong$ we prefer using the canonical representation $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$. In particular, we make use of this when defining term graph valued functions. We usually define such functions in terms of $\mathcal{G}^\infty(\Sigma)$. Such a function $f\colon A \to \mathcal{G}^\infty(\Sigma)$ can be uniquely modified to a function $f'\colon A \to \mathcal{G}^\infty_\mathcal{C}(\Sigma)$ by defining $f'(g) = \mathcal{C}(f(g))$. Usually, we also use the same notation, i.e., we identify f and f'. We use the same convention when defining operations on term graphs, i.e. functions $f\colon \mathcal{G}^\infty(\Sigma) \to \mathcal{G}^\infty(\Sigma)$. Such an operation can then be modified in the same way to an operation $f'\colon \mathcal{G}^\infty_\mathcal{C}(\Sigma) \to \mathcal{G}^\infty_\mathcal{C}(\Sigma)$.

An example for this approach is the following unravelling operation:

Definition 4.3.10 (unravelling)
Let $g = (N, \mathsf{lab}, \mathsf{suc}, r) \in \mathcal{G}^\infty(\Sigma)$. The *unravelling* of g, denoted $\mathcal{U}(g)$, is the term graph $(\mathcal{P}(g), \mathsf{lab}', \mathsf{suc}', \varepsilon)$, where

$$\begin{aligned} \mathsf{lab}'(\pi) &= \mathsf{lab}(\mathsf{node}_g(\pi)) \\ \mathsf{suc}'_i(\pi) &= \pi \cdot i \quad \text{for all } 0 \le i < \mathsf{ar}_g(\pi) \end{aligned}$$

From the definition, it is clear that the unravelling of a term graph is a term tree as each node has precisely one occurrence (which is, by construction, the node itself).

Remark 4.3.11. By identifying isomorphic term graphs, we can make the correspondence between term trees and terms explicit: Let T be the set of term trees over Σ. Then there is a one-to-one correspondence between $T/\cong$ and $\mathcal{T}^\infty(\Sigma)$: For each term $t \in \mathcal{T}^\infty(\Sigma)$, we can define a unique term tree $g = (\mathcal{P}(t), \mathsf{lab}, \mathsf{suc}, \varepsilon)$, with

$$\begin{aligned} \mathsf{lab}(\pi) &= t(\pi) \\ \mathsf{suc}_i(\pi) &= \pi \cdot i \quad \text{for all } 0 \le i < \mathsf{ar}_g(\pi) \end{aligned}$$

Similarly, following Remark 2.3.17, each canonical representative of an equivalence class in $T/\cong$ uniquely determines a term. Therefore, justified by this correspondence, we will identify the sets $\mathcal{T}^\infty(\Sigma)$ and $T/\cong$ (and the set of canonical term trees).

Having this, we can consider the unraveling operation on canonical term graphs as a function $\mathcal{U}\colon \mathcal{G}^\infty_\mathcal{C}(\Sigma) \to \mathcal{T}^\infty(\Sigma)$. By means of this operation, term graphs can be used as a compact representation of terms. This is the central motivation for considering term graphs. That is particularly the case for this thesis as term graphs allow to finitely represent infinite terms.

It can be easily seen that the internal structure of the nodes of a canonical term graph suffices to determine the external graph structure up to the labelling. This observation is detailed in the following lemma:

Lemma 4.3.12 (internal structure of nodes defines external structure)
Let $g = (N, \mathsf{lab}, \mathsf{suc}, r)$ be a canonical term graph. Then the following holds:

$$\left.\begin{aligned} n = r \quad &\textit{iff} \quad \varepsilon \in n \\ \mathsf{suc}_i(n) = n' \quad &\textit{iff} \quad \exists \pi \in n.\ \pi \cdot i \in n' \end{aligned}\right. \qquad \textit{for all } n, n' \in N,\ 0 \le i < \mathsf{ar}_g(n).$$

Proof. Immediate consequence of the definition of reachability in term graphs. □

But note that not any arbitrary partition and not any arbitrary set of occurrences define a canonical term graph structure as outlined in the lemma above. The next lemma presents some of the properties of the set of occurrences and of the equivalence relation on it, which are in fact characteristic.

Lemma 4.3.13 (internal node structure of canonical term graphs)
Let $g \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$. Then, for all $\pi, \pi' \in \mathcal{P}(g)$ and $i \in \mathbb{N}$, the following holds:

$$\pi \cdot i \in \mathcal{P}(g) \implies \pi \in \mathcal{P}(g) \quad \textit{and} \quad i < \mathsf{ar}_g(\pi) \qquad \text{(reachablity)}$$

$$\pi \sim_g \pi' \implies \begin{cases} g(\pi) = g(\pi') & \textit{and} \\ \pi \cdot j \sim_g \pi' \cdot j & \textit{for all } j < \mathsf{ar}_g(\pi) \end{cases} \qquad \text{(congruence)}$$

Proof. Straightforward. □

This observation motivates the following definition of an alternative representation of canonical term graphs.

Definition 4.3.14 (occurrence representation)
Let Σ be a signature.

(i) An *occurrence representation* over Σ is a tuple $(P, l, \sim)$ consisting of

- a non-empty set $P \subseteq \mathbb{N}^*$,
- a function $l\colon P \to \Sigma$, and
- an equivalence relation $\sim$ on P

satisfying the following conditions for all $\pi, \pi' \in P$ and $i \in \mathbb{N}$:

$$\pi \cdot i \in P \quad\Longrightarrow\quad \pi \in P \quad\text{and}\quad i < \mathsf{ar}(l(\pi)) \qquad\text{(reachablity)}$$

$$\pi \sim \pi' \quad\Longrightarrow\quad \begin{cases} l(\pi) = l(\pi') & \text{and} \\ \pi \cdot j \sim \pi' \cdot j & \text{for all } \; j < \mathsf{ar}(l(\pi)) \end{cases} \qquad\text{(congruence)}$$

(ii) Let $o = (P, l, \sim)$ be an occurrence representation. A canonical term graph $g \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$ *matches* o if

$$\begin{aligned} \mathcal{P}(g) &= P, \\ g(\pi) &= l(\pi) \quad \text{for all } \pi \in P, \text{ and} \\ P/{\sim} &= N^g \end{aligned}$$

There is a one-to-one correspondence between occurrence representations and canonical term graphs. The following lemma describes this in more detail.

Lemma 4.3.15 (occurrence representations are unique)

(i) Each term graph $g \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$ matches exactly one occurrence representation over Σ.

(ii) For each occurrence representation o over Σ, there is exactly one canonical term graph $g \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$ that matches o.

Proof. (i) Let $g = (N, \mathsf{suc}, \mathsf{lab}, r)$. Define the function $l \colon \mathcal{P}(g) \to \Sigma$ by $l(\pi) = g(\pi)$. By Lemma 4.3.13, we have that $o = (\mathcal{P}(g), l, \sim_g)$ is an occurrence representation. By construction, g matches o. It is immediate from the definition that any two occurrence representation matched by a common term graph have to coincide,

(ii) Let $o = (P, l, \sim)$. Define a term graph $g = (N, \mathsf{lab}, \mathsf{suc}, r)$ by

$$\begin{aligned} N &= P/\sim \\ \mathsf{lab}(n) = f \quad &\text{iff} \quad \exists \pi \in n.\, l(\pi) = f \\ \mathsf{suc}_i(n) = n' \quad &\text{iff} \quad \exists \pi \in n.\, \pi \cdot i \in n' \\ r = n \quad &\text{iff} \quad \varepsilon \in n \end{aligned}$$

lab and suc are well-defined because of the congruence condition satisfied by o. Since P is non-empty and closed under prefixes, it contains ε. Hence, r is well-defined. Moreover, by the reachability condition, each node in N is reachable from the root node. Thus, g is a well-defined canonical term graph. By construction. g matches o. Whenever there are two canonical term graphs matching o, they are isomorphic due to Corollary 4.3.6 and, therefore, have to be identical. □

This alternative representation of canonical term graphs will turn out to be quite useful for constructing term graphs in order to show that the yet to be defined partial order on canonical term graphs is a complete semilattice.

4.4 A Partial Order on Term Graphs

In this section we want to define and study a partial order on term graphs. The goal is to obtain a generalisation of the partial order defined on terms (cf. Section 2.3.2) to the setting of term graphs. As already pointed out in the previous section, we do not want to distinguish isomorphic term graphs. In fact, this is crucial for the partial order. Therefore, the order will be defined on canonical term graphs. An additional important requirement for the partial order, at least for our purposes, is that it admits the limit inferior for any sequence of canonical term graphs. In order to achieve this, it is sufficient, according to

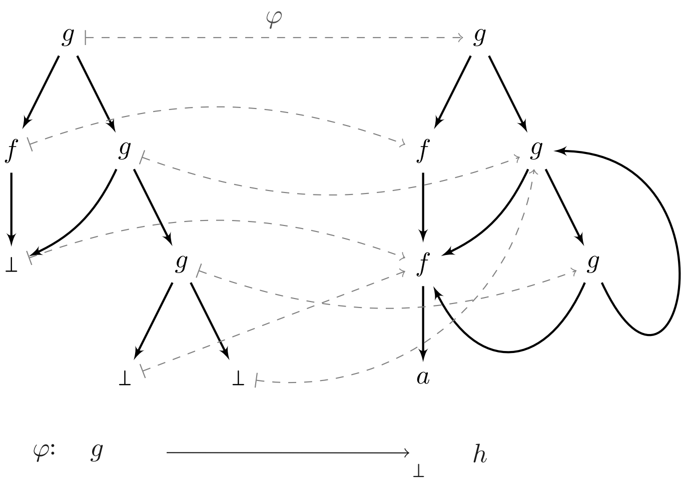

Figure 4.4: Example of a $\perp$-homomorphism injective in non-$\perp$-nodes.

Proposition 2.1.33, to have a complete semilattice. And indeed, the partial order $\leq_\perp$ on terms is a complete semilattice as mentioned in Proposition 2.3.20. Thus, our goal is to obtain a complete semilattice on canonical term graphs.

Just as in the case for terms, we use an additional nullary symbol $\perp$ which is supposed to denoted "undefinedness". That is, we consider term graphs over the extended signature $\Sigma_\perp = \Sigma \uplus \{\perp/0\}$. Term graphs over $\Sigma_\perp$ are also referred to as *partial term graphs*. If it is necessary to make it explicit, we refer to the term graphs in $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ as *total term graphs*. Now we do need the flexibility that the notion of Δ-homomorphisms gives us. We are going to define the partial order $\leq_\perp$ on term graphs using $\perp$-homomorphisms, i.e. Δ-homomorphisms with $\Delta = \{\perp\}$.

Before getting to the technical details of the definition of $\leq_\perp$ on term graphs let us discuss the intention of this partial order. It ought to capture two concepts: Firstly, $g \leq_\perp h$ should state that g and h are consistent, i.e. the information contained in g and h do not contradict each other. Secondly, $g \leq_\perp h$ is supposed to say that h contains more information than g does.

A first attempt to formalise this intuition is to use injective $\perp$-homomorphisms: $g \leq_\perp h$ iff there is an injective $\perp$-homomorphism $\varphi\colon g \to_\perp h$. Alternatively, one could also weaken this slightly by only requiring φ to be injective for non-$\perp$-nodes. Both approaches do capture the intended intuition of $\leq_\perp$. Let us focus on the less restrictive latter alternative. The following considerations also apply to the former: The requirement of φ being injective and homomorphic in non-$\perp$-nodes yields that the non-$\perp$-part of g is isomorphic to an initial fragment of h. Hence, g and h do not contradict each other. Secondly, since φ does not have to be homomorphic for $\perp$-nodes, the images of these nodes w.r.t. φ can be non-$\perp$-nodes. That is, in those places where g is undefined, indicated by the label $\perp$, h might have more information. An example for this approach is illustrated in Figure 4.4. It shows two term graphs g and h and a $\perp$-homomorphism φ which is injective in non-$\perp$-nodes, i.e. no two nodes not labelled with $\perp$ are mapped to a common image.

In fact, both approaches, injective $\perp$-homomorphisms and its weakened variant, define a complete partial order. Yet, both fail to admit glbs of arbitrary non-empty sets and, therefore, are not complete semilattices. Glbs might not even exists for finite non-empty Sets. An example for this is shown in Figure 4.5. Four term graphs are depicted: g_1, g_2, g_3, g_4. Neither of the proposed orders admits a glb for the set $\{g_1, g_2\}$. To appreciate why this is the case, take a look at the term graphs g_3 and g_4. For the weakened variant of the partial order, these are two maximal lower bounds of $\{g_1, g_2\}$ as one can easily verify. For the

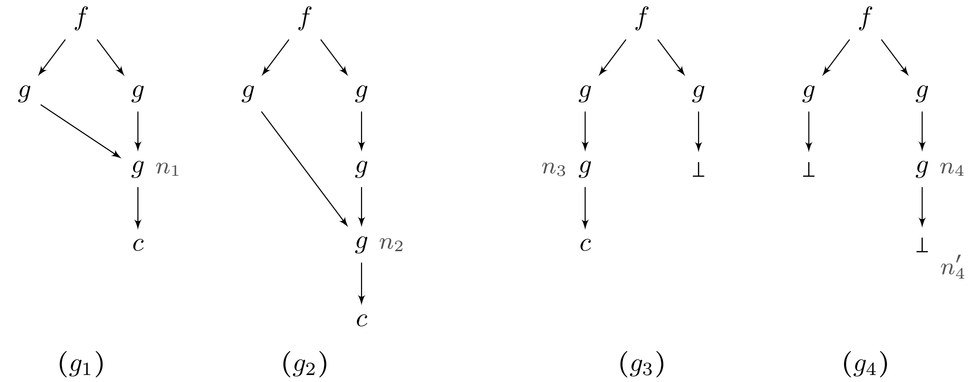

Figure 4.5: Counterexample for candiate partial orders.

injective $\perp$-homomorphism definition of the partial order, one can also derive two maximal lower bounds: Modify g_3 by replacing the right successor of the f-node with a single $\perp$-node and modify g_4 by replacing the left successor of the f-node with a single $\perp$ node, too. Hence, for both orders, there is no greatest lower bound of $\{g_1, g_2\}$ – only two maximal ones, respectively.

The example illustrates the underlying problem quite nicely. In both g_1 and g_2 the f-node spawns two "branches" which are ultimately joined into a single branch in the node n_1 resp. n_2. Yet, the right branch in g_2 is longer than in g_1. Hence, the junction of the two branches cannot be part of a common lower bound. The two different maximal lower bounds of g_1 and g_2 represent two different choices of which of the two branches to "prefer". g_3 prefers the left branch whereas g_4 prefers the right branch.

A solution to this problem, the one that we adopt here, is to strengthen the order $\leq_\perp$ such that g_3 and g_4 are no longer lower bounds of g_1 and g_2. Note that the nodes n_1 and n_2 are shared, i.e. they have two predecessors. Yet, the corresponding nodes in g_3 and g_4, viz. n_3 and n_4 resp. n'_4 are not shared. Hence, we add another restriction that says that corresponding nodes have to have the same "sharing behaviour". What this means will become clear as soon as we will come to the formal definition. Furthermore, note that the sharing that we have in the example is an *acyclic* one (also called *horizontal sharing*). And as it will turn out, only acyclic sharing causes the problem illustrated by this example. *Cyclic sharing* (also called *vertical sharing*), on the other hand, is harmless. The formal approach to this intuition is to take into account the occurrences of the nodes. The different lengths of the right branches in the two term graphs g_1 and g_2 causes the nodes n_1 and n_2 to have different occurrences: n_1 has the occurrences $0 \cdot 0$ and $1 \cdot 0$ whereas n_2 has the occurrences $0 \cdot 0$ and $1 \cdot 0 \cdot 0$. Moreover, as already mentioned, only acyclic sharing is of relevance in this setting. Thus, we restrict our attention to acyclic occurrences:

Notation 4.4.1. Recall that an occurrence π in a term graph g is called cyclic iff there are occurrences π_1, π_2 with $\pi_1 < \pi_2 \leq \pi$ such that $\mathsf{node}_g(\pi_1) = \mathsf{node}_g(\pi_2)$. Otherwise it is called acyclic. We will use the notation $\mathcal{P}^a(g)$ for the set of all acyclic occurrences in g and $\mathcal{P}^a_g(n)$ for the set of all acyclic occurrences of a node n in g. For an acyclic term graph g, a fortiori a term tree, $\mathcal{P}_g(n)$ and $\mathcal{P}^a_g(n)$ coincide, of course.

Definition 4.4.2 (preservation of sharing, strong Δ-homomorphism)
Let $g, h \in \mathcal{G}^\infty(\Sigma)$ and $\varphi\colon g \to_\Delta h$.

(i) Let $n \in N^g$. φ is said to *preserve the sharing* of n if it satisfies the equation

$$\mathcal{P}^a_g(n) = \mathcal{P}^a_h(\varphi(n)) \qquad \text{(preservation of sharing)}$$

(ii) φ is called *strong* if it preserves the sharing of all $n \in N^g$ with $g(n) \notin \Delta$.

Note that since in a term tree each node has a unique occurrence, each Δ-homomorphism between term trees is trivially strong.

We now have strengthened Δ-homomorphism in the way we need it to be able to define the partial order. In fact, the notion of strong Δ-homomorphisms subsumes injectivity for non-Δ-nodes:

Lemma 4.4.3 (strong Δ-homomorphisms are injective for non-Δ-nodes)
Let $g, h \in \mathcal{G}^\infty(\Sigma)$ and $\varphi\colon g \to_\Delta h$ strong. Then φ is injective for all non-Δ-nodes in g. That is, for two nodes $n, m \in N^g$ with $\mathsf{lab}^g(n), \mathsf{lab}^g(m) \notin \Delta$ we have that $\varphi(n) = \varphi(m)$ implies $n = m$.

Proof. Let $n, m \in N^g$ with $\mathsf{lab}^g(n), \mathsf{lab}^g(m) \notin \Delta$ and $\varphi(n) = \varphi(m)$. Since φ is strong, it preserves the sharing of n and m. That is, in particular we have $\mathcal{P}^a_h(\varphi(n)) \subseteq \mathcal{P}_g(n)$ and $\mathcal{P}^a_h(\varphi(m)) \subseteq \mathcal{P}_g(m)$. Moreover, because $\mathcal{P}^a_h(\varphi(n)) = \mathcal{P}^a_h(\varphi(m)) \neq \emptyset$, we can conclude that $\mathcal{P}_g(n) \cap \mathcal{P}_g(m) \neq \emptyset$ and, therefore, $m = n$. ☐

Before we turn to the definition of the partial order $\leq_\perp$ on term graphs and its properties, we want to study strong Δ-homomorphisms.

One can quite easily see that the depth of a node can be defined in terms of its acyclic occurrences.

Lemma 4.4.4 (depth in terms of acyclic occurrences)
Let $g \in \mathcal{G}^\infty(\Sigma)$ and $n \in N^g$. Then $\mathsf{depth}_g(n) = \min\left\{|\pi| \,\middle|\, \pi \in \mathcal{P}^a_g(n)\right\}$.

Proof. Recall that $\mathsf{depth}_g(n) = \min\{|\pi| \mid \pi \in \mathcal{P}_g(n)\}$. Hence, we immediately have the inequation $\mathsf{depth}_g(n) \leq \min\left\{|\pi| \,\middle|\, \pi \in \mathcal{P}^a_g(n)\right\}$. Suppose, that $\mathsf{depth}_g(n) < \min\left\{|\pi| \,\middle|\, \pi \in \mathcal{P}^a_g(n)\right\}$. Then there is some $\pi \in \mathcal{P}_g(n) \setminus \mathcal{P}^a_g(n)$ with $|\pi| \leq |\pi'|$ for all $\pi' \in \mathcal{P}_g(n)$. Since π is cyclic, there are paths π_1, π_2, π_3 with $\pi_2 \neq \varepsilon$, $\pi = \pi_1 \cdot \pi_2 \cdot \pi_3$ and $\mathsf{node}_g(\pi_1) = \mathsf{node}_g(\pi_1 \cdot \pi_2)$. Consequently, $\pi_1 \cdot \pi_3 \in \mathcal{P}_g(n)$ and $|\pi_1 \cdot \pi_3| < |\pi_1 \cdot \pi_2 \cdot \pi_3| = |\pi|$. This is contradicts that $|\pi| \leq |\pi'|$ for all $\pi' \in \mathcal{P}_g(n)$. ☐

This observation then immediately gives us the result that the preservation of sharing of a node also yields a preservation of its depth:

Corollary 4.4.5 (depth preservation of strong Δ-homomorphisms)
Let $g, h \in \mathcal{G}^\infty(\Sigma)$ and $\varphi\colon g \to_\Delta h$ a strong Δ-homomorphism. Then $\mathsf{depth}_g(n) = \mathsf{depth}_h(\varphi(n))$ for all $n \in N^g$ with $\mathsf{lab}^g(n) \notin \Delta$.

Proof. This follows immediately from Lemma 4.4.4 since $\mathsf{lab}^g(n) \notin \Delta$ implies $\mathcal{P}^a_g(n) = \mathcal{P}^a_h(\varphi(n))$ for the strong Δ-homomorphisms φ. ☐

The following lemma provides an equivalent characterisation of strong Δ-homomorphisms that reduces the proof obligations necessary to show that a Δ-homomorphism is strong.

Lemma 4.4.6 (preservation of sharing)
Let $g, h \in \mathcal{G}^\infty(\Sigma)$, $\varphi\colon g \to_\Delta h$. Then φ is strong iff $\mathcal{P}^a_h(\varphi(n)) \subseteq \mathcal{P}_g(n)$ for all $n \in N^g$ with $\mathsf{lab}^g(n) \notin \Delta$.

Proof. The "only if" direction is trivial. For the "if" direction, suppose that φ satisfies $\mathcal{P}^a_h(\varphi(n)) \subseteq \mathcal{P}_g(n)$ for all $n \in N^g$ with $\mathsf{lab}^g(n) \notin \Delta$. In order to prove that φ is strong, we will show that $\mathcal{P}^a_h(\varphi(n)) = \mathcal{P}^a_g(n)$ holds for each $n \in N^g$ with $\mathsf{lab}^g(n) \notin \Delta$.

We first show the inclusion $\mathcal{P}^a_h(\varphi(n)) \subseteq \mathcal{P}^a_g(n)$. For this purpose, let $\pi \in \mathcal{P}^a_h(\varphi(n))$. Due to the hypothesis, this implies that $\pi \in \mathcal{P}_g(n)$. Now suppose that π is cyclic in g, i.e. there are two occurrences π_1, π_2 of a node $m \in N^g$ with $\pi_1 < \pi_2 \leq \pi$. By Lemma 4.3.2, we can conclude that $\pi_1, \pi_2 \in \mathcal{P}_h(\varphi(m))$. This is a contradiction to the assumption that π is acyclic in h. Hence, $\pi \in \mathcal{P}^a_g(n)$.

For the other inclusion, assume some $\pi \in \mathcal{P}^a_g(n)$. Using Lemma 4.3.2 we obtain that $\pi \in \mathcal{P}_h(\varphi(n))$. It remains to be shown that π is acyclic in h. Suppose that this is not true, i.e. there are two occurrences π_1, π_2 of a node $m \in N^h$ with $\pi_1 < \pi_2 \leq \pi$. Note that since $\pi \in \mathcal{P}(g)$, also $\pi_1, \pi_2 \in \mathcal{P}(g)$. Let $m_i = \mathsf{node}_g(\pi_i)$, $i = 1, 2$. According to Lemma 4.3.2, we have that $\varphi(m_1) = m = \varphi(m_2)$. Moreover, observe that $g(\pi_1), g(\pi_2) \notin \Delta$: $g(\pi_1)$ cannot be a nullary symbol because $\pi_1 < \pi \in \mathcal{P}(g)$. The same argument applies for the case that $\pi_2 < \pi$. If this is not the case, then $\pi_2 = \pi$ and $g(\pi) \notin \Delta$ follows from the assumption that $\mathsf{lab}^g(n) \notin \Delta$. Thus, we can apply Lemma 4.4.3 to conclude that $m_1 = m_2$. Consequently, π is cyclic in g, which contradicts the assumption. Hence, $\pi \in \mathcal{P}^a_h(\varphi(n))$. □

From this we obtain that Δ-isomorphisms are, in fact, also strong Δ-homomorphisms.

Corollary 4.4.7 (Δ-isomorphisms are strong)
Let $g, h \in \mathcal{G}^\infty(\Sigma)$. If $\varphi\colon g \cong_\Delta h$, then φ is a strong Δ-homomorphism.

Proof. This follows from Corollary 4.3.6 and Lemma 4.4.6. □

Again with the strengthened notion of homomorphism we get a categorical structure.

Proposition 4.4.8 (category of strong Δ-homomorphisms)
$\mathcal{G}^\infty(\Sigma)$ together with the strong Δ-homomorphisms on $\mathcal{G}^\infty(\Sigma)$ forms a category.

Proof. It is clear that the identity homomorphism is strong. Therefore, by Proposition 4.2.3, it remains to be shown that, for two strong Δ-homomorphisms $\varphi\colon g_1 \to_\Delta g_2$ and $\psi\colon g_2 \to_\Delta g_3$, also the composition $\psi \circ \varphi$ is strong. To this end, let $n \in N^{g_1}$ with $\mathsf{lab}^{g_1}(n) \notin \Delta$. From the labelling condition for φ, we obtain that $\mathsf{lab}^{g_2}(\varphi(n)) \notin \Delta$. Hence, φ preserves the sharing of n and ψ preserves the sharing of $\varphi(n)$. We can combine this to get the equation

$$\mathcal{P}^a_{g_3}(\psi(\varphi(n))) = \mathcal{P}^a_{g_2}(\varphi(n)) = \mathcal{P}^a_{g_1}(n).$$

That is, $\psi \circ \varphi$ preserves the sharing of n. □

Having established the most important properties of strong Δ-homomorphisms, we can now define the partial order $\leq_\bot$ on term graphs.

Definition 4.4.9 (partial order on term graphs)
Let $g, h \in \mathcal{G}^\infty(\Sigma_\bot)$. Then we define

$$g \leq_\bot h \quad \text{iff} \quad \text{there is a strong } \bot\text{-homomorphism } \varphi\colon g \to_\bot h$$

By abuse of notation, we also use $\bot$ to denote the canonical term graph whose only node is labelled with $\bot$. Moreover, $\leq_\bot$ can be lifted to $\mathcal{G}^\infty(\Sigma_\bot)/\cong$ by defining

$$[g]_\cong \leq_\bot [h]_\cong \quad \text{iff} \quad g \leq_\bot h$$

Remark 4.4.10. The extension of $\leq_\bot$ to equivalence classes is easily seen to be well-defined: Suppose that $g \leq_\bot h$ and $g' \cong g$ and $h' \cong h$. Since, by Corollary 4.4.7, isomorphisms are also strong ($\bot$-)homomorphisms, we have two strong $\bot$-homomorphisms $\varphi_1\colon g' \to_\bot g$ and $\varphi_2\colon h \to_\bot h'$. Since $g \leq_\bot h$, there is also a strong $\varphi\colon g \to_\bot h$. Hence, by Proposition 4.4.8, $\varphi_2 \circ \varphi \circ \varphi_1$ is a strong $\bot$-homomorphism from g' to h', i.e. $g' \leq_\bot h'$.

Furthermore, we immediately get that the ordered set $(\mathcal{G}^\infty(\Sigma_\bot)/\cong, \leq_\bot)$ is isomorphic to the ordered set $(\mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot), \leq_\bot)$. We will employ this fact by switching between these structures to be able to use the respective structure that is more convenient for the given setting.

Using the properties of strong Δ-homomorphisms we easily obtain that $\leq_\bot$ is a partial order on $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot)$:

Corollary 4.4.11 ($\leq_\bot$ is a partial order on $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot)$)
$\leq_\bot$ is a partial order on $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot)$.

Proof. Reflexivity and transitivity of $\leq_\perp$ follow immediately from Proposition 4.4.8. For the antisymmetry, assume $g \leq_\perp h$ and $h \leq_\perp g$. By Lemma 4.2.8, this implies $g \cong_\perp h$. Lemma 4.2.6 then yields that $g \cong h$. Hence, according to Proposition 4.3.7, $g = h$. □

This result also shows that on the set $\mathcal{G}^\infty(\Sigma_\perp)$ of all partial term graphs $\leq_\perp$ is a quasi-order.

In order to work with $\leq_\perp$ in a convenient way, we need some more insight into the properties of strong Δ-homomorphisms on canonical term graphs. Again we want to establish a characterisation in terms of occurrences and the equivalence relations $\sim_g$ induced by term graphs g:

Lemma 4.4.12 (characterisation of strong Δ-homomorphisms)
Let $g, h \in \mathcal{G}^\infty(\Sigma)$ and $\varphi\colon g \to_\Delta h$. Then φ is strong iff

$$\pi \sim_h \pi' \quad\Longrightarrow\quad \pi \sim_g \pi' \quad \textit{for all } \pi \in \mathcal{P}(g) \textit{ with } g(\pi) \notin \Delta \textit{ and } \pi' \in \mathcal{P}^a(h).$$

Proof. For the "only if" direction, assume that φ is strong. Moreover, let $\pi \in \mathcal{P}(g)$ with $g(\pi) \notin \Delta$ and $\pi' \in \mathcal{P}^a(h)$ such that $\pi \sim_h \pi'$, and let $n = \mathsf{node}_g(\pi)$. According to Lemma 4.3.2, we get that $\pi \in \mathcal{P}_h(\varphi(n))$. Because of $\pi \sim_h \pi'$, also $\pi' \in \mathcal{P}_h(\varphi(n))$. Since, by assumption, π' is acyclic in h, we know in particular that $\pi' \in \mathcal{P}^a_h(\varphi(n))$. Since φ is strong and $\mathsf{lab}^g(n) \notin \Delta$, we know that φ preserves the sharing of n which yields that $\pi' \in \mathcal{P}_g(n)$. Hence, $\pi \sim_g \pi'$.

For the converse direction, let $n \in N^g$ with $\mathsf{lab}^g(n) \notin \Delta$. We need to show that φ preserves the sharing of n. Due to Lemma 4.4.6, it suffices to show that $\mathcal{P}^a_h(\varphi(n)) \subseteq \mathcal{P}_g(n)$. Since $\mathcal{P}_g(n) \neq \emptyset$, we can choose some $\pi^* \in \mathcal{P}_g(n)$. Then, according to Lemma 4.3.2, also $\pi^* \in \mathcal{P}_h(\varphi(n))$. Let $\pi \in \mathcal{P}^a_h(\varphi(n))$. Then $\pi^* \sim_h \pi$ holds. Since π is acyclic in h and $g(\pi^*) \notin \Delta$, we can use the hypothesis to obtain that $\pi^* \sim_g \pi$ holds which shows that $\pi \in \mathcal{P}_g(n)$. □

By combining this lemma with the corresponding findings on Δ-homomorphisms we are able to obtain an alternative characterisation for $\leq_\perp$ by means of occurrences:

Corollary 4.4.13 (characterisation of $\leq_\perp$)
Let $g, h \in \mathcal{G}^\infty_\mathcal{C}(\Sigma_\perp)$. Then $g \leq_\perp h$ iff all of the following conditions are met:

(a) $\pi \sim_g \pi' \quad\Longrightarrow\quad \pi \sim_h \pi' \quad$ *for all* $\pi, \pi' \in \mathcal{P}(g)$

(b) $\pi \sim_h \pi' \quad\Longrightarrow\quad \pi \sim_g \pi' \quad$ *for all* $\pi \in \mathcal{P}(g)$ *with* $g(\pi) \neq \perp$ *and* $\pi' \in \mathcal{P}^a(h)$

(c) $g(\pi) = h(\pi) \quad$ *for all* $\pi \in \mathcal{P}(g)$ *with* $g(\pi) \neq \perp$.

Proof. This follows immediately from the definition of $\leq_\perp$, Lemma 4.3.5 and Lemma 4.4.12. □

As terms form a special case of canonical term graphs, we are able to derive a characterisation of $\leq_\perp$ on terms from the corollary above. Due to the simpler structure of terms, this characterisation is considerably more succinct:

Corollary 4.4.14 (characterisation of $\leq_\perp$ on $\mathcal{T}^\infty(\Sigma_\perp, \mathcal{V})$)
Let $s, t \in \mathcal{T}^\infty(\Sigma_\perp, \mathcal{V})$. Then $s \leq_\perp t$ iff $s(\pi) = t(\pi)$ holds for all $\pi \in \mathcal{P}(s)$ with $s(\pi) \neq \perp$.

Proof. The "only if" direction follows immediately from Corollary 4.4.13. Also for the "if" direction, we can apply Corollary 4.4.13 as (b) is trivially true for all pairs of terms, and the requirement that $s(\pi) = t(\pi)$ for all $\pi \in \mathcal{P}(s)$ with $s(\pi) \neq \perp$ implies both (a) and (c). □

Recall that we observed that there is a one-to-one correspondence between terms and canonical term trees. From the above corollary, it is easy to see that the order $\leq_\perp$ defined on terms coincides with the order $\leq_\perp$ defined on the corresponding canonical term trees. Hence, identifying terms and canonical term trees will cause no confusion w.r.t. $\leq_\perp$.

Now we have prepared all the necessary tools in order to prove that $\leq_\perp$ is a complete partial order.

Proposition 4.4.15 ($\leq_\bot$ is a cpo)
$\leq_\bot$ is a complete partial order on $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot)$.

Proof. The least element of $\leq_\bot$ is obviously $\bot$. Hence, it remains to be shown that each each directed subset of $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot)$ has a least upper bound. To this end, suppose that G is a directed subset of $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot)$. We define a canonical term graph $\overline{g}$ by giving an occurrence representation $(P, l, \sim)$ where

$$P = \bigcup_{g \in G} \mathcal{P}(g)$$
$$\sim = \bigcup_{g \in G} \sim_g$$
$$l(\pi) = \begin{cases} f & \text{if } f \neq \bot \text{ and } \exists g \in G.\, g(\pi) = f \\ \bot & \text{otherwise} \end{cases}$$

We will make extensive use of Corollary 4.4.13 in order to show that $\overline{g}$ is the lub of G. Therefore, we use (a), (b), (c) to refer to the conditions mentioned there.

At first we need to show that l is indeed well-defined. For this purpose, let $g_1, g_2 \in G$ and $\pi \in \mathcal{P}(g_1) \cap \mathcal{P}(g_2)$ with $g_1(\pi), g_2(\pi) \neq \bot$. Since G is directed, there is some $g \in G$ such that $g_1, g_2 \leq_\bot g$. By (c), we can conclude $g_1(\pi) = g(\pi) = g_2(\pi)$.

Next we show that $(P, l, \sim)$ is indeed an occurrence representation. Recall that $\sim$ needs to be an equivalence relation. For the reflexivity, assume that $\pi \in P$. Then there is some $g \in G$ with $\pi \in \mathcal{P}(g)$. Since $\sim_g$ is an equivalence relation, $\pi \sim_g \pi$ must hold and, therefore, $\pi \sim \pi$. For the symmetry, assume that $\pi_1 \sim \pi_2$. Then there is some $g \in G$ such that $\pi_1 \sim_g \pi_2$. Hence, we get $\pi_2 \sim_g \pi_1$ and, consequently, $\pi_2 \sim \pi_1$. In order to show transitivity, assume that $\pi_1 \sim \pi_2, \pi_2 \sim \pi_3$. That is, there are $g_1, g_2 \in G$ with $\pi_1 \sim_{g_1} \pi_2$ and $\pi_2 \sim_{g_2} \pi_3$. Since G is directed, we find some $g \in G$ such that $g_1, g_2 \leq_\bot g$. By (a), this implies that also $\pi_1 \sim_g \pi_2$ and $\pi_2 \sim_g \pi_3$. Hence, $\pi_1 \sim_g \pi_3$ and, therefore, $\pi_1 \sim \pi_3$.

For the reachability condition, let $\pi \cdot i \in P$. That is, there is a $g \in G$ with $\pi \cdot i \in \mathcal{P}(g)$. Lemma 4.3.13 yields $\pi \in \mathcal{P}(g)$ which in turn implies $\pi \in P$. Moreover, $\pi \cdot i \in \mathcal{P}(g)$ implies that $i < \mathsf{ar}(g(\pi))$. Since $g(\pi)$ cannot be a nullary symbol and in particular not $\bot$, we obtain that $l(\pi) = g(\pi)$. Hence, $i < \mathsf{ar}(l(\pi))$.

For the congruence condition, assume that $\pi_1 \sim \pi_2$ and that $l(\pi_1) = f$. If $f \neq \bot$, then there are $g_1, g_2 \in G$ with $\pi_1 \sim_{g_1} \pi_2$ and $g_2(\pi_1) = f$. Since G is directed, there is some $g \in G$ such that $g_1, g_2 \leq_\bot g$. Hence, by (a) resp. (c), we have $\pi_1 \sim_g \pi_2$ and $g(\pi_1) = f$. Using Lemma 4.3.13 we can conclude that $g(\pi_2) = g(\pi_1) = f$ and that $\pi_1 \cdot i \sim_g \pi_2 \cdot i$ for all $< \mathsf{ar}(g(\pi_1))$. Because $g \in G$, it holds that $l(\pi_2) = f$ and that $\pi_1 \cdot i \sim \pi \cdot i$ for all $i < \mathsf{ar}(l(\pi_1))$. If $f = \bot$, then also $l(\pi_2) = \bot$, for if $l(\pi_2) = f'$ for some $f' \neq \bot$, then, by the symmetry of $\sim$ and the above argument (for the case $f \neq \bot$), we would obtain $f = f'$ and, therefore, a contradiction. Since $\bot$ is a nullary symbol, the remainder of the condition is vacuously satisfied.

This shows that $(P, l, \sim)$ is an occurrence representation which, by Lemma 4.3.15, uniquely defines a canonical term graph. Next we show that the thus obtained term graph $\overline{g}$ is an upper bound for G. To this end, let $g \in G$. We will show that $g \leq_\bot \overline{g}$ by establishing (a),(b) and (c). (a) and (c) are an immediate consequence of the construction. For (b), assume that $\pi_1 \in \mathcal{P}(g)$, $g(\pi_1) \neq \bot$, $\pi_2 \in \mathcal{P}^a(\overline{g})$ and $\pi_1 \sim \pi_2$. We will show that then also $\pi_1 \sim_g \pi_2$ holds. Since $\pi_1 \sim \pi_2$, there is some $g' \in G$ with $\pi_1 \sim_{g'} \pi_2$. Because G is directed, there is some $g^* \in G$ with $g, g' \leq_\bot g^*$. Using (a), we then get that $\pi_1 \sim_{g^*} \pi_2$. Note that since π_2 is acyclic in $\overline{g}$, it is also acyclic in g^*: Suppose that this is not the case, i.e. there are occurrences π_3, π_4 with $\pi_3 < \pi_4 \leq \pi_2$ and $\pi_3 \sim_{g^*} \pi_4$. But then we also have $\pi_3 \sim \pi_4$ which contradicts the assumption that π_2 is acyclic in $\overline{g}$. With this knowledge we are able to apply (b) to $\pi_1 \sim_{g^*} \pi_2$ in order to obtain $\pi_1 \sim_g \pi_2$.

In the final part of this proof, we will show that $\overline{g}$ is the least upper bound of G. For this purpose, let $\hat{g}$ be an upper bound of G, i.e. $g \leq_\bot \hat{g}$ for all $g \in G$. We will show that

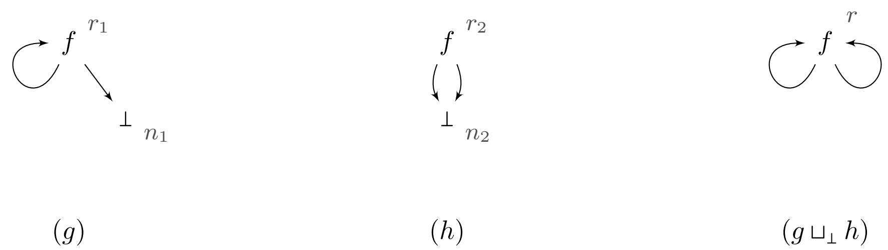

Figure 4.6: Least upper bound $g \sqcup_\perp h$ of compatible term graphs g and h.

$\overline{g} \leq_\perp \hat{g}$ by establishing (a), (b) and (c). For (a), assume that $\pi_1 \sim \pi_2$. Hence, there is some $g \in G$ with $\pi_1 \sim_g \pi_2$. Since, by assumption, $g \leq_\perp \widehat{g}$, we can conclude $\pi_1 \sim_{\widehat{g}} \pi_2$ using (a). For (b), assume $\pi_1 \in P$, $l(\pi_1) \neq \perp$, $\pi_2 \in \mathcal{P}^a(\widehat{g})$ and $\pi_1 \sim_{\widehat{g}} \pi_2$. That is, there is some $g \in G$ with $g(\pi_1) \neq \perp$. Together with $g \leq_\perp \widehat{g}$ this implies $\pi_1 \sim_g \pi_2$ by (b). $\pi_1 \sim \pi_2$ follows immediately. For (c), assume $\pi \in P$ and $l(\pi) = f \neq \perp$. Then there is some $g \in G$ with $g(\pi) = f$. Applying (c) then yields $\widehat{g}(\pi) = f$ since $g \leq_\perp \widehat{g}$. $\square$

From the construction in the previous proof, we immediately get the following corollary:

Corollary 4.4.16 (lub of directed sets)
Let G be a directed subset of $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\perp)$ and $\overline{g} = \bigsqcup^\perp G$. Then the following holds:

(i) $\mathcal{P}(\overline{g}) = \bigcup_{g \in G} \mathcal{P}(g)$, *and*

(ii) $\overline{g}(\pi) = f \neq \perp$ *iff* $\exists g \in G.\ g(\pi) = f$.

Next we will prove that also compatible term graphs have a lub. Recall that compatible elements in a partially ordered set are elements that have an upper bound. The issue that makes the construction of the lub of compatible elements a bit more complicated than in the case of directed sets is illustrated in Figure 4.6. Note that the lub $g \sqcup_\perp h$ of the term graphs g and h has an additional cycle. The fact that in $g \sqcup_\perp h$ the second successor of r has to be r itself is enforced by g saying that the first successor of r_1 is r_1 itself and by h saying that the first and the second successor of r_2 must be identical. Because of the additional cycle in $g \sqcup_\perp h$, we have that the set of occurrences in $g \sqcup_\perp h$ is a proper superset of the union of the sets of occurrences in g and h. This makes the construction of $g \sqcup_\perp h$ using an occurrence representation inappropriate.

A possible strategy to construct the lub is to take all the nodes of the two term graphs in question and identify those nodes that have a common occurrence. In our example, we have four nodes r_1, n_1, r_2 and n_2. At first r_1 and r_2 have to be identified as both have the occurrence ε. Next, r_1 and n_2 are identified as they share the occurrence 0. And eventually, also n_2 and n_1 are identified since they share the occurrence 1. Hence, all four nodes have to be identified. The result is, therefore, a term graph with a single node r. The following lemma and its proof show that, for any two compatible term graphs, this construction always leads to their lub.

Lemma 4.4.17 (compatible elements have lub)
Let $g_1, g_2 \in \mathcal{G}^\infty_\mathcal{C}(\Sigma_\perp)$ be compatible (w.r.t. $\leq_\perp$). Then $\{g_1, g_2\}$ has a least upper bound.

Proof. Instead of the order $\leq_\perp$ on $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\perp)$ we consider the order $\leq_\perp$ on $\mathcal{G}^\infty(\Sigma_\perp)/\cong$. As already mentioned, this is justified by Proposition 4.3.7 and Proposition 4.3.8. We will construct a term graph $\overline{g}$ such that $[\overline{g}]_\cong$ is the least upper bound of $\{[g_1]_\cong, [g_2]_\cong\}$. Let $g_j = (N^j, \mathsf{suc}^j, \mathsf{lab}^j, r^j)$, $j = 1, 2$. Since we are dealing with isomorphism classes, we can assume w.l.o.g. that the nodes in g_j are of the form n^j for $j = 1, 2$. This is only a technical trick for the purpose of reducing the necessary notation and to ensure that N^1 and N^2 are disjoint. As g_1, g_2 are compatible, there is an upper bound $\widehat{g} = (\widehat{N}, \widehat{\mathsf{lab}}, \widehat{\mathsf{suc}}, \widehat{r})$ for these term graphs. That is, there are two strong $\bot$-homomorphisms $\varphi_j \colon g_j \to_\bot \widehat{g}$, $j = 1, 2$.

Let $\overline{M} = N^1 \uplus N^2$. Define the relation $\sim$ on $\overline{M}$ as follows:

$$n^j \sim m^k \quad \text{iff} \quad \mathcal{P}_{g_j}(n^j) \cap \mathcal{P}_{g_k}(m^k) \neq \emptyset$$

$\sim$ is clearly reflexive and symmetric. Hence, its transitive closure $\sim^+$ is an equivalence relation on $\overline{M}$. Now define the term graph $\overline{g} = (\overline{N}, \overline{\mathsf{lab}}, \overline{\mathsf{suc}}, \overline{r})$ as follows:

$$\begin{aligned}
\overline{N} &= \overline{M}/\sim^+ \\
\overline{\mathsf{lab}}(N) &= \begin{cases} f & \text{if } f \neq \bot, \exists n^j \in N.\ \mathsf{lab}^j(n^j) = f \\ \bot & \text{otherwise} \end{cases} \\
\overline{\mathsf{suc}}_i(N) = N' &\quad \text{iff} \quad \exists n^j \in N.\ \mathsf{suc}^j_i(n^j) \in N' \\
\overline{r} &= [r^1]_{\sim^+}
\end{aligned}$$

Note that since $\varepsilon \in \mathcal{P}_{g_1}(r^1) \cap \mathcal{P}_{g_2}(r^2)$, we also have $\overline{r} = [r^2]_{\sim^+}$.

Before we argue about the well-definedness of $\overline{g}$ we need to establish some auxiliary claims:

$$n^j \sim^+ m^k \implies \varphi_j(n^j) = \varphi_k(m^k) \qquad \text{for all } n^j, m^k \in \overline{M} \tag{1}$$

$$\varphi_j(n^j) = \varphi_k(m^k) \implies n^j \sim m^k \qquad \begin{array}{l}\text{for all } n^j, m^k \in \overline{M} \\ \text{with } \mathsf{lab}^j(n^j), \mathsf{lab}^k(m^k) \neq \bot\end{array} \tag{1'}$$

We show (1) by proving that $n^j \sim^p m^k$ implies $\varphi_j(n^j) = \varphi_k(m^k)$ by induction on $p > 0$. If $p = 1$, then $n^j \sim m^k$. Hence, $\mathcal{P}_{g_j}(n^j) \cap \mathcal{P}_{g_k}(m^k) \neq \emptyset$. Additionally, from Lemma 4.3.2 we obtain both $\mathcal{P}_{g_j}(n^j) \subseteq \mathcal{P}_{\widehat{g}}(\varphi_j(n^j))$ and $\mathcal{P}_{g_k}(m^k) \subseteq \mathcal{P}_{\widehat{g}}(\varphi_k(m^k))$. Consequently, we also have that $\mathcal{P}_{\widehat{g}}(\varphi_j(n^j)) \cap \mathcal{P}_{\widehat{g}}(\varphi_k(m^k)) \neq \emptyset$, i.e. $\varphi_j(n^j) = \varphi_k(m^k)$. If $p = q + 1 > 1$, then there is some $o^l \in \overline{M}$ with $n^j \sim o^l$ and $o^l \sim^q m^k$. Applying the induction hypothesis immediately yields $\varphi_j(n^j) = \varphi_l(o^l) = \varphi_k(m^k)$.

For (1'), let $n^j, m^k \in \overline{M}$ with $\mathsf{lab}^j(n^j), \mathsf{lab}^k(m^k) \neq \bot$ and $\varphi_j(n^j) = \varphi_k(m^k)$. Since φ_j and φ_k are strong $\bot$-homomorphisms, we have the following equations:

$$\mathcal{P}^a_{g_j}(n^j) = \mathcal{P}^a_{\widehat{g}}(\varphi_j(n^j)) = \mathcal{P}^a_{\widehat{g}}(\varphi_k(m^k)) = \mathcal{P}^a_{g_k}(m^k).$$

Hence, $\mathcal{P}_{g_j}(n^j) \cap \mathcal{P}_{g_k}(m^k) \neq \emptyset$ and, therefore, $n^j \sim m^k$.

Next we show that $\overline{\mathsf{lab}}$ is well-defined. To this end, let $N \in \overline{N}$ and $n^j, m^k \in N$ such that $\mathsf{lab}^j(n^j) = f_1 \neq \bot$ and $\mathsf{lab}^k(m^k) = f_2 \neq \bot$. We need to show that $f_1 = f_2$. By (1), we have that $\varphi_j(n^j) = \varphi_k(m^k)$. Since $f_1, f_2 \neq \bot$, we can employ the labelling condition for φ_j and φ_k in order to obtain that

$$f_1 = \widehat{\mathsf{lab}}(\varphi_j(n^j)) = \widehat{\mathsf{lab}}(\varphi_k(m^k)) = f_2.$$

To argue that $\overline{\mathsf{suc}}$ is well-defined, we first have to show for all $N \in \overline{N}$ that $\overline{\mathsf{suc}}_i(N)$ is defined iff $i < \mathsf{ar}(\overline{\mathsf{lab}}(N))$. Suppose that $\overline{\mathsf{suc}}_i(N)$ is defined. Then there is some $n^j \in N$ such that $\mathsf{suc}^j_i(n^j)$ is defined. Hence, $i < \mathsf{ar}(\mathsf{lab}^j(n^j))$. Since then also $\mathsf{lab}^j(n^j) \neq \bot$, we have $\overline{\mathsf{lab}}(N) = \mathsf{lab}^j(n^j)$. Therefore, $i < \mathsf{ar}(\overline{\mathsf{lab}}(N))$. If, conversely, there is some $i \in \mathbb{N}$ with $i < \mathsf{ar}(\overline{\mathsf{lab}}(N))$, then we know that $\overline{\mathsf{lab}}(N) = f \neq \bot$. Hence, there is some $n^j \in N$ with $\mathsf{lab}^j(n^j) = f$. Hence, $i < \mathsf{ar}(\mathsf{lab}^j(n^j))$ and, therefore, $\mathsf{suc}^j_i(n^j)$ is defined. Hence, $\overline{\mathsf{suc}}_i(N)$ is defined.

To finish the argument showing that $\overline{\mathsf{suc}}$ is well-defined, we have to show that, for all $N, N_1, N_2 \in \overline{N}$ and $n^j, m^k \in N$ such that $\mathsf{suc}^j_i(n^j) \in N_1$ and $\mathsf{suc}^k_i(m^k) \in N_2$, we indeed have $N_1 = N_2$. As $n^j, m^k \in N$, we have $n^j \sim^+ m^k$ and, therefore, $\varphi_j(n^j) = \varphi_k(n^k)$ according to (1). Since both $\mathsf{suc}^j_i(n^j)$ and $\mathsf{suc}^k_i(m^k)$ are defined, we have $\mathsf{lab}^j(n^j), \mathsf{lab}^k(m^k) \neq \bot$. By (1') we then have $n^j \sim m^k$, i.e. there is some $\pi \in \mathcal{P}_{g_j}(n^j) \cap \mathcal{P}_{g_k}(m^k)$. Consequently, $\pi \cdot i \in \mathcal{P}_{g_j}(\mathsf{suc}^j_i(n^j)) \cap \mathcal{P}_{g_k}(\mathsf{suc}^k_i(m^k))$. Hence, $\mathsf{suc}^j_i(n^j) \sim \mathsf{suc}^k_i(m^k)$ and, therefore, $N_1 = N_2$.

Before we begin the main argument we need establish the following auxiliary claims:

$$\mathcal{P}_{g_j}(n^j) \subseteq \mathcal{P}_{\overline{g}}([n^j]_{\sim^+}) \qquad \text{for all } n^j \in \overline{M} \tag{2}$$

$$\forall \pi \in \mathcal{P}^a_{\overline{g}}(N)\ \exists n^j \in N.\ \mathsf{lab}^j(n^j) \neq \bot, \pi \in \mathcal{P}^a_{g_j}(n^j) \qquad \text{for all } N \in \overline{N} \text{ with } \overline{\mathsf{lab}}(N) \neq \bot \tag{3}$$

$$n^j \sim^+ m^k \implies \mathcal{P}^a_{g_j}(n^j) = \mathcal{P}^a_{g_k}(m^k) \qquad \begin{array}{l}\text{for all } n^j, m^j \in \overline{M} \\ \text{with } \mathsf{lab}^j(n^j), \mathsf{lab}^k(m^k) \neq \bot\end{array} \tag{4}$$

For (2), we will show that $\pi \in \mathcal{P}_{g_j}(n^j)$ implies $\pi \in \mathcal{P}_{\overline{g}}([n^j]_{\sim^+})$ by induction on the length of π. If $\pi = \varepsilon$, then $\varepsilon \in \mathcal{P}_{g_j}(n^j)$, i.e. $n^j = r^j$. Recall that $[r^j]_{\sim^+} = \overline{r}$. Hence, $\varepsilon \in \mathcal{P}_{\overline{g}}([n^j]_{\sim^+})$. If $\pi = \pi' \cdot i$, then $\pi' \cdot i \in \mathcal{P}_{g_j}(n^j)$, i.e., for $m^j = \mathsf{node}_{g_j}(\pi')$, we have $\mathsf{suc}^j_i(m^j) = n^j$. Employing the induction hypothesis, we obtain $\pi' \in \mathcal{P}_{\overline{g}}([m^j]_{\sim^+})$. Additionally, according to the construction of $\overline{g}$, we have $\overline{\mathsf{suc}}_i([m^j]_{\sim^+}) = [n^j]_{\sim^+}$. Consequently, $\pi' \cdot i \in \mathcal{P}_{\overline{g}}([n^j]_{\sim^+})$ holds.

Similarly, we also show (3) by induction on the length of π. If $\pi = \varepsilon$, then we have $\varepsilon \in \mathcal{P}^a_{\overline{g}}(N)$, i.e. $N = \overline{r}$. Since, by assumption, $\overline{\mathsf{lab}}(\overline{r}) \neq \bot$ holds, there is some $j \in \{1, 2\}$ such that $\mathsf{lab}^j(r^j) \neq \bot$. Moreover, we clearly have $\varepsilon \in \mathcal{P}^a_{g_j}(r^j)$. If $\pi = \pi' \cdot i$, then we have $\pi' \cdot i \in \mathcal{P}^a_{\overline{g}}(N)$. Let $N' = \mathsf{node}_{\overline{g}}(\pi')$. Since $\pi' \cdot i$ is acyclic in $\overline{g}$, so is π', i.e. $\pi' \in \mathcal{P}^a_{\overline{g}}(N')$. Moreover, we have that $\overline{\mathsf{suc}}_i(N')$ is defined, i.e. $\overline{\mathsf{lab}}(N')$ is not nullary and in particular not $\bot$. Thus, we can apply the induction hypothesis to obtain some $n^j \in N'$ with $\mathsf{lab}^j(n^j) \neq \bot$ and $\pi' \in \mathcal{P}^a_{g_j}(n^j)$. Hence, according to the construction of $\overline{g}$, we have $\mathsf{lab}^j(n^j) = \overline{\mathsf{lab}}(N')$, i.e. $\mathsf{suc}^j_i(n^j) = m^j$ is defined. Furthermore, we then get $m^j \in N$. Note that $\pi' \cdot i \in \mathcal{P}_{g_j}(m^j)$. Thus, it remains to be shown that $\pi' \cdot i$ is acyclic in g_j. Suppose that $\pi' \cdot i$ is cyclic in g_j. As π' is acyclic in g_j, this means that there is some occurrence $\pi^* < \pi' \cdot i$ with $\pi^* \in \mathcal{P}_{g_j}(m^j)$. Using (2), we obtain that $\pi^* \in \mathcal{P}_{\overline{g}}(N)$. This contradicts the assumption of $\pi' \cdot i$ being acyclic in $\overline{g}$. Hence, $\pi' \cdot i \in \mathcal{P}^a_{g_j}(m^j)$ holds.

For (4), suppose that $n^j \sim^+ m^k$ holds with $\mathsf{lab}^j(n^j), \mathsf{lab}^k(m^k) \neq \bot$. From (1), we obtain $\varphi_j(n^j) = \varphi_k(n^k)$. Moreover, since both n^j and m^k are not labelled with $\bot$, we know that φ_j and φ_k preserve the sharing of n^j and m^k, respectively, which yields the equations

$$\mathcal{P}^a_{g_j}(n^j) = \mathcal{P}^a_{\overline{g}}(\varphi_j(n^j)) = \mathcal{P}^a_{\overline{g}}(\varphi_k(m^k)) = \mathcal{P}^a_{g_k}(m^k).$$

Next we show that $[g_1]_{\cong}, [g_1]_{\cong} \leq_\bot [\overline{g}]_{\cong}$ holds by giving two strong $\bot$-homomorphisms $\psi_j \colon g_j \to_\bot \overline{g}$, $j = 1, 2$. Define $\psi_j \colon N^j \to \overline{N}$ by $n^j \mapsto [n^j]_{\sim^+}$. From (2) and the fact that, according to the construction, $\mathsf{lab}^j(n^j) \neq \bot$ implies $\mathsf{lab}^j(n^j) = \overline{\mathsf{lab}}([n^j]_{\sim^+})$, we immediately get that ψ_j is a $\bot$-homomorphism by applying Lemma 4.3.2. In order to argue that ψ_j is strong, assume that $n^j \in N^j$ with $\mathsf{lab}^j(n^j) \neq \bot$. According to Lemma 4.4.6, it suffices to show that $\mathcal{P}^a_{\overline{g}}(\psi_j(n^j)) \subseteq \mathcal{P}_{g_j}(n^j)$. Suppose that $\pi \in \mathcal{P}^a_{\overline{g}}(\psi_j(n^j))$. Note that, by construction, also $\psi_j(n^j)$ is not labelled with $\bot$. Hence, we can apply (3) to obtain some $m^k \in \psi_j(n^j)$ with $\mathsf{lab}^k(m^k) \neq \bot$ and $\pi \in \mathcal{P}^a_{g_k}(m^k)$. By definition, $m^k \in \psi_j(n^j)$ is equivalent to $n^j \sim^+ m^k$. Therefore, we can employ (4), which yields $\mathcal{P}^a_{g_k}(m^k) = \mathcal{P}^a_{g_j}(n^j)$. Hence, $\pi \in \mathcal{P}^a_{g_j}(n^j)$.

Note that the construction of $\overline{g}$ did not depend on $\widehat{g}$, viz., for any other upper bound $[\widehat{h}]_{\cong}$ of $[g_1]_{\cong}, [g_2]_{\cong}$, we get the same term graph $\overline{g}$. Hence, it is still just an arbitrary upper bound which means that in order to show that $[\overline{g}]_{\cong}$ is the least upper bound, it suffices to show $[\overline{g}]_{\cong} \leq_\bot [\widehat{g}]_{\cong}$. For this purpose, we will devise a strong $\bot$-homomorphism $\psi \colon \overline{g} \to_\bot \widehat{g}$. Define $\psi \colon \overline{N} \to \widehat{N}$ by $[n^j]_{\sim^+} \mapsto \varphi_j(n^j)$. (1) shows that ψ is well-defined. The root condition for ψ follows from the root condition for φ_1:

$$\psi(\overline{r}) = \psi([r^1]_{\sim^+}) = \varphi_1(r^1) = \widehat{r}.$$

For the labelling condition, assume that $\overline{\mathsf{lab}}(N) = f \neq \bot$ for some $N \in \overline{N}$. Then there is some $n^j \in N$ with $\mathsf{lab}^j(n^j) = f$. Therefore, the labelling condition for φ_j yields

$$\widehat{\mathsf{lab}}(\psi(N)) = \widehat{\mathsf{lab}}(\varphi_j(n^j)) = \overline{\mathsf{lab}}(N) = f$$

For the successor condition, let $\overline{\mathsf{suc}}_i(N) = N'$ for some $N, N' \in \overline{N}$. Then there is some $n^j \in N$ with $\mathsf{suc}_i^j(n^j) \in N'$. Therefore, the successor condition for ψ follows from the successor condition for φ_j as follows:

$$\begin{aligned}\psi(\overline{\mathsf{suc}}_i(N)) = \psi(N') = \psi([\mathsf{suc}_i^j(n^j)]_{\sim^+}) &= \varphi_j(\mathsf{suc}_i^j(n^j)) \\ &= \widehat{\mathsf{suc}}_i(\varphi_j(n^j)) = \widehat{\mathsf{suc}}_i(\psi([n^j]_{\sim^+})) = \widehat{\mathsf{suc}}_i(\psi(N))\end{aligned}$$

Finally, we show that ψ is strong. To this end, let $N \in \overline{N}$ with $\overline{\mathsf{lab}}(N) \neq \bot$. That is, there is some $n^j \in N$ with $\mathsf{lab}^j(n^j) \neq \bot$. Recall, that we have shown that $\psi_j\colon g^j \to_\bot \overline{g}$ is strong. That is, we have

$$\mathcal{P}^a_{g_j}(n^j) = \mathcal{P}^a_{\overline{g}}(\psi_j(n^j)) = \mathcal{P}^a_{\overline{g}}([n^j]_{\sim^+}).$$

Analogously, we have $\mathcal{P}^a_{\widehat{g}}(\varphi_j(n^j)) = \mathcal{P}^a_{g_j}(n^j)$ as φ_j is strong. Using this we can obtain the following equations:

$$\mathcal{P}^a_{\widehat{g}}(\psi(N)) = \mathcal{P}^a_{\widehat{g}}(\psi([n^j]_{\sim^+})) = \mathcal{P}^a_{\widehat{g}}(\varphi_j(n^j)) = \mathcal{P}^a_{g_j}(n^j) = \mathcal{P}^a_{\overline{g}}([n^j]_{\sim^+}) = \mathcal{P}^a_{\overline{g}}(N)$$

Hence, ψ is a strong $\bot$-homomorphism from $\overline{g}$ to $\widehat{g}$. ☐

With this lemma we can conclude that $\leq_\bot$ is – as desired – a complete semilattice.

Proposition 4.4.18 ($\leq_\bot$ is a complete semilattice)
$\leq_\bot$ is a complete semilattice on $\mathcal{G}^\infty_{\mathcal{C}}(\Sigma_\bot)$.

Proof. This can be obtained by applying Proposition 2.1.7 to both Proposition 4.4.15 and Lemma 4.4.17. ☐

The practical use of term graphs is chiefly the representation of terms as described by the unravelling operation $\mathcal{U}$ presented in Definition 4.3.10. The following proposition shows that unravelling preserves the ordering of $\leq_\bot$.

Proposition 4.4.19 (unravelling preserves ordering)
Let $g, h \in \mathcal{G}^\infty(\Sigma_\bot)$. If $g \leq_\bot h$, then $\mathcal{U}(g) \leq_\bot \mathcal{U}(h)$.

Proof. Suppose that $g \leq_\bot h$. Note that $\mathcal{U}(g)$ and $\mathcal{U}(h)$ are term trees and we can assume w.l.o.g. that they are canonical. Hence, according to Corollary 4.4.14, it suffices to show that $\mathcal{U}(g)(\pi) = \mathcal{U}(h)(\pi)$ holds for all $\pi \in \mathcal{P}(\mathcal{U}(g))$ with $\mathcal{U}(g)(\pi) \neq \bot$. This, however, follows immediately from $g \leq_\bot h$, according to Corollary 4.4.13, as, for each term graph g', it holds that $\mathcal{U}(g')(\pi) = g'(\pi)$ for all $\pi \in \mathcal{P}(g')$. ☐

Unfortunately, such preservation by the unravelling operation does not hold for lubs and glbs. That is, in general the equations $\mathcal{U}(\bigsqcup G) = \bigsqcup \mathcal{U}(G)$ and $\mathcal{U}(\bigsqcap G) = \bigsqcap \mathcal{U}(G)$ are not valid. The following two examples illustrate this.

Example 4.4.20

(i) Let $G = \{g, h\}$ where g and h are the terms depicted in Figure 4.6. Then we have $\mathcal{U}(G) = \{f(f(f(\dots,\bot),\bot),\bot), f(\bot,\bot)\}$, where $f(f(f(\dots,\bot),\bot),\bot)$ is the term

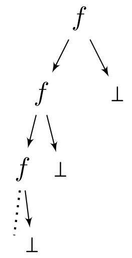

The lub of G is depicted in Figure 4.6. Its unravelling is the term

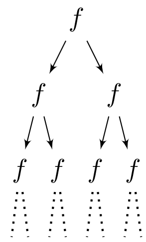

The lub of $\mathcal{U}(G)$ is, however, the term $f(f(f(\dots,\bot),\bot),\bot)$ itself.

(ii) Let G be the set consisting of the term graphs

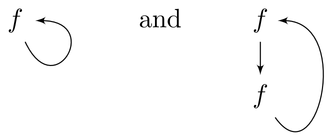

Since the unravelling of both term graphs is the term f^ω, we have that $\mathcal{U}(G) = \{f^\omega\}$. The glb of G is the term $f(\bot)$ whose unravelling is also $f(\bot)$. On the other hand, however, the glb of $\mathcal{U}(G)$ is, of course, f^ω.

The underlying problem causing the preservation of lubs and glbs to fail is that the preservation of the ordering by unravellings only holds for one direction. That is, in general $\mathcal{U}(g) \le_\bot \mathcal{U}(h)$ does not imply $g \le_\bot h$. This is, however, not an issue of the particular partial order $\le_\bot$. In fact, no partial order can satisfy this implication as it would violate its antisymmetry or its reflexivity property. To appreciate this, consider two distinct canonical term graphs g and h that have the same unravelling, i.e. $\mathcal{U}(g) = \mathcal{U}(h)$. For example, take the two term graphs of Example 4.4.20(ii). Due to reflexivity, both $\mathcal{U}(g) \le \mathcal{U}(h)$ and $\mathcal{U}(h) \le \mathcal{U}(g)$ must hold for any partial order $\le$ on canonical term graphs. If $\le$ satisfies the implication mentioned above, then we would obtain $g \le h$ and $h \le g$. This, however, would violate the antisymmetry of $\le$ as h and g are distinct.

Intuitively, partial term graphs represent partial results of computations where $\bot$-nodes act as placeholders denoting the uncertainty or ignorance of the actual "value" at that position. On the other hand, total term graphs do contain all the information of a result of a computation – they have the maximally possible information content. In other words, they are the maximal elements w.r.t. $\le_\bot$. The following proposition confirms this intuition.

Proposition 4.4.21 (total term graphs are the maximal elements)
Let Σ be a non-empty signature. Then $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$ is the set of maximal elements in $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot)$ w.r.t. $\le_\bot$.

Proof. At first we need to show that each element in $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$ is maximal. For this purpose, let $g \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$ and $h \in \mathcal{G}^\infty_\mathcal{C}(\Sigma_\perp)$ such that $g \leq_\perp h$. We have to show that then $g = h$. Since $g \leq_\perp h$, there is a strong $\perp$-homomorphism $\varphi\colon g \to_\perp h$. As g does not contain any $\perp$-node, φ is even a strong homomorphism. By Lemma 4.4.3, φ is injective and, therefore, according to Lemma 4.2.7, an isomorphism. Hence, we obtain that $g \cong h$ and, consequently, using Proposition 4.3.7, that $g = h$.

Secondly, we need to show that $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\perp)$ does not contain any other maximal elements besides those in $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$. Suppose there is a term graph $g \in \mathcal{G}^\infty_\mathcal{C}(\Sigma_\perp) \setminus \mathcal{G}^\infty_\mathcal{C}(\Sigma)$ which is maximal in $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\perp)$. Hence, there is a node $n^* \in N^g$ with $\mathsf{lab}^g(n^*) = \perp$. Let $\overline{n}$ be a fresh node (i.e. $\overline{n} \notin N^g$) and f some k-ary symbol in Σ. Define the term graph h by

$$N^h = N^g \uplus \{\overline{n}\} \qquad\qquad r^h = r^g$$

$$\mathsf{lab}^h(n) = \begin{cases} f & \text{if } n = n^* \\ \perp & \text{if } n = \overline{n} \\ \mathsf{lab}^g(n) & \text{otherwise} \end{cases} \qquad \mathsf{suc}^h(n) = \begin{cases} \overline{n}\cdot\ldots\cdot\overline{n} & \text{if } n = n^* \\ \varepsilon & \text{if } n = \overline{n} \\ \mathsf{suc}^g(n) & \text{otherwise} \end{cases}$$

That is, h is obtained from g by relabelling $n^\star$ with f and setting the $\perp$-labelled node $\overline{n}$ as the target of all outgoing edges of $n^\star$. We assume that $\overline{n}$ was chosen such that h is canonical (i.e. $\overline{n} = \mathcal{P}_h(\overline{n})$). Obviously, g and h are distinct. Define $\varphi\colon N^g \to N^h$ by $n \mapsto n$ for all $n \in N^g$. Clearly, φ defines a strong $\perp$-homomorphism from g to h. Hence, $g \leq_\perp h$. This contradicts the assumption of g being maximal. Consequently, no element in $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\perp) \setminus \mathcal{G}^\infty_\mathcal{C}(\Sigma)$ is maximal. □

Since $\leq_\perp$ forms a complete semilattice on $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\perp)$, it admits a limit inferior for every sequence of partial canonical term graphs. The following lemma explains the intuition behind the limit inferior on canonical term graphs:

Lemma 4.4.22 (limit inferior)
Let $(g_\iota)_{\iota<\alpha}$ be a sequence in $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\perp)$, $\overline{g} = \liminf_{\iota\to\alpha} g_\iota$ and $\pi \in \mathcal{P}(\overline{g})$.

(i) There is some $\beta < \alpha$ such that $\pi \in \mathcal{P}(g_\iota)$ for all $\beta \leq \iota < \alpha$.

(ii) If $\overline{g}(\pi) \neq \perp$, then there is some $\beta < \alpha$ such that $\overline{g}(\pi) = g_\iota(\pi)$ for all $\beta \leq \iota < \alpha$.

Proof. Let $G_\gamma = \{g_\iota \mid \gamma \leq \iota < \alpha\}$ and $h_\gamma = \bigsqcap^\perp G_\gamma$. Then $\overline{g} = \bigsqcup^\perp{}_{\gamma<\alpha} h_\gamma$. Since $\{h_\gamma \mid \gamma < \alpha\}$ is a directed set, we can employ Corollary 4.4.16 which yields that $\mathcal{P}(\overline{g}) = \bigcup_{\gamma<\alpha} \mathcal{P}(h_\gamma)$. That is, there is some $\beta < \alpha$ with $\pi \in \mathcal{P}(h_\beta)$. Because $h_\beta \leq_\perp g_\iota$ for all $\beta \leq \iota < \alpha$, this implies $\pi \in \mathcal{P}(g_\iota)$ for all $\beta \leq \iota < \alpha$ according to Corollary 4.4.13.

Moreover, if $g(\pi) \neq \perp$, according to Corollary 4.4.16, there is some $\beta < \alpha$ with $\overline{g}(\pi) = h_\beta(\pi)$. Consequently, the fact that $h_\beta \leq_\perp g_\iota$ for all $\beta \leq \iota < \alpha$ now implies $\overline{g}(\pi) = g_\iota(\pi)$ for all $\beta \leq \iota < \alpha$ according to Corollary 4.4.13. □

Note that, according to Proposition 2.1.8, the semilattice structure of $\leq_\perp$ also entails that each non-empty set has a glb. The following lemma provides an intuition of glbs – at least for sets of terms:

Lemma 4.4.23 (glb for terms)
Let T be a subset of $\mathcal{T}^\infty(\Sigma_\perp)$ and P a set of occurrences closed under prefixes such that all terms in T coincide in all occurrences in P, i.e. $s(\pi) = t(\pi)$ for all $\pi \in P$ and $s, t \in T$. Then the glb $\overline{t} = \bigsqcap^\perp T$ also coincides with all terms in T in all occurrences in P.

Proof. Construct a term $\widehat{t}$ such that it coincides with all terms in T in all occurrences in P and in all other reachable nodes has the label $\perp$. Then $\widehat{t}$ is a lower bound of T. By construction, $\widehat{t}$ coincides with all terms in T in all occurrences in P. Since $\widehat{t} \leq_\perp \overline{t}$, this property carries over to $\overline{t}$ according to Corollary 4.4.13. □

4.5 A Metric on Term Graphs

In our endeavour to extend infinitary rewriting to the setting of term graph rewriting, we also need to find an appropriate concept of a complete metric space for (canonical) term graphs. Just as for the partial order we want to obtain an extension of the metric space on terms. That is, our aim is to define a metric on canonical term graphs that – when restricted to canonical term trees – coincides with the metric on terms.

Similar to the metric on terms, the metric on canonical term graphs will be defined by means of the least depth where the term graphs differ. We will call this measure *similarity*. To define the similarity of two term graphs, we will employ the partial order that we have studied in the previous section. The greatest lower bound $g \sqcap_\perp h$ of two term graphs g and h allows us to formalise the intuitive notion of similarity.

Definition 4.5.1 (similarity, distance)
Let $g, h \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$ and $\perp$ a fresh nullary symbol (i.e. $\perp \notin \Sigma$). We define the *similarity* of g and h as

$$\mathsf{sim}(g,h) = \perp\text{-}\mathsf{depth}(g \sqcap_\perp h).$$

Recall that $g \sqcap_\perp h$ denotes the greatest lower bound of $\{g,h\}$ w.r.t. the partial order $\leq_\perp$ on $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\perp)$. Using this we can define the distance function $\mathbf{d}$ on $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$ as follows:

$$\mathbf{d}(g,h) = 2^{-\mathsf{sim}(g,h)},$$

where we interpret $2^{-\infty}$ as 0.

Before we will prove that this distance function $\mathbf{d}$ on $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$ indeed defines an ultrametric, we want to establish an alternative characterisation in the style of [AN80]. In order to achieve this, we define a truncation operation on term graphs. This operation removes nodes at a certain depth and fills each of the resulting holes with a fresh $\perp$-node. Let us at first look at the formal definition.

Definition 4.5.2 (truncation of term graphs)
Let $g \in \mathcal{G}^\infty(\Sigma_\perp)$ and $d \in \mathbb{N}$.

(i) Let $n, m \in N^g$. m is an *acyclic predecessor* of n in g if there is an acyclic occurrence $\pi \cdot i \in \mathcal{P}^a_g(n)$ such that $\pi \in \mathcal{P}_g(m)$. The set of all acyclic predecessors of n in g is denoted as $\mathsf{Pre}^a_g(n)$.

(ii) The *truncation nodes* of N^g at d, denoted $N^g_{<d}$, is the least set M satisfying the following conditions for all $n \in N^g$:

$$\mathsf{depth}_g(n) < d \quad \Longrightarrow \quad n \in M \tag{T1}$$

$$n \in M \quad \Longrightarrow \quad \mathsf{Pre}^a_g(n) \subseteq M \tag{T2}$$

(iii) The *fringe nodes* of N^g at d, denoted $N^g_{=d}$, is defined as $\{r^g\}$ if $d = 0$ and as

$$\left\{ n^i \;\middle|\; \begin{array}{l} n \in N^g_{<d}, \mathsf{depth}_g(n) \geq d-1, 0 \leq i < \mathsf{ar}_g(n), \text{ such that} \\ n \notin \mathsf{Pre}^a_g(\mathsf{suc}^g_i(n)) \quad \text{or} \quad \mathsf{suc}^g_i(n) \notin N^g_{<d} \end{array} \right\}$$

if $d > 0$. In the case of $d > 0$, we assume the elements in $N^g_{=d}$ to be pairwise distinct, fresh nodes, i.e. $N^g_{=d} \cap N^g = \emptyset$.

(iv) The *truncation* of g at d, denoted $g|d$, is a term graph defined by

$$N^{g|d} = N^g_{<d} \uplus N^g_{=d} \qquad\qquad r^{g|d} = r^g$$

$$\mathsf{lab}^{g|d}(n) = \begin{cases} \mathsf{lab}^g(n) & \text{if } n \in N^g_{<d} \\ \perp & \text{if } n \in N^g_{=d} \end{cases} \qquad \mathsf{suc}^{g|d}_i(n) = \begin{cases} \mathsf{suc}^g_i(n) & \text{if } n^i \notin N^g_{=d} \\ n^i & \text{if } n^i \in N^g_{=d} \end{cases}$$

Additionally, we define $g|\infty$ to be the term graph g itself.

Before discussing the intuition behind this definition of truncation let us have a look at the rôle of truncation and fringe nodes: The truncation nodes, i.e. the nodes in the set $N^g_{<d}$ are the nodes that are preserved by the truncation. All other nodes in $N^g \setminus N^g_{<d}$ are cut off. The "holes" that are thus created are filled by the fringe nodes, i.e. the nodes in the set $N^g_{=d}$. This is expressed in the condition

$$\mathsf{suc}^g_i(n) \notin N^g_{<d}$$

that has to be fulfilled in order to create a fringe node n^i. A fresh fringe node is inserted for each successor of a truncation node that is not a truncation node.

But there is another circumstance that can give rise to a fringe node. This is encoded in the alternative condition

$$\mathsf{depth}_g(n) \geq d-1 \text{ and } n \notin \mathsf{Pre}^a_g(\mathsf{suc}^g_i(n))$$

that also produces a fringe node n^i. This condition is satisfied when an outgoing edge from a truncation node closes a cycle. An example is depicted in Figure 4.8. For the depth $d = 2$, the node n in the term graph g is a fringe node whose 0-th successor is the root node r. This edge closes a cycle. Hence the truncation at depth 2 contains the fringe node n^0 which is now the 0-th successor of n. The reason for including this is of rather technical nature. We will discuss this later when it is actually needed in Lemma 4.5.11.

The reason for defining the truncation of term graphs in this way is the following: Our goal for the truncation is to make it "compatible" with the definition of the partial order $\leq_\perp$ on term graphs. That is, first of all, the truncation of a term graph is supposed to yield a smaller term graph w.r.t. $\leq_\perp$, viz. $g|d \leq_\perp g$. Hence, whenever a node is kept in the truncation (as opposed to being cut off), also its acyclic occurrences have to be maintained. To do so, with each node also its acyclic predecessors have to be kept in the truncation, i.e. in $N^g_{<d}$. That is the reason for having the closure condition (T2) which is enforced solely for this purpose.

To see this, consider Figure 4.7. It shows a term graph g and the truncation of g at depth 2 once without the closure condition (T2), shown in the middle, and once including (T2), shown on the right. The grey area highlights the nodes that are at depth smaller than 2, i.e. the nodes contained in $N^g_{<2}$ due to (T1) only. The nodes within the area surrounded by a dashed line are all the nodes in $N^g_{<2}$. One can easily observe that with the alternative definition without (T2) we do not have $g|2 \leq_\perp g$. The reason in this particular example is the bottommost h-node whose acyclic sharing in g differs from that in the truncation $g|2$ as one of its predecessors was removed due to the truncation. This effect is avoided in our definition of truncation which always includes all acyclic predecessors of a node. This can be seen in the term graph on the right. It includes both predecessors of the bottommost h-node.

If the truncation construction is applied to term trees, then the result is also a term tree and is equal to the *truncation of terms* as defined in [AN80]. Since in a term tree every node has at most one predecessor, the truncation nodes of a truncation of a term tree t at d are exactly the nodes of t at depth smaller than d, and the fringe nodes are the nodes at depth d. Therefore, the truncation $t|d$ of t at d is a term tree satisfying

$$t|d(\pi) = \begin{cases} t(\pi) & \text{if } |\pi| < d \\ \perp & \text{if } |\pi| = d \\ \text{undefined} & \text{if } |\pi| > d \end{cases}$$

Remark 4.5.3.

(i) In order to argue that the construction of $g|d$ yields a well-defined term graph, one has to show for each $n \in N^g_{<d}$ that $\mathsf{suc}^g_i(n) \in N^g_{<d}$ whenever $n^i \notin N^g_{=d}$: If $\mathsf{depth}_g(n) < d-1$, then $\mathsf{depth}_g(\mathsf{suc}^g_i(n)) < d$ and, hence, $\mathsf{suc}^g_i(n) \in N^g_{<d}$ by (T1). If $\mathsf{depth}_g(n) \geq d-1$, then $n^i \notin N^g_{=d}$ implies that $\mathsf{suc}^g_i(n) \in N^g_{<d}$.

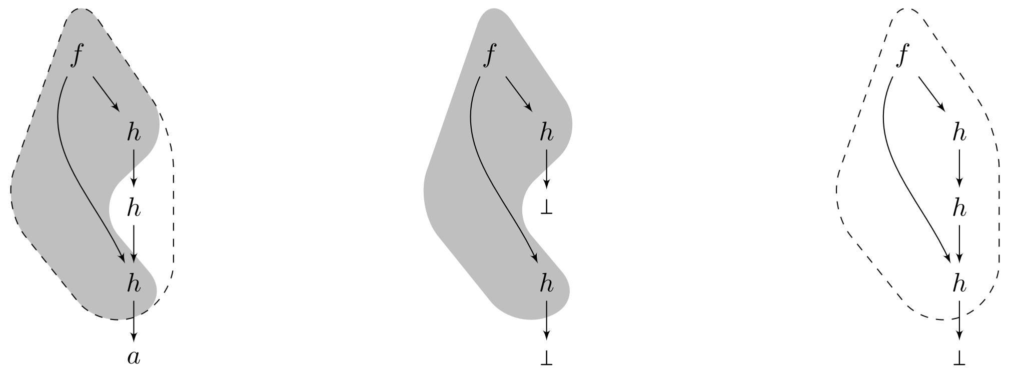

(term graph g) (alternative truncation $g|2$) (actual truncation $g|2$)

Figure 4.7: Example for truncation.

(ii) Note that, for each term graph g, we have $g|0 = \bot$. Hence, for most of the properties of the truncation operation we are going to prove, the case for $d = 0$ is trivial.

The following fact follows immediately from the definition of truncation:

Fact 4.5.4 (truncation preserves labelling up to truncation depth)
Let $g \in \mathcal{G}^\infty(\Sigma_\bot)$ and $d \in \mathbb{N}$. Then $g|d$ and g coincide in all occurrences of depth smaller than d.

The following lemma confirms that we were indeed successful in making the truncation of term graphs compatible with the partial order $\leq_\bot$:

Lemma 4.5.5 (truncation yields a smaller term graph)
Let $g \in \mathcal{G}^\infty(\Sigma_\bot)$ and $d \in \mathbb{N}$. Then $g|d \leq_\bot g$.

Proof. For $d = 0$, this is obvious. Assume $d > 0$. Define the function φ as follows:

$$\begin{aligned}\varphi\colon N^{g|d} &\to N^g\\ N^g_{<d} \ni n &\mapsto n\\ N^g_{=d} \ni n^i &\mapsto \mathsf{suc}^g_i(n)\end{aligned}$$

We will show that φ is a strong $\bot$-homomorphism from $g|d$ to g and, thereby, $g|d \leq_\bot g$.

Since $r^{g|d} = r^g$ and $r^{g|d} \in N^g_{<d}$, we have $\varphi(r^{g|d}) = r^g$ and, therefore, the root condition. Note that all nodes in $N^g_{=d}$ are labelled with $\bot$ in $g|d$. Hence, all non-$\bot$-nodes are in $N^g_{<d}$. Thus, the labelling condition is trivially satisfied as for all $n \in N^g_{<d}$ we have

$$\mathsf{lab}^{g|d}(n) = \mathsf{lab}^g(n) = \mathsf{lab}^g(\varphi(n)).$$

For the successor condition, let $n \in N^g_{<d}$. If $n^i \in N^g_{=d}$, then $\mathsf{suc}^{g|d}_i(n) = n^i$. Hence, we have

$$\varphi(\mathsf{suc}^{g|d}_i(n)) = \varphi(n^i) = \mathsf{suc}^g_i(n) = \mathsf{suc}^g_i(\varphi(n)).$$

If, on the other hand, $n^i \notin N^g_{=d}$, then $\mathsf{suc}^{g|d}_i(n) = \mathsf{suc}^g_i(n) \in N^g_{<d}$. Hence, we have

$$\varphi(\mathsf{suc}^{g|d}_i(n)) = \varphi(\mathsf{suc}^g_i(n)) = \mathsf{suc}^g_i(n) = \mathsf{suc}^g_i(\varphi(n)).$$

This shows that φ is a $\bot$-homomorphism. In order to prove that φ is strong, we will show that $\mathcal{P}^a_g(\varphi(n)) \subseteq \mathcal{P}_{g|d}(n)$ for all $n \in N^g_{<d}$, which is sufficient according to Lemma 4.4.6.

Note that we can replace $\varphi(n)$ by n since $n \in N^g_{<d}$. Therefore, we can show this statement by proving

$$\forall e \in \mathbb{N} \forall n \in N^g_{<d} \forall \pi \in \mathcal{P}^a_g(n).\ (|\pi| = e \implies \pi \in \mathcal{P}_{g|d}(n))$$

by induction on e. If $e = 0$, then $\pi = \varepsilon$. Hence, $n = r^g$ and, therefore, $\pi \in \mathcal{P}_{g|d}(n)$. If $e > 0$, then there is some occurrence π' and natural number i with $\pi = \pi' \cdot i$. Let $m = \mathsf{node}_g(\pi')$. Then we have $m \in \mathsf{Pre}^a_g(n)$ and, therefore, $m \in N^g_{<d}$ by the closure property (T2). And since $\pi' \in \mathcal{P}^a_g(m)$, we can apply the induction hypothesis to obtain that $\pi' \in \mathcal{P}_{g|d}(m)$. Moreover, because $\mathsf{suc}^g_i(m) = n$, this implies that $m^i \notin N^g_{=d}$. Thus, $\mathsf{suc}^{g|d}_i(m) = n$ and, therefore, $\pi' \cdot i \in \mathcal{P}_{g|d}(n)$. □

The gaps that are caused by a truncation due to the removal of nodes are filled by fresh $\bot$-nodes. The following lemma provides a lower bound for the depth of the introduced $\bot$-nodes.

Lemma 4.5.6 ($\bot$-depth in truncated term graphs)
Let Σ be a signature not containing $\bot$, $g \in \mathcal{G}^\infty(\Sigma)$ and $d \in \mathbb{N}$.

(i) $\bot\text{-}\mathsf{depth}(g|d) \geq d$.

(ii) If $d > \mathsf{depth}(g) + 1$, then $g|d = g$, i.e. $\bot\text{-}\mathsf{depth}(g|d) = \infty$.

Proof. (i) From the proof of Lemma 4.5.5, we obtain a strong $\bot$-homomorphism $\varphi\colon g|d \to_\bot g$. Note that the only $\bot$-nodes in $g|d$ are those in $N^g_{=d}$. Each of these nodes has only a single predecessor, a node $n \in N^g_{<d}$ with $\mathsf{depth}_g(n) \geq d - 1$. By Corollary 4.4.5, we also have $\mathsf{depth}_{g|d}(n) \geq d - 1$ for these nodes since φ is strong, n is not labelled with $\bot$ and $\varphi(n) = n$. Hence, we have $\mathsf{depth}_{g|d}(m) \geq d$ for each node $m \in N^g_{=d}$. Consequently, it holds that $\bot\text{-}\mathsf{depth}(g|d) \geq d$.

(ii) Note that if $d > \mathsf{depth}(g) + 1$, then $N^g_{<d} = N^g$ and $N^g_{=d} = \emptyset$. Hence, $g|d = g$. □

Remark 4.5.7. Note that the condition for the statement of clause (ii) in the lemma above reads $d > \mathsf{depth}(g) + 1$ rather than $d > \mathsf{depth}(g)$ as one might expect. The reason for this is that a truncation might cut off an edge that emanates from a node at depth $d - 1$ and closes a cycle. For an example of this phenomenon, take a look at Figure 4.8. It shows a term graph g of depth 1 and its truncation at depth 2. Even though there is no node at depth 2 the truncation introduces a $\bot$-node.

On the other hand, although a term graph has depth more than d the truncation at depth d might still preserve the whole term graph. An example for this behaviour is the family of term graphs h_n, $n > 0$, depicted in Figure 4.8. Each of the term graphs h_n has depth n. Yet, the truncation at depth 2 preserves the whole term graph h_n for each $n > 0$. Even though there might be f-nodes which are at depth ≥ 2 these nodes are directly or indirectly acyclic predecessors of the a-node and are, thus, included in $N^{h_n}_{<2}$.

Lemma 4.5.8 (isomorphic truncations and similarity)
Let $g, h \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$ and $d \in \mathbb{N}$. If $g|d \cong h|d$, then $\mathsf{sim}(g, h) \geq d$.

Proof. W.l.o.g. we can assume that Σ does not contain $\bot$. Assume $g|d \cong h|d$. Then Proposition 4.3.7 yields $\mathcal{C}(g|d) = \mathcal{C}(h|d)$. By Lemma 4.5.5, we have $g|d \leq_\bot g$ and $h|d \leq_\bot h$. Hence, $\mathcal{C}(g|d) \leq_\bot g$ and $\mathcal{C}(h|d) \leq_\bot h$ (cf. Remark 4.4.10). That is, $\mathcal{C}(g|d)$ is a lower bound for g and h. Therefore, $\mathcal{C}(g|d) \leq_\bot g \sqcap_\bot h$. Since this means that there is a $\bot$-homomorphism from $\mathcal{C}(g|d)$ to $g \sqcap_\bot h$ (and, therefore, also from $g|d$ to $g \sqcap_\bot h$), we can employ Lemma 4.2.11 to obtain that $\bot\text{-}\mathsf{depth}(g|d) \leq \bot\text{-}\mathsf{depth}(g \sqcap_\bot h)$. According to Lemma 4.5.6, we have $d \leq \bot\text{-}\mathsf{depth}(g|d)$ which means that we can conclude that $d \leq \bot\text{-}\mathsf{depth}(g \sqcap_\bot h)$ and, thus, $d \leq \mathsf{sim}(g, h)$. □

The lemma below will serve as a tool for the two lemmas that are to follow afterwards.

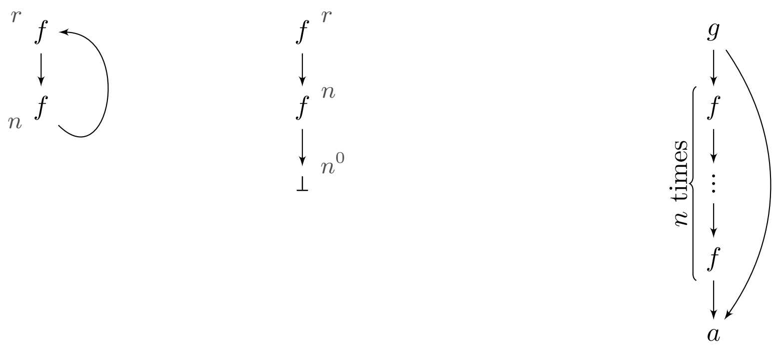

(g) $(g|2)$ $(h_n = h_n|2)$

Figure 4.8: $\bot$-depth in truncated term graphs.

Lemma 4.5.9 (labelling)
Let $g \in \mathcal{G}^\infty(\Sigma)$, $\Delta \subseteq \Sigma^{(0)}$ and $d \in \mathbb{N}$. If $\Delta\text{-}\mathsf{depth}(g) \geq d$, then $\mathsf{lab}^g(n) \notin \Delta$ for all $n \in N^g_{<d}$.

Proof. We will show that $N_\nabla = \{n \in N^g \mid \mathsf{lab}^g(n) \notin \Delta\}$ satisfies the properties (T1) and (T2) of Definition 4.5.2 for the term graph g and depth d. Since $N^g_{<d}$ is the least such set, we then obtain $N^g_{<d} \subseteq N_\nabla$ and, thereby, the claimed statement.

For (T1), let $n \in N^g$ with $\mathsf{depth}_g(n) < d$. Since $\Delta\text{-}\mathsf{depth}(g) \geq d$, we have $\mathsf{lab}^g(n) \notin \Delta$ and, therefore, $n \in N_\nabla$. For (T2), let $n \in N_\nabla$ and $m \in \mathsf{Pre}^a_g(n)$. Then m cannot be labelled with a nullary symbol, a fortiori $\mathsf{lab}^g(m) \notin \Delta$. Hence, we have $m \in N_\nabla$. □

The following two lemmas a rather technical. They state that Δ-homomorphisms preserve truncation nodes and in a stricter sense also fringe nodes.

Lemma 4.5.10 (preservation of truncation nodes)
Let $g, h \in \mathcal{G}^\infty(\Sigma)$, $d \in \mathbb{N}$, $\varphi\colon g \to_\Delta h$ strong, and $\Delta\text{-}\mathsf{depth}(g) \geq d$. Then $\varphi(N^g_{<d}) = N^h_{<d}$.

Proof. Let $N_\nabla = \{n \in N^g \mid \mathsf{lab}^g(n) \notin \Delta\}$. At first we will show that $\varphi(N^g_{<d}) \subseteq N^h_{<d}$. To this end, we will show that $\varphi^{-1}(N^h_{<d}) \cap N_\nabla$ satisfies (T1) and (T2) of Definition 4.5.2 for term graph g and depth d. Since $N^g_{<d}$ is the least such set, we then obtain $N^g_{<d} \subseteq \varphi^{-1}(N^h_{<d}) \cap N_\nabla$ and, a fortiori, $N^g_{<d} \subseteq \varphi^{-1}(N^h_{<d})$. This is equivalent to $\varphi(N^g_{<d}) \subseteq N^h_{<d}$.

For (T1), let $n \in N^g$ with $\mathsf{depth}_g(n) < d$. By Lemma 4.2.9, we then have $\mathsf{depth}_h(\varphi(n)) < d$. Hence, $\varphi(n) \in N^h_{<d}$ by (T1). Moreover, since $\Delta\text{-}\mathsf{depth}(g) \geq d$, we have $\mathsf{lab}^g(n) \notin \Delta$. That is, $n \in \varphi^{-1}(N^h_{<d}) \cap N_\nabla$.

For (T2), let $n \in \varphi^{-1}(N^h_{<d}) \cap N_\nabla$. That is, we have $\varphi(n) \in N^h_{<d}$ and $\mathsf{lab}^g(n) \notin \Delta$. Hence, by (T2), it holds that $\mathsf{Pre}^a_h(\varphi(n)) \subseteq N^h_{<d}$. We have to show now that $\mathsf{Pre}^a_g(n) \subseteq \varphi^{-1}(N^h_{<d}) \cap N_\nabla$. Let $m \in \mathsf{Pre}^a_g(n)$. That is, there is some $\pi \cdot i \in \mathcal{P}^a_g(n)$ with $\pi \in \mathcal{P}_g(m)$. As $\mathsf{lab}^g(n) \notin \Delta$ and φ is strong, φ preserves the sharing of n. Consequently, $\pi \cdot i \in \mathcal{P}^a_h(\varphi(n))$. Moreover, we have $\pi \in \mathcal{P}_h(\varphi(m))$ by Lemma 4.3.2. Hence, $\varphi(m) \in \mathsf{Pre}^a_g(\varphi(n))$ and, therefore, $\varphi(m) \in N^h_{<d}$ by (T2). Additionally, as m has a successor in g it cannot be labelled with a symbol in Δ. Hence, $m \in \varphi^{-1}(N^h_{<d}) \cap N_\nabla$.

In order to prove the converse inclusion $\varphi(N^g_{<d}) \supseteq N^h_{<d}$, we will show that $\varphi(N^g_{<d})$ satisfies (T1) and (T2) for term graph h and depth d. This will prove the abovementioned inclusion since $N^h_{<d}$ is the least such set.

For (T1), let $n \in N^h$ with $\mathsf{depth}_h(n) < d$. By Lemma 4.2.10, there is some $m \in N^g$ with $\mathsf{depth}_g(m) < d$ and $\varphi(m) = n$. Hence, according to (T1), we have $m \in N^g_{<d}$ and, therefore, $n \in \varphi(N^g_{<d})$.

For (T2), let $n \in \varphi(N^g_{<d})$. That is, there is some $m \in N^g_{<d}$ with $\varphi(m) = n$. By (T2), we have $\mathsf{Pre}^a_g(m) \subseteq N^g_{<d}$. We must show that $\mathsf{Pre}^a_h(n) \subseteq \varphi(N^g_{<d})$. Let $n' \in \mathsf{Pre}^a_h(n)$. That is, there

is some $\pi \cdot i \in \mathcal{P}^a_h(n)$ with $\pi \in \mathcal{P}_h(n')$. Since $m \in N^g_{<d}$, we have $\mathsf{lab}^g(m) \notin \Delta$ by Lemma 4.5.9. Consequently, φ preserves the sharing of m which yields that $\pi \cdot i \in \mathcal{P}^a_g(m)$. Note that then also $\pi \in \mathcal{P}(g)$. Let $m' = \mathsf{node}_g(\pi)$. Thus, $m' \in \mathsf{Pre}^a_g(m)$ and, therefore, $m' \in N^g_{<m}$ according to (T2). Moreover, because $\pi \in \mathcal{P}_g(m') \cap \mathcal{P}_h(n')$, we are able to obtain from Lemma 4.3.2 that $\varphi(m') = n'$. Hence, $n' \in \varphi(N^g_{<d})$. □

Lemma 4.5.11 (preservation of fringe nodes)
Let $g, h \in \mathcal{G}^\infty(\Sigma)$, $\varphi\colon g \to_\Delta h$ strong, $d \in \mathbb{N}^+$, $\Delta\text{-}\mathsf{depth}(g) \geq d$, $n \in N^g$, and $0 \leq i < \mathsf{ar}_g(n)$. Then $n^i \in N^g_{=d}$ iff $\varphi(n)^i \in N^h_{=d}$.

Proof. Note that, by Lemma 4.5.9, we have that $\mathsf{lab}^g(n) \notin \Delta$ for all nodes $n \in N^g_{<d}$. Additionally, by Lemma 4.5.10, we obtain $\varphi(N^g_{<d}) = N^h_{<d}$ and, therefore, according to the labelling condition for φ, we get that $\mathsf{lab}^h(n) \notin \Delta$ for all $n \in N^h_{<d}$.

At first we will show the "only if" direction. To this end, let $n^i \in N^g_{=d}$. By definition, we then have $\mathsf{depth}_g(n) \geq d - 1$. Hence, by Corollary 4.4.5, $\mathsf{depth}_h(\varphi(n)) \geq d - 1$. Furthermore, we have that $\mathsf{suc}^g_i(n) \notin N^g_{<d}$ or $n \notin \mathsf{Pre}^a_g(\mathsf{suc}^g_i(n))$. We show now that in either case we can conclude $\varphi(n)^i \in N^h_{=d}$.

Let $\mathsf{suc}^g_i(n) \notin N^g_{<d}$. If we have $\mathsf{suc}^h_i(\varphi(n)) \notin N^h_{<d}$, then $\varphi(n)^i \in N^h_{=d}$. So suppose $\mathsf{suc}^h_i(\varphi(n)) \in N^h_{<d}$. Then, by the successor condition for φ, we have $\varphi(\mathsf{suc}^g_i(n)) \in N^h_{<d} = \varphi(N^g_{<d})$. Hence, there is some $m \in N^g_{<d}$ with $\varphi(m) = \varphi(\mathsf{suc}^g_i(n))$. In the following, we will show that this implies $\varphi(n) \notin \mathsf{Pre}^a_h(\mathsf{suc}^h_i(\varphi(n)))$. Suppose this would not be true, i.e. that $\varphi(n) \in \mathsf{Pre}^a_h(\mathsf{suc}^h_i(\varphi(n)))$. Note that we have the following equations:

$$\mathsf{Pre}^a_h(\mathsf{suc}^h_i(\varphi(n))) = \mathsf{Pre}^a_h(\varphi(\mathsf{suc}^g_i(n))) = \mathsf{Pre}^a_h(\varphi(m)).$$

Consequently, there is some $\pi \cdot i \in \mathcal{P}^a_h(\varphi(m))$ with $\pi \in \mathcal{P}^a_h(\varphi(n))$. Since $n, m \in N^g_{<d}$, we have that φ preserves the sharing of m and n. Hence, we have $\pi \cdot i \in \mathcal{P}^a_g(m)$ and $\pi \in \mathcal{P}^a_g(n)$ which implies that $m = \mathsf{suc}^g_i(n)$. This, however, violates the assumption that $\mathsf{suc}^g_i(n) \notin N^g_{<d}$. Thus, we indeed have $\varphi(n) \notin \mathsf{Pre}^a_h(\mathsf{suc}^h_i(\varphi(n)))$ and, consequently, $\varphi(n)^i \in N^h_{=d}$.

Let $n \notin \mathsf{Pre}^a_g(\mathsf{suc}^g_i(n))$. If $\varphi(n) \notin \mathsf{Pre}^a_h(\mathsf{suc}^h_i(\varphi(n)))$, then $\varphi(n)^i \in N^h_{=d}$. So suppose that $\varphi(n) \in \mathsf{Pre}^a_h(\mathsf{suc}^h_i(\varphi(n)))$. Hence, $\varphi(n) \in \mathsf{Pre}^a_h(\varphi(\mathsf{suc}^g_i(n)))$. If $\mathsf{lab}^g(\mathsf{suc}^g_i(n)) \notin \Delta$, then φ preserves the sharing of $\mathsf{suc}^g_i(n)$ and we would also get $n \in \mathsf{Pre}^a_g(\mathsf{suc}^g_i(n))$ which contradicts the assumption. Hence, $\mathsf{lab}^g(\mathsf{suc}^g_i(n)) \in \Delta$ and, therefore, $\mathsf{suc}^g_i(n) \notin N^g_{<d}$. Thus, we can employ the argument for this case that we have already given above.

We now turn to the converse direction. For this purpose, let $\varphi(n)^i \in N^h_{=d}$. Then $\mathsf{depth}_h(\varphi(n)) \geq d - 1$ and, consequently $\mathsf{depth}_g(n) \geq d - 1$ by Corollary 4.4.5. Additionally, we also have $\mathsf{suc}^h_i(\varphi(n)) \notin N^h_{<d}$ or $\varphi(n) \notin \mathsf{Pre}^a_h(\mathsf{suc}^h_i(\varphi(n)))$. Again we will show that in either case we can conclude $n^i \in N^g_{=d}$.

If $\mathsf{suc}^h_i(\varphi(n)) \notin N^h_{<d}$, then $\varphi(\mathsf{suc}^g_i(n)) \notin N^h_{<d}$ and, therefore, $\varphi(\mathsf{suc}^g_i(n)) \notin \varphi(N^g_{<d})$ according to Lemma 4.5.10. Consequently, $\mathsf{suc}^g_i(n) \notin N^g_{<d}$ which implies that $n^i \in N^g_{=d}$.

Let $\varphi(n) \notin \mathsf{Pre}^a_h(\mathsf{suc}^h_i(\varphi(n)))$. If $n \notin \mathsf{Pre}^a_g(\mathsf{suc}^g_i(n))$, then we get $n^i \in N^g_{=d}$ immediately. So assume that $n \in \mathsf{Pre}^a_g(\mathsf{suc}^g_i(n))$. If $\mathsf{lab}^g(\mathsf{suc}^g_i(n)) \notin \Delta$, then φ would preserve the sharing of $\mathsf{suc}^g_i(n)$. Thereby, we would get $\varphi(n) \in \mathsf{Pre}^a_h(\varphi(\mathsf{suc}^g_i(n)))$ which contradicts the assumption. Hence, $\mathsf{lab}^g(\mathsf{suc}^g_i(n)) \in \Delta$. Consequently, $\mathsf{suc}^g_i(n) \notin N^g_{<d}$ and, therefore, $n^i \in N^g_{=d}$. □

As we have mentioned in the discussion about the definition of the truncation operation, the above lemma depends upon the peculiar definition of fringe nodes – in particular those fringe nodes that are due to the condition

$$\mathsf{depth}_g(n) \geq d - 1 \text{ and } n \notin \mathsf{Pre}^a_g(\mathsf{suc}^g_i(n)).$$

Recall that this condition produces a fringe node for each edge from a truncation node that closes a cycle. Let us have a look at the term graph h depicted in Figure 4.9. If the abovementioned alternative condition for fringe nodes would not be present, then the set $N^h_{=2}$ would be empty (and, thus, $h|2 = h$). Then, however, the strong $\bot$-homomorphism φ

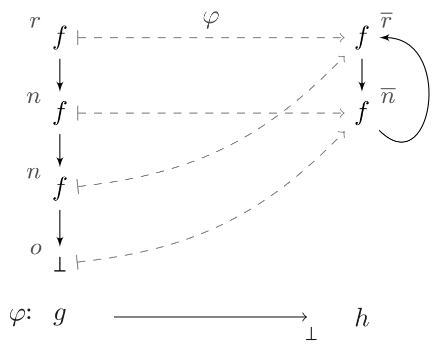

Figure 4.9: Fringe nodes and strong $\bot$-homomorphisms.

illustrated in Figure 4.9 would violate Lemma 4.5.11. Since the node m is cut off from g in the truncation $g|2$, there is a fringe node n^0 in $g|2$. On the other hand, there would be no fringe node $\overline{n}^0$ in $h|2$ if not for the alternative condition above.

Intuitively, the following lemma states that a strong $\bot$-homomorphism has the properties of an isomorphism up to the depth of the shallowest $\bot$-node:

Lemma 4.5.12 ($\leq_\bot$ and truncation)
Let $g, h \in \mathcal{G}^\infty(\Sigma_\bot)$, $g \leq_\bot h$, $d \in \mathbb{N}$ and $\bot\text{-}\mathsf{depth}(g) \geq d$. Then $g|d \cong h|d$.

Proof. For $d = 0$, this is trivial. So assume $d > 0$. Since $g \leq_\bot h$, there is a strong $\bot$-homomorphism $\varphi\colon g \to_\bot h$. Define the function ψ as follows:

$$\begin{gathered}\psi\colon N^{g|d} \to N^{h|d}\\ N^g_{<d} \ni n \mapsto \varphi(n)\\ N^g_{=d} \ni n^i \mapsto \varphi(n)^i\end{gathered}$$

At first we have to argue that ψ is well-defined. For this purpose, we first need that $\varphi(N^g_{<d}) \subseteq N^{g|d}$. Lemma 4.5.10 confirms this. Secondly, we need that $n^i \in N^g_{=d}$ implies $\varphi(n)^i \in N^{g|d}$. This is asserted by Lemma 4.5.11.

Next we show that ψ is a homomorphism from $g|d$ to $h|d$. The root condition is inherited from φ as $r^{g|d} \in N^g_{<d}$. Note that, according to Lemma 4.5.9, we have $\mathsf{lab}^g(n) \neq \bot$ for all $n \in N^g_{<d}$. Hence, φ is homomorphic in $N^g_{<d}$ which means that the labelling condition for nodes in $N^g_{<d}$ is also inherited from φ. For nodes $n^i \in N^g_{=d}$, we have $\mathsf{lab}^{g|d}(n^i) = \bot$. Since, by definition, $\psi(n^i) \in N^h_{=d}$, we can conclude $\mathsf{lab}^{h|d}(\psi(n^i)) = \bot$.

The successor condition is trivially satisfied by nodes in $N^g_{=d}$ as they do not have any successors. Let $n \in N^g_{<d}$ and $0 \leq i < \mathsf{ar}_{g|d}(n)$. We distinguish two cases: At first assume that $n^i \notin N^g_{=d}$. Hence, $\mathsf{suc}^{g|d}_i(n) = \mathsf{suc}^g_i(n) \in N^g_{<d}$. Since, by Lemma 4.5.11, also $\varphi(n)^i \notin N^h_{=d}$, we additionally have $\mathsf{suc}^{h|d}_i(\varphi(n)) = \mathsf{suc}^h_i(\varphi(n))$. Hence, using the successor condition for φ, we can reason as follows:

$$\psi(\mathsf{suc}^{g|d}_i(n)) = \psi(\mathsf{suc}^g_i(n)) = \varphi(\mathsf{suc}^g_i(n)) = \mathsf{suc}^h_i(\varphi(n)) = \mathsf{suc}^{h|d}_i(\varphi(n)) = \mathsf{suc}^{h|d}_i(\psi(n))$$

If, on the other hand, $n^i \in N^g_{=d}$, then $\mathsf{suc}^{g|d}_i(n) = n^i$. Moreover, since also $\varphi(n)^i \in N^h_{=d}$ by Lemma 4.5.11, we have $\mathsf{suc}^{h|d}_i(\varphi(n)) = \varphi(n)^i$, too. Hence, we can reason as follows:

$$\psi(\mathsf{suc}^{g|d}_i(n)) = \psi(n^i) = \varphi(n)^i = \mathsf{suc}^{h|d}_i(\varphi(n)) = \mathsf{suc}^{h|d}_i(\psi(n))$$

This shows that ψ is a homomorphism. Note that, according to Lemma 4.4.3, φ is injective in $N^g_{<d}$. Then also ψ is injective in $N^g_{<d}$. For the same reason, ψ is also injective in

$N^g_{=d}$. Moreover, we have $\psi(N^g_{<d}) \subseteq N^h_{<d}$ and $\psi(N^g_{=d}) \subseteq N^h_{=d}$, i.e. $\psi(N^g_{<d}) \cap \psi(N^g_{=d}) = \emptyset$. Hence, ψ is injective which implies, by Lemma 4.2.7, that ψ is an isomorphism from $g|d$ to $h|d$. $\square$

We can use the above findings in order to obtain the following properties of truncations that one would intuitively expect from a truncation operation:

Corollary 4.5.13 (smaller truncations)
Let $g, h \in \mathcal{G}^\infty(\Sigma)$, $e, d \in \mathbb{N} \cup \{\infty\}$ with $e \leq d$ and $g|d \cong h|d$.

(i) $g|e \cong (g|d)|e$

(ii) $g|d \cong h|d \quad \Longrightarrow \quad g|e \cong h|e$

Proof. We assume w.l.o.g. that $\bot \notin \Sigma$.

(i) For $d = \infty$, this is trivial. Suppose $d \in \mathbb{N}$. From Lemma 4.5.5, we obtain $g|d \leq_\bot g$. Moreover, by Lemma 4.5.6, we have $\bot\text{-}\mathsf{depth}(g|d) \geq d$ and, a fortiori, $\bot\text{-}\mathsf{depth}(g|d) \geq e$. Hence, we can employ Lemma 4.5.12 to get $g|e \cong (g|d)|e$.

(ii) Since $g|d \cong h|d$, we also have $(g|d)|e \cong (h|d)|e$, as the construction of the truncation only depends on the structure of the term graphs. Hence, using

we can conclude

$$g|e \cong (g|d)|e \cong (h|d)|e \cong h|e.$$

$\square$

Lemma 4.5.14 (similarity and isomorphic truncation)
Let $g, h \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$ and $d \in \mathbb{N}$. $\mathsf{sim}(g, h) \geq d$ implies $g|d \cong h|d$.

Proof. We assume w.l.o.g. that $\bot \notin \Sigma$. Let $g^* = g \sqcap_\bot h$. Then $\bot\text{-}\mathsf{depth}(g^*) = \mathsf{sim}(g, h) \geq d$. Since $g^* \leq_\bot g, h$, we can apply Lemma 4.5.12 twice in order to obtain $g|d \cong g^*|d \cong h|d$. $\square$

The previous lemmas stated various details about the connection between truncations and the partial order $\leq_\bot$. The following proposition summarises this by giving an alternative characterisation of similarity.

Proposition 4.5.15 (alternative characterisation of similarity)
Let $g, h \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$. Then $\mathsf{sim}(g, h) = \max\{d \in \mathbb{N} \cup \{\infty\} \mid g|d \cong h|d\}$.

Proof. We assume w.l.o.g. that $\bot \notin \Sigma$. Furthermore, we will use $\mathsf{sim}'(g, h)$ as a shorthand for $\max\{d \in \mathbb{N} \cup \{\infty\} \mid g|d \cong h|d\}$. At first assume that $g = h$. Hence, $g \sqcap_\bot h = g$ and, consequently $\mathsf{sim}(g, h) = \infty$ as g does not contain any $\bot$. On the other hand, this implies $g|\infty \cong h|\infty$, and, therefore, $\mathsf{sim}'(g, h) = \infty$, too. If $g \neq h$, then $g \ncong h$ by Proposition 4.3.7. Moreover, according to Proposition 4.4.21, $g \sqcap_\bot h$ has to contain some $\bot$. Hence, we have both $\mathsf{sim}(g, h) \in \mathbb{N}$ and $\mathsf{sim}'(g, h) \in \mathbb{N}$. We prove that $\mathsf{sim}(g, h) = \mathsf{sim}'(g, h)$ by showing that both $\mathsf{sim}(g, h) \leq \mathsf{sim}'(g, h)$ and $\mathsf{sim}(g, h) \geq \mathsf{sim}'(g, h)$ hold. In order to show the former, let $d = \mathsf{sim}(g, h)$. Then, by Lemma 4.5.14, $g|d \cong h|d$ and, therefore, $\mathsf{sim}'(g, h) \geq d$. To show the latter, let $d = \mathsf{sim}'(g, h)$. Hence, $g|d \cong h|d$. Furthermore, by Lemma 4.5.5, we have both $g|d \leq_\bot g$ and $h|d \leq_\bot h$. Note that, for the canonical representation, we then have $\mathcal{C}(g|d) = \mathcal{C}(h|d)$, $\mathcal{C}(g|d) \leq_\bot g$ and $\mathcal{C}(h|d) \leq_\bot h$ (cf. Proposition 4.3.7 resp. Remark 4.4.10). That is, $\mathcal{C}(g|d)$ is a lower bound of g and h. Thus, $\mathcal{C}(g|d) \leq_\bot g \sqcap_\bot h$ and we can reason as follows:

$$\begin{aligned}
d &\leq \bot\text{-}\mathsf{depth}(g|d) && \text{(Lem. 4.5.6)}\\
&= \bot\text{-}\mathsf{depth}(\mathcal{C}(g|d)) && \text{(Cor. 4.4.5, Cor. 4.4.7)}\\
&\leq \bot\text{-}\mathsf{depth}(g \sqcap_\bot h) && (\mathcal{C}(g|d) \leq_\bot g \sqcap_\bot h\text{, Lem. 4.2.11})\\
&= \mathsf{sim}(g, h)
\end{aligned}$$

$\square$

With this alternative characterisation of similarity proving the distance function on canonical term graphs to be an ultrametric is straightforward:

Proposition 4.5.16 (ultrametric on term graphs)
$(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d})$ *is an ultrametric space.*

Proof. We assume here w.l.o.g. that $\bot \notin \Sigma$. It needs to be shown that $\mathbf{d}$ satisfies the identity, symmetry and strong triangle condition. To this end, we use the alternative characterisation of $\mathsf{sim}(\cdot,\cdot)$ provided by Proposition 4.5.15. The identity condition is met as the following equivalences show:

$$\mathbf{d}(g,h) = 0 \overset{\text{Prop. 4.5.15}}{\iff} \mathsf{sim}(g,h) = \infty \iff g \cong h \overset{\text{Prop. 4.3.7}}{\iff} g = h$$

The symmetry condition is satisfied by

$$\mathbf{d}(g,h) = \bot\text{-}\mathsf{depth}(g \sqcap_\bot h) = \bot\text{-}\mathsf{depth}(h \sqcap_\bot g) = \mathbf{d}(h,g).$$

For the strong triangle condition, we have to show that

$$\mathbf{d}(g_1, g_3) \leq \max\{\mathbf{d}(g_1,g_2), \mathbf{d}(g_2,g_3)\}$$

This is easily seen to be equivalent to

$$\mathsf{sim}(g_1,g_3) \geq \min\{\mathsf{sim}(g_1,g_2), \mathsf{sim}(g_2,g_3)\}.$$

By symmetry, we can assume w.l.o.g. that $\mathsf{sim}(g_1,g_2) \leq \mathsf{sim}(g_2,g_3)$. Let $d = \mathsf{sim}(g_1,g_2)$. We have to show that $\mathsf{sim}(g_1,g_3) \geq d$. Since $\mathsf{sim}(g_1,g_2) = d$, we have $g_1|d \cong g_2|d$ according to Proposition 4.5.15. Because $\mathsf{sim}(g_2,g_3) \geq d$, we have $g_2|d \cong g_3|d$ according to Lemma 4.5.14. Hence, $g_1|d \cong g_3|d$ which yields $\mathsf{sim}(g_1,g_3) \geq d$. ☐

Remark 4.5.17. From now on, we are not dealing with the concrete construction of truncations $g|d$ of term graphs g. Therefore, we will rather use the canonical representation $\mathcal{C}(g|d)$ of $g|d$. In order to avoid the notational overhead, we also adopt the convention outlined in Remark 4.3.9 and write $g|d$ instead of $\mathcal{C}(g|d)$.

The next steps is to show that the obtained ultrametric on canonical term graphs is indeed complete. The following proposition states even more: The limit of Cauchy sequences in the metric space equals the corresponding limit inferior in the partially ordered set. This is also the first step towards proving that the partial order extends the metric in the sense of Definition 3.3.10.

Proposition 4.5.18 (metric limit equals limit inferior)
Let Σ be a signature not containing $\bot$ and $(g_\iota)_{\iota<\alpha}$ a non-empty Cauchy sequence in the metric space $(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d})$. Then $\lim_{\iota\to\alpha} g_\iota = \liminf_{\iota\to\alpha} g_\iota$.

Proof. If α is a successor ordinal, this is trivial, as the limit and the limit inferior are then $g_{\alpha-1}$. Assume that α is a limit ordinal and let $\overline{g}$ be the limit inferior of $(g_\iota)_{\iota<\alpha}$. By Proposition 2.1.33 and Proposition 4.4.18, $\overline{g}$ is well-defined. Since $(g_\iota)_{\iota<\alpha}$ is Cauchy, we obtain that, for each $e \in \mathbb{R}^+$, there is a $\beta < \alpha$ such that, for all ι, ι' with $\beta < \iota, \iota' < \alpha$, we have $\mathbf{d}(g_\iota, g_{\iota'}) < e$. A fortiori, we get that, for each $e \in \mathbb{R}^+$, there is a $\beta < \alpha$ such that, for all ι with $\beta < \iota < \alpha$, we have $\mathbf{d}(g_\beta, g_\iota) < e$. By definition of $\mathbf{d}$, this is equivalent to $2^{-\mathsf{sim}(g_\beta,g_\iota)} < e$. Consequently, we have, for each $d \in \mathbb{N}$, a $\beta < \alpha$ such that $\mathsf{sim}(g_\beta, g_\iota) > d$ for all $\beta < \iota < \alpha$. Due to Lemma 4.5.14, $\mathsf{sim}(g_\beta,g_\iota) > d$ implies $g_\beta|d = g_\iota|d$ which in turn implies $g_\beta|d \leq_\bot g_\iota$ according to Lemma 4.5.5. Hence, $g_\beta|d$ is a lower bound for $G_\beta = \{g_\iota \mid \beta \leq \iota < \alpha\}$. As $\bigsqcap^\bot G_\beta$ is the greatest lower bound of G_β, we get that $g_\beta|d \leq_\bot \bigsqcap^\bot G_\beta$. Moreover, by the definition of the limit inferior, it holds that $\bigsqcap^\bot G_\beta \leq_\bot \overline{g}$. Consequently, $g_\beta|d \leq_\bot \overline{g}$, i.e. we have

$$\forall d \in \mathbb{N} \exists \beta < \alpha\colon \quad g_\beta|d \leq_\bot \overline{g} \tag{1}$$

Applying Lemma 4.5.6 and Lemma 4.5.12 yields $g_\beta|d = \overline{g}|d$. Hence, $\mathsf{sim}(\overline{g}, g_\beta) \geq d$. That is, we have shown that

$$\forall d \in \mathbb{N} \exists \beta < \alpha: \quad \mathsf{sim}(\overline{g}, g_\beta) \geq d$$

Since, for each $e \in \mathbb{R}^+$, we find a $d \in \mathbb{N}$ with $2^{-d} < e$, this implies

$$\forall e \in \mathbb{R}^+ \exists \beta < \alpha: \quad \mathbf{d}(\overline{g}, g_\beta) < e$$

This shows that $(g_\iota)_{i<\alpha}$ converges to $\overline{g}$. Now it remains to be shown that $\overline{g}$ is indeed in $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$, i.e. it does not contain any $\bot$. Suppose that $\overline{g}$ does contain a node labelled with $\bot$. Then $\bot\text{-}\mathsf{depth}(\overline{g}) \in \mathbb{N}$. Let $d = \bot\text{-}\mathsf{depth}(\overline{g}) + 1$. By (1), there is a β with $g_\beta|d \leq_\bot \overline{g}$. By applying Lemma 4.5.6 and Lemma 4.2.11, we then get

$$\bot\text{-}\mathsf{depth}(\overline{g}) + 1 = d \leq \bot\text{-}\mathsf{depth}(g_\beta|d) \leq \bot\text{-}\mathsf{depth}(\overline{g}).$$

This is a contradiction. Hence, $\overline{g}$ is indeed in $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$. □

This result has two obvious but important consequences: Firstly, whenever the limit (w.r.t. the metric) of a sequence of canonical term graphs exists, it is equal to the limes inferior (w.r.t. the partial order) of this sequence. Secondly, this shows that the metric space $(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d})$ is complete:

Proposition 4.5.19 (completeness of metric on term graph)
The metric space $(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d})$ is complete.

Proof. Immediate consequence of Proposition 4.5.18. □

From Proposition 4.5.18, we know that the limit inferior in the partially ordered set is at least as powerful as the limit in the metric space. The following proposition shows that the limit inferior restricted to total term graphs is not more powerful than the limit.

Proposition 4.5.20 (total limit inferior equals limit)
Let $(g_\iota)_{\iota<\alpha}$ be a non-empty sequence in $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$ with $\bot \notin \Sigma$. If $\liminf_{\iota\to\alpha} g_\iota \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$, then $\liminf_{\iota\to\alpha} g_\iota = \lim_{\iota\to\alpha} g_\iota$.

Proof. If α alpha is a successor ordinal, then both the limit and the limit inferior are equal to $g_{\alpha-1}$. Let α be a limit ordinal. According to Proposition 4.5.18, in order to show that limit and limit inferior coincide, it suffices to prove that $(g_\iota)_{\iota<\alpha}$ is Cauchy. For this purpose, assume that $(g_\iota)_{\iota<\alpha}$ is not Cauchy. Then there is some $e \in \mathbb{R}^+$ such that, for all $\beta < \alpha$, there are $\beta < \iota, \iota' < \alpha$ with $\mathbf{d}(g_\iota, g_{\iota'}) \geq e$. Take some $d \in \mathbb{N}$ with $e \geq 2^{-d}$. Then we have, for each $\beta < \alpha$, some $\beta < \iota, \iota' < \alpha$ with $\mathsf{sim}(g_\iota, g_{\iota'}) \leq d$, i.e. $\bot\text{-}\mathsf{depth}(g_\iota \sqcap_\bot g_{\iota'}) \leq d$. Define $G_\beta = \{g_\iota \,|\, \beta \leq \iota < \alpha\}$ and $h_\beta = \bigsqcap^\bot G_\beta$ for each $\beta < \alpha$. Note that for two ι, ι' with $\beta < \iota, \iota' < \alpha$ we have $h_\beta \leq_\bot g_\iota \sqcap_\bot g_{\iota'}$ since $g_\iota, g_{\iota'} \in G_\beta$. Thus, by employing Lemma 4.2.11, we obtain $\bot\text{-}\mathsf{depth}(h_\beta) \leq \bot\text{-}\mathsf{depth}(g_\iota \sqcap_\bot g_{\iota'})$. Since there are, for each $\beta < \alpha$, some $\beta < \iota, \iota' < \alpha$ with $\bot\text{-}\mathsf{depth}(g_\iota \sqcap_\bot g_{\iota'}) \leq d$, we, therefore, have $\bot\text{-}\mathsf{depth}(h_\beta) \leq d$ for each $\beta < \alpha$. That is,

$$\text{for each } \beta < \alpha \text{ there is some } \pi \in \mathcal{P}(h_\beta) \text{ with } |\pi| \leq d \text{ such that } h_\beta(\pi) = \bot. \tag{1}$$

Let $\overline{g} = \liminf_{\iota\to\alpha} g_\iota$. Note that $\overline{g} = \bigsqcup^\bot{}_{\beta<\alpha} h_\beta$. Since $\{h_\beta \,|\, \beta < \alpha\}$ is a directed set, we can employ Corollary 4.4.16 which yields that $\mathcal{P}(\overline{g}) = \bigcup_{\beta<\alpha} \mathcal{P}(h_\beta)$. Therefore, we can rephrase (1) in order to obtain that, for all $\beta < \alpha$, there is a $\pi \in \mathcal{P}(\overline{g})$ with $|\pi| \leq d$ such that $h_\beta(\pi) = \bot$. According to Lemma 4.1.13, there are only finitely many occurrences in $\overline{g}$ of length at most d. Hence, as α is a limit ordinal, there is some occurrence π^* in $\overline{g}$ such that

$$\text{for any } \beta < \alpha, \text{ there is some } \beta \leq \gamma < \alpha \text{ with } h_\gamma(\pi^*) = \bot. \tag{2}$$

Note that $(h_\iota)_{\iota<\alpha}$ is a $\leq_\bot$-chain. From Corollary 4.4.13, we know that whenever there are two term graphs g, h with $g \leq_\bot h$ and $h(\pi) = \bot$, then also $g(\pi) = \bot$ provided $\pi \in \mathcal{P}(g)$. We now show that

$$h_\beta(\pi^*) = \bot \text{ for any } \beta < \alpha \text{ with } \pi^* \in \mathcal{P}(h_\beta). \tag{3}$$

Let $\beta < \alpha$ with $\pi^* \in \mathcal{P}(h_\beta)$. Due to (2), there is some $\beta \leq \gamma < \alpha$ with $h_\gamma(\pi^*) = \bot$. As $(h_\iota)_{\iota<\alpha}$ is a $\leq_\bot$-chain, we then have $h_\beta \leq_\bot h_\gamma$ and, therefore, $h_\beta(\pi^*) = \bot$. This proves (3). From (3), we obtain, according to Corollary 4.4.16, that $\overline{g}(\pi^*) = \bot$. This is a contradiction to the assumption that $\overline{g} \in \mathcal{G}^\infty(\Sigma)$. Hence, $(g_\iota)_{\iota<\alpha}$ is Cauchy. □

Note that Proposition 4.5.20 depends on the finiteness of the arity of the symbols in the signature, just as Lemma 4.1.13 does – which is used in the proof above. This observation also holds for terms as the following example shows:

Example 4.5.21
Let $\Sigma = \{f/\omega, a/0, b/0\}$ and $(g_i)_{i<\omega}$ a sequence with

$$\begin{aligned} g_0 &= f(a, a, a, a, a \dots), \\ g_1 &= f(b, a, a, a, a \dots), \\ g_2 &= f(b, b, a, a, a \dots), \\ g_3 &= f(b, b, b, a, a \dots), \\ &\vdots \end{aligned}$$

$(g_i)_{i<\omega}$ has the limit inferior $f(b, b, b, b, b, \dots)$. On the other hand, the sequence is not even Cauchy since, for each $i \neq j$, we have $\mathsf{sim}(g_i, g_j) = 1$ and, therefore, $\mathbf{d}(g_i, g_j) = \frac{1}{2}$.

Note that in conjunction Proposition 4.5.18 and Proposition 4.5.20 state that limits and limits inferior coincide on total term graphs. Since, according to Proposition 4.4.21, total term graphs are precisely the maximal term graphs, we obtain that the partial order on partial term graphs extends the metric space on total term graphs:

Proposition 4.5.22 (partial order extends metric on term graphs)
$(\mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot), \leq_\bot)$ *extends* $(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d})$.

Proof. This is an immediate consequence of Proposition 4.5.18, Proposition 4.5.20 and Proposition 4.4.21. □

The following proposition shows that the limit inferior is invariant to truncations that are performed at increasing depths.

Proposition 4.5.23 (limit inferior of truncations)
Let $(t_\iota)_{\iota<\lambda}$ be a sequence in $\mathcal{T}^\infty(\Sigma_\bot)$ and $(d_\iota)_{\iota<\lambda}$ a sequence in $\mathbb{N}$ such that λ is a limit ordinal and $(d_\iota)_{\iota<\lambda}$ tends to infinity. Then $\liminf_{\iota\to\lambda} t_\iota = \liminf_{\iota\to\lambda} t_\iota|d_\iota$.

Proof. Let $\overline{t} = \liminf_{\iota\to\lambda} t_\iota|d_\iota$ and $\widehat{t} = \liminf_{\iota\to\lambda} t_\iota$. By Lemma 4.5.5, we have that $t_\iota|d_\iota \leq_\bot t_\iota$ for each $\iota < \lambda$. Hence, we also have that $\overline{t} \leq_\bot \widehat{t}$. Thus, it remains to be shown that also $\widehat{t} \leq_\bot \overline{t}$ holds. That is, according to Corollary 4.4.11, we have to show that $\widehat{t}(\pi) = \overline{t}(\pi)$ holds for all $\pi \in \mathcal{P}_{\setminus\bot}(\widehat{t})$.

Let $\pi \in \mathcal{P}_{\setminus\bot}(\widehat{t})$. That is, $\widehat{t}(\pi) = f \neq \bot$. Hence, by Corollary 4.4.16, there is some $\alpha < \lambda$ with $(\bigsqcap^\bot{}_{\alpha\leq\iota<\lambda} t_\iota)(\pi) = f$. Let $P = \{\pi' \mid \pi' \leq \pi\}$ be the set of all prefixes of π. Note that it holds that $\bigsqcap^\bot{}_{\alpha\leq\iota<\lambda} t_\iota \leq_\bot t_\iota$ for all $\iota < \lambda$. Hence, we can apply Corollary 4.4.13 to obtain that $\bigsqcap^\bot{}_{\alpha\leq\iota<\lambda} t_\iota$ and t_ι coincide in all occurrences in P in particular for all $\alpha \leq \iota < \lambda$. Because $(d_\iota)_{\iota<\lambda}$ tends to infinity, there is some $\alpha \leq \beta < \lambda$ such that $d_\iota > |\pi|$ for all $\beta \leq \iota < \lambda$. Consequently, since $t_\iota|d_\iota$ and t_ι coincide in all occurrences of length smaller than d_ι for all $\iota < \lambda$, we have that $t_\iota|d_\iota$ and t_ι coincide in all occurrences in P for all $\beta \leq \iota < \lambda$. Hence, $t_\iota|d_\iota$ and $\bigsqcap^\bot{}_{\alpha\leq\iota<\lambda} t_\iota$ coincide in all occurrences in P for all $\beta \leq \iota < \lambda$. Hence, according to Lemma 4.4.23, $\bigsqcap^\bot{}_{\alpha\leq\iota<\lambda} t_\iota$ and $\bigsqcap^\bot{}_{\beta\leq\iota<\lambda} t_\iota|d_\iota$ coincide in all occurrences in P. Particularly, it holds that $(\bigsqcap^\bot{}_{\beta\leq\iota<\lambda} t_\iota|d_\iota)(\pi) = f$ which in turn implies by Corollary 4.4.16 that $\overline{t}(\pi) = f$. □

Chapter 5

Infinitary Term Rewriting

This chapter is concerned with the theory of infinitary term rewriting. That is, we present and investigate properties of transfinite reductions induced by infinitary term rewriting systems. Compared to the well-established area of finitary term rewriting this means dropping two major restrictions: Firstly, instead of finite terms we have to consider possibly infinite terms. Secondly, reduction sequences of length beyond ω have to be taken into account. The former generalisation, seemingly innocent at first glance, turns out to have considerable consequences – even to the behaviour of finitary properties. A detailed discussion of these phenomena (and their absence in some cases) is conducted in Section 5.1. The latter generalisation concerning the reduction sequences under consideration was already investigated – however, from a chiefly abstract point of view – in Chapter 3. Providing an insight into the properties of transfinite reductions on terms is the main goal of this chapter.

Chapter 3 has introduced a number of different notions of transfinite reductions. A priori there is no “natural” choice of transfinite reductions that seems superior to the others. One can argue that weak convergence is more elegant due to its relative simplicity compared to strong convergence. Yet, because of its more restrictive character, strong convergence is considerably more well-behaved as we will see. The same can be said about the choice between the MRS and the PRS model of transfinite rewriting. As we will learn in Section 5.2, reductions in the former model are merely a well-defined special case of reductions in the latter model. Unlike weak convergence, however, the generalisation that the PRS model allows does lead in some cases to stronger properties as the investigations performed in Section 5.5 will confirm.

Section 5.3 and Section 5.4 summarise the most important results already known for the established notions of weakly resp. strongly convergent reductions in the MRS model. The intention pursued in these two sections is to give an impression of the properties of infinitary rewriting and to provide a comparison of weak and strong convergence. In particular, some of the shortcomings of weak convergence compared to strong convergence are revealed.

Finally, in Section 5.5 strongly convergent reductions in the PRS model are investigated. As this model of transfinite reductions was just introduced in this thesis, there are only few properties known for it. Blom [Blo04] has considered a model of transfinite reductions quite similar to the PRS model. However, he has investigated λ-calculi and their – in our terminology – strongly convergent reductions. We, on the other hand, restrict our attention to term rewriting systems and try to reproduce results known for the MRS model as presented in Section 5.4. The analysis is mainly concerned with confluence properties. Yet, along the way we also find other properties and, most importantly, an equivalence to certain reduction systems, called Böhm reductions. The latter result provides a deep insight into the essential difference between the MRS and the PRS model of infinitary term rewriting.

5.1 Finitary Properties on Infinite Terms

In Section 3.3, we have studied the difference between finitary properties, like SN, and their infinitary counterparts, like SN^∞. But not only the length of the reduction sequences under consideration is significant. The choice of the underlying set of terms that are allowed for forming a reduction sequence does also make a difference. For example, we will see in Section 5.2 that the MRS model and the PRS model of infinitary term rewriting coincide when we restrict the attention to the subset of total terms. The discussion in this section is concerned with the ramifications of extending the scope of term rewriting from finite terms to possibly infinite terms. The theory of (finitary) term rewriting is usually concerned with rewriting on the set $\mathcal{T}(\Sigma, \mathcal{V})$ of finite terms, whereas in infinitary term rewriting we consider the set of possibly infinite terms $\mathcal{T}^\infty(\Sigma, \mathcal{V})$. We will learn that this extension of the domain has an impact even to finitary properties. First, we will have a look at termination properties. Afterwards, we will consider confluence properties.

5.1.1 Termination Properties

One might already anticipate that the shift to possibly infinite terms has severe consequences for finite termination properties. Indeed, virtually everything we know about termination properties on finite terms becomes invalid when infinite terms are considered as well. Of course, if a systems is terminating on $\mathcal{T}^\infty(\Sigma, \mathcal{V})$, it is terminating on $\mathcal{T}(\Sigma, \mathcal{V})$, too. The converse, however, is not true. The following example illustrates this:

Example 5.1.1
Consider the TRS $\mathcal{R}$ with the single rule $f(x) \to x$. On finite terms, $\mathcal{R}$ is, of course, SN and, thus, also WN. Yet, if, additionally, infinite terms are considered, $\mathcal{R}$ is neither SN nor WN as the infinite term f^ω has no normal form. We only have the infinite reduction

$$f^\omega \to f^\omega \to f^\omega \to \ \dots$$

For infinite terms, $\mathcal{R}$ does even violate several weaker variants of termination, viz. *non-loopingness* and *acyclicity*: A TRS is *non-looping* if it does not allow reductions of the form $t \to^+ C[t\sigma]$, and it is called *acyclic* if it does not allow reductions of the form $t \to^+ t$. Both properties are obviously violated by $\mathcal{R}$ on infinite terms.

In a different sense, also some stronger variants of termination are affected such as ω-*termination*, *polynomial termination*, and *simple termination* (cf. [Ter03]). These properties do only depend on the set of rules of a TRS and are, thus, independent of the choice of the underlying set of terms. However, all three properties imply termination for reductions on finite terms. $\mathcal{R}$ can easily be shown to be polynomially terminating and, thus, is also ω-, and simply terminating. Yet, $\mathcal{R}$ is not terminating for reductions on $\mathcal{T}^\infty(\Sigma, \mathcal{V})$. Hence, the implications that we have for reductions on finite terms break when considering infinite terms. The problem that arises here is that these properties are based on proof techniques for termination on finite terms, which fail for reductions on infinite terms.

Rewrite strategies are also affected. For example, consider the innermost reduction strategy: A reduction sequence is said to be *innermost* if every reduction step in it contracts an innermost redex. A redex is called *innermost* if it does not contain a proper subterm which is a redex. The essential difference between finite and infinite terms is that in finite terms the existence of a redex also implies the existence of an innermost redex. In other words: A finite term is a normal form iff it is a normal form w.r.t. innermost reduction. This is not true for infinite terms. Returning to the TRS $\mathcal{R}$ of Example 5.1.1, we can see that the term f^ω is not a normal form. However, f^ω does not have an innermost redex. Each redex in f^ω is of the form f^ω and, thus, has proper subterms which are redexes themselves. Hence, f^ω is a normal form w.r.t. innermost reduction although it not a normal form.

In particular, this has ramifications for termination properties defined w.r.t. innermost reduction: *Innermost termination* (SIN) is termination w.r.t. innermost reductions and *innermost normalisation* (WIN) is normalisation w.r.t. innermost reductions. For reductions on finite terms, we have the implications

$$\mathsf{SN} \implies \mathsf{SIN} \implies \mathsf{WIN} \implies \mathsf{WN}$$

For reductions on possibly infinite terms, the fist two implications also hold true, i.e. it holds that

$$\mathsf{SN} \implies \mathsf{SIN} \implies \mathsf{WIN}$$

This is simply due to the fact that the set of innermost reductions is a subset of the set of all reductions. On the other hand, the implications from WIN to WN and from SIN to WN do not hold in general: The TRS $\mathcal{R}$ of Example 5.1.1 is not WN on $\mathcal{T}^\infty(\Sigma, \mathcal{V})$. However, it is certainly SIN as it is SN and, thus, SIN on finite terms, and the only infinite term, viz. f^ω, is also SIN since it is in normal form w.r.t. innermost reduction. Again, this suggests that innermost reductions are not meaningful in the setting of infinite terms.

5.1.2 Confluence Properties

Compared to termination properties the situation is completely different for confluence properties: Almost every confluence property that one might care about is invariant under the inclusion of infinite terms. The key argument to prove this was given in [Luc01]. It is based on an abstraction technique: For every finite reduction sequence $S\colon s \to^\star t$ in an TRS on $\mathcal{T}^\infty(\Sigma, \mathcal{V})$, there is an upper bound d on the depth where rewrite rules are applied, simply by the fact that S contains only finitely many steps. By the same argument, the depth of the left-hand sides of the rules applied in S has an upper bound h. Hence, we can replace subterms in S at depth $h+d+1$ with fresh variables $x_1, \dots, x_n$ without affecting the applicability of the rules applied in S. We do this in a consistent way, i.e. subterms that are equal are replaced by the same variable, and different subterms are replaced by different variables. Thus, we obtain a reduction sequence S' on $\mathcal{T}(\Sigma, \mathcal{V})$. Moreover, this defines a substitution σ such that $x_i\sigma$ is the term that was replaced by variable x_i for all $1 \le i \le n$. By applying σ to each term in S', we can obtain the original reduction sequence S.

We can use this construction for the CR property as follows: Suppose we have a TRS $\mathcal{R}$ that is CR for reductions on finite terms. In order to show that $\mathcal{R}$ is also CR on infinite terms, suppose we have two reductions $S\colon t \to^\star t_1$ and $T\colon t \to^\star t_2$ on possibly infinite terms. By using the above construction, we can devise two reduction sequences $S'\colon t' \to^\star t'_1$ and $T'\colon t' \to^\star t'_2$ on finite terms and a single substitution σ such that the original reduction sequences can be obtained from S' and T' by applying σ to each term of the respective reduction sequence. Since $\mathcal{R}$ is CR for finite terms, there are two reduction sequences $U'\colon t'_1 \to^\star t'_3$ and $V'\colon t'_2 \to^\star t'_3$. One can show that by applying σ to each term of U' and V', we obtain two reduction sequences $U\colon t_1 \to^\star t_3$ and $V\colon t_2 \to^\star t_3$. This shows that $\mathcal{R}$ is also CR on possibly infinite terms. The converse direction is, of course, trivial since $\mathcal{T}(\Sigma, \mathcal{V}) \subseteq \mathcal{T}^\infty(\Sigma, \mathcal{V})$. Therefore, CR on finite terms is equivalent to CR on infinitary terms.

This argument can also be applied to other confluence properties such as weak and strong confluence. But also for properties that involve normal forms such as NF, UN, and $\mathsf{UN}_\to$, the above argument can still be used. The key observation is, that if a term t is a normal form, then also all consistent abstractions of t, as they are constructed in the proof sketched above, are normal forms. Moreover, the argument can also be applied to the properties of consistency and consistency w.r.t. reduction.

However, we can certainly not use this argument for the property of *ground confluence* (GCR), i.e. confluence of reductions on $\mathcal{T}(\Sigma)$. In fact, GCR depends on whether finite or potentially infinite terms are considered:

Example 5.1.2
Let $\mathcal{R}$ be the TRS over the signature $\Sigma = \{f, g, a\}$ consisting of the rules

$$\rho_1\colon f(x) \to x$$
$$\rho_2\colon g(x) \to x$$
$$\rho_3\colon g(x) \to a$$

$\mathcal{R}$ is GCR on $\mathcal{T}(\Sigma)$ since every term in $\mathcal{T}(\Sigma)$ is reducible to a. Also all infinite terms except f^ω are reducible to a. In particular, we can reduce the term $g(f^\omega)$ to a and to f^ω in a single step using rule ρ_3 and ρ_2, respectively. However, a is a normal form and f^ω can only be reduced to itself. Hence, $\mathcal{R}$ is not GCR for possibly infinite terms.

5.2 MRS vs. PRS Model of Infinitary Term Rewriting

As seen in Chapter 3 there are several ways of defining transfinite reductions for a term rewriting system. We have considered two dimensions of choices. At first, one can choose between the MRS semantics and the PRS semantics for term rewriting systems. And secondly, one can choose between a weak and a strong variant of transfinite rewriting. In order to indicate whether weakly convergent or strongly convergent reductions are considered, we use the arrows $\hookrightarrow$ and $\twoheadrightarrow$, respectively. Since, for term rewriting systems, we also have the choice of either considering its MRS or its PRS model of transfinite reductions, we need a notation to indicate which one we choose. To this end, we use the superscripts ^{m} and ^{p} to indicate that the MRS semantics respectively the PRS semantics is considered. That is, we write $\hookrightarrow^m$ and $\twoheadrightarrow^m$ for reductions w.r.t. the MRS semantics (also called *MRS reductions*) and $\hookrightarrow^p$ and $\twoheadrightarrow^p$ for reductions w.r.t. the PRS semantics (also called *PRS reductions*). However, if it is clear from the context which semantics is meant, the superscript is dropped. Particularly, in Section 5.3 and Section 5.4 we implicitly assume MRS reductions, whereas in Section 5.5 PRS reductions are assumed.

The purpose of this section is to establish that the PRS semantics always extends the MRS semantics of an ITRS in the sense of Definition 3.3.10. As a consequence, we will obtain that the MRS model of transfinite reductions in ITRSs yields the same reductions as the PRS model restricted to the set $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ of total terms.

Proposition 5.2.1 (partial order extends metric on terms)
$(\mathcal{T}^\infty(\Sigma_\bot, \mathcal{V}), \leq_\bot)$ *extends* $(\mathcal{T}^\infty(\Sigma, \mathcal{V}), \mathbf{d})$.

Proof. This is a special case of Proposition 4.5.22 which states that this holds also for term graphs. □

Proposition 5.2.2 (PRS semantics of ITRSs extends MRS semantics)
For each ITRS $\mathcal{R}$, its induced PRS $\mathcal{P}_\mathcal{R}$ extends its induced MRS $\mathcal{M}_\mathcal{R}$.

Proof. Let $\mathcal{M}_\mathcal{R} = (A, \Phi, \mathsf{src}, \mathsf{tgt}, \mathbf{d}, \mathsf{hgt})$ and $\mathcal{P}_\mathcal{R} = (B, \Phi', \mathsf{src}', \mathsf{tgt}', \leq, \mathsf{cxt})$. In the following, we drop the subscript $\mathcal{R}$ and simply write $\mathcal{M}$ and $\mathcal{P}$. We have to show the clauses (1) – (5) of Definition 3.3.10(ii). (1) holds by Proposition 5.2.1. Clauses (2) – (4) follow immediately from Definition 2.3.24.

For the "only if" direction of (5), assume that $S = (\varphi_\iota\colon t_\iota \to_{c_\iota} t_{\iota+1})_{\iota<\lambda}$ is a total open reduction sequence strongly converging to t_λ in $\mathcal{P}$. We will prove by induction on λ that then $S\colon t_0 \twoheadrightarrow_\mathcal{M} t_\lambda$.

Since $S\colon t_0 \twoheadrightarrow_\mathcal{P} t_\lambda$ is total, it is also a reduction sequence in $\mathcal{M}$ due to (3) and (4). Moreover, by Proposition 3.2.12, we also have that $S|_{[0,\lambda')}\colon t_0 \twoheadrightarrow_\mathcal{P} t_{\lambda'}$ for each limit ordinal $\lambda' < \lambda$. Applying the induction hypothesis then yields that $S|_{[0,\lambda')}\colon t_0 \twoheadrightarrow_\mathcal{M} t_{\lambda'}$ for each limit ordinal $\lambda' < \lambda$. Consequently, by Proposition 3.1.17, it holds that $S\colon t_0 \twoheadrightarrow_\mathcal{M} \dots$. It remains to be shown that S also strongly converges to t_λ in $\mathcal{M}$. Note that $t_\lambda = \liminf_{\iota\to\lambda} c_\iota$. By definition, $c_\iota \leq_\bot t_\iota$ for all $\iota < \lambda$. Hence, also $\liminf_{\iota\to\lambda} c_\iota \leq_\bot \liminf_{\iota\to\lambda} t_\iota$. Since $t_\lambda = \liminf_{\iota\to\lambda} c_\iota$ is maximal w.r.t. $\leq_\bot$, this implies $\liminf_{\iota\to\lambda} c_\iota = \liminf_{\iota\to\lambda} t_\iota$. By (1), we then have $\liminf_{\iota\to\lambda} t_\iota = \lim_{\iota\to\lambda} t_\iota$ as all t_ι and also their limit inferior are maximal. Hence, S weakly converges to t_λ in $\mathcal{M}$. In order to show that S also strongly converges in $\mathcal{M}$, suppose that it does not. According to Proposition 5.4.2, which is proven later, this means that there is a position π such that, for each $\alpha < \lambda$, there is some $\alpha \leq \beta < \lambda$ such that the step φ_β takes place at position π. By Lemma 5.5.4, which is also proven later, this contradicts the fact that t_λ is a total term.

For the converse direction of (5), assume that $S = (\varphi_\iota \colon t_\iota \to t_{\iota+1})_{\iota<\lambda}$ is an open reduction sequence strongly converging to t_λ in $\mathcal{M}$. By performing a transfinite induction proof on λ, we will show that then $S\colon t_0 \twoheadrightarrow_{\mathcal{P}} t_\lambda$ is total.

By (1), (2) and (3), S is a total reduction sequence in $\mathcal{P}$. Moreover, by applying Proposition 3.2.12, we obtain that $S|_{[0,\lambda')}\colon t_0 \twoheadrightarrow_{\mathcal{M}} t_{\lambda'}$ for each limit ordinal $\lambda' < \lambda$. By applying the induction hypothesis, we obtain that $S|_{[0,\lambda')}\colon t_0 \twoheadrightarrow_{\mathcal{P}} t_{\lambda'}$ for each limit ordinal $\lambda' < \lambda$. Hence, $S\colon t_0 \twoheadrightarrow_{\mathcal{P}} \dots$ by Proposition 3.2.12. It remains to be shown that S strongly converges to t_λ in $\mathcal{P}$ as well. By definition, it holds that $t_\lambda = \lim_{\iota\to\lambda} t_\iota$. Additionally, by (1), we have $t_\lambda = \liminf_{\iota\to\lambda} t_\iota$. We obtain the desired result if we can show that $\liminf_{\iota\to\lambda} t_\iota = \liminf_{\iota\to\lambda} c_\iota$, where $c_\iota = \mathsf{cxt}(\varphi_\iota)$ for all $\iota < \lambda$. Let d_ι be the depth of the reduction step φ_ι for each $\iota < \lambda$. Since S is strongly convergent in $\mathcal{M}_\mathcal{R}$, the sequence $(d_\iota)_{\iota<\lambda}$ tends to infinity. By Proposition 4.5.23, this means that

$$\liminf_{\iota\to\lambda} t_\iota = \liminf_{\iota\to\lambda} t_\iota|d_\iota \quad \text{and} \quad \liminf_{\iota\to\lambda} c_\iota = \liminf_{\iota\to\lambda} c_\iota|d_\iota.$$

By definition, it holds that $t_\iota|d_\iota = c_\iota|d_\iota$ and we can conclude that

$$\liminf_{\iota\to\lambda} t_\iota = \liminf_{\iota\to\lambda} t_\iota|d_\iota = \liminf_{\iota\to\lambda} c_\iota|d_\iota = \liminf_{\iota\to\lambda} c_\iota.$$

□

Now we can apply the theory established in Section 3.3.2 in order to identify MRS reductions with total PRS reductions.

Corollary 5.2.3 (total PRS reductions = MRS reductions)
Let $\mathcal{R}$ be an ITRS. Then the following holds for reductions in $\mathcal{R}$:

(i) $S\colon s \hookrightarrow^p \dots$ *is total iff* $S\colon s \hookrightarrow^m \dots$.

(ii) $S\colon s \hookrightarrow^p t$ *is total iff* $S\colon s \hookrightarrow^m t$.

(iii) $S\colon s \twoheadrightarrow^p \dots$ *is total iff* $S\colon s \twoheadrightarrow^m \dots$.

(iv) $S\colon s \twoheadrightarrow^p t$ *is total iff* $S\colon s \twoheadrightarrow^m t$.

Proof. Follows immediately from Proposition 5.2.2 and Proposition 3.3.11. □

This relation between MRS and PRS reductions will become particularly useful when comparing PRS reductions and so-called Böhm reductions in Section 5.5.3.

5.3 Weakly Convergent MRS Reductions

The purpose of this section is to provide an overview of the properties of weakly convergent reductions of ITRSs w.r.t. the traditional MRS semantics. The MRS semantics of infinitary term rewriting, given in Definition 3.1.3, uses the usual ultrametric on terms in order to formalise the limit behaviour of transfinite reductions. In Section 3.1, the intuition behind transfinite rewriting in this setting was already presented and in conjunction with Section 3.1

and Section 3.3.1 we have seen those of its properties that are already observable in the abstract case. For weakly convergent reductions, only few interesting properties have been established. In Section 5.3.1, we present some criteria which ensure that a reduction can be performed within at most ω steps. In Section 5.3.2, some confluence properties are shown. And in Section 5.3.3, criteria are given that guarantee the existence of a corresponding strongly convergent reduction in case a weakly convergent one is available.

5.3.1 Compression and Approximation

In the following, we want to present criteria which allow to simulate transfinite reductions by reductions with only finitely many steps or at least with ω steps. Moreover, we will show that it is possible to approximate the final term of a transfinite reduction arbitrarily precise by a finite reduction.

The following theorem states that transfinite reductions can be "compressed" to a length of at most ω if the system under consideration is left-linear and top-terminating:

Theorem 5.3.1 (Compression Lemma for top-terminating systems, [DKP91])
Let $\mathcal{R}$ be a left-linear top-terminating ITRS. Then $s \hookrightarrow t$ implies $s \hookrightarrow^{\leq\omega} t$.

Proof. The theorem's phrasing differs slightly from that of Theorem 1 in [DKP91]. Here we allow the system's right-hand sides as well as the term s to be infinite. The only part of the proof in [DKP91], which is affected by this generalisation, is the argument that $s \hookrightarrow^{\omega+1} t$ implies $s \hookrightarrow^{\omega} t$. Due to the top-termination, this argument can be performed in the same way as in Lemma 5.1 from [KKSdV95a] which allows rules with infinite right-hand sides. □

Primarily, this serves as a tool for proofs on transfinite reductions as reductions of length beyond ω are harder to work with. Note that the restriction to left-linear systems is necessary:

Example 5.3.2 ([DKP91])
Consider the following top-terminating TRS $\mathcal{R}$:

$$a \to g(a), \qquad b \to g(b), \qquad f(x,x) \to c$$

$\mathcal{R}$ allows the strongly convergent $(\omega+1)$-reduction sequence

$$f(a,b) \to^2 f(g(a),g(b)) \to^2 f(g^2(a),g^2(b)) \to^2 \dots f(g^\omega,g^\omega) \to c$$

Yet, $\mathcal{R}$ does not allow a reduction $f(a,b) \hookrightarrow^{\leq\omega} c$ of length at most ω.

Also the requirement of top-termination is vital for compression:

Example 5.3.3 ([FW90])
Consider the following left-linear TRS $\mathcal{R}$:

$$a \to b, \qquad f(x,a) \to f(g(x),a)$$

$\mathcal{R}$ allows the weakly convergent $(\omega+1)$-reduction sequence

$$f(c,a) \to f(g(c),a) \to f(g^2(c),a) \to \dots f(g^\omega,a) \to f(g^\omega,b)$$

Yet, $\mathcal{R}$ does not allow a reduction $f(c,a) \hookrightarrow^{\leq\omega} f(g^\omega,b)$ of length at most ω.

Another variant of the above compression property requires (finitary) confluence instead of top-termination. But it is restricted to reductions converging to constructor terms:

Theorem 5.3.4 (Compression Lemma for confluent systems, [Luc01])
Let $\mathcal{R}$ be a left-linear (finitarily) confluent TRS, $s \in \mathcal{T}^\infty(\Sigma,\mathcal{V})$, and $t \in \mathcal{T}^\infty(\mathcal{C},\mathcal{V})$. If $s \hookrightarrow t$, then $s \hookrightarrow^{\leq\omega} t$.

The following example shows that confluence is crucial for the above theorem:

Example 5.3.5
Let $\mathcal{R}$ be the TRS:

$$f(x) \to f(g(x)), \quad f(x) \to x$$

$\mathcal{R}$ is not confluent since the term $f(x)$ can be rewritten to the two distinct normal forms x and $g(x)$. Moreover, $\mathcal{R}$ allows the weakly convergent $(\omega + 1)$-reduction sequence:

$$f(x) \to f(g(x)) \to f(g^2(x)) \to \ldots f(g^\omega) \to g^\omega$$

Yet, there is no reduction $f(x) \hookrightarrow^{\leq\omega} g^\omega$ of length at most ω.

It is well-known that weakly orthogonal TRSs are (finitarily) confluent. Hence, we can obtain the following corollary:

Corollary 5.3.6 (Compression Lemma for weakly orthogonal systems)
Let $\mathcal{R}$ be a weakly orthogonal TRS, $s \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$, and $t \in \mathcal{T}^\infty(\mathcal{C}, \mathcal{V})$. If $s \hookrightarrow t$, then $s \hookrightarrow^{\leq\omega} t$.

Proof. Follows from Theorem 2.3.31 and Theorem 5.3.4. □

The next theorem shows that it is possible to approximate the result of a transfinite reduction by a finite reduction with arbitrary precision:

Theorem 5.3.7 (finite approximation, [Luc01])
Let $\mathcal{R}$ be a left-linear TRS and $s, t \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$. If $s \hookrightarrow t$, then for each depth $d \in \mathbb{N}$, there is some $t' \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$ such that $s \to^* t'$ and t and t' coincide up do depth d, i.e. $\mathsf{sim}(t, t') > d$.

Of course, for finite terms, this means that they can be computed by a finite reduction:

Corollary 5.3.8 (finite reductions to finite terms, [Luc01])
Let $\mathcal{R}$ be a left-linear TRS, $s \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$, and $t \in \mathcal{T}(\Sigma, \mathcal{V})$. If $s \hookrightarrow t$, then $s \to^* t$.

The assumption of left-linearity is crucial for both Theorem 5.3.7 and Corollary 5.3.8. Example 5.3.2 illustrates this.

5.3.2 Confluence

Obtaining infinitary confluence for weakly convergent reductions is hard. It is, for example, not possible to establish an adequate theory of *complete developments* as we will argue in Section 5.4.2. This was also one of the most important reasons for considering strong convergence (cf. [KKSdV95a]). There are some criteria which relate weakly convergent reductions to strongly convergent reductions. We will discuss them in Section 5.3.3. With these we are able to hijack some of the confluence results for strongly convergent reductions.

However, there are some results concerning variants of infinitary confluence. This section summarises the most important ones.

The following Theorem provides criteria of a property called semi-ω-confluence (cf. [DKP91]). Its structure is depicted in Figure 5.1.

Theorem 5.3.9 (semi-ω-confluence, [DKP91])
Let $\mathcal{R}$ be a weakly ω-convergent orthogonal TRS. For any two coinitial reductions $t \to^* t_1$ and $t \hookrightarrow^{\leq\omega} t_2$, there are two reductions $t_1 \hookrightarrow^{\leq\omega} t_3$ and $t_2 \hookrightarrow^{\leq\omega} t_3$.

It is not known whether full orthogonality and ω-convergence is really necessary for the theorem. However, left-linearity is indeed needed as the following example shows:

$$
\begin{tikzcd}
& t \arrow[dl, "*"'] \arrow[dr, hook, "\leq\omega"] & \\
t_1 \arrow[dr, hook, dashed, "\leq\omega"'] & & t_2 \arrow[dl, hook', dashed, "\leq\omega"] \\
& t_3 &
\end{tikzcd}
$$

Figure 5.1: Semi-ω-confluence of Theorem 5.3.9.

Example 5.3.10 ([DKP91])
Let $\mathcal{R}$ be the TRS given by the following rules:

$$a \to g(a), \quad b \to g(b), \quad f(x,x) \to c$$

In $\mathcal{R}$ we have the reductions

$$f(f(a,b),f(a,b)) \to c \quad \text{and} \quad f(f(a,b),f(a,b)) \hookrightarrow^{\omega} f(f(a,g^{\omega}),f(g^{\omega},b)).$$

c is a normal form and every reduction $f(f(a,g^{\omega}),f(g^{\omega},b)) \hookrightarrow c$ has a length of at least $\omega + 1$.

The following two theorems are concerned with a weaker variant of $\mathsf{UN}^{\infty}_{\to}$. Instead of normal forms, constructor terms are considered:

Theorem 5.3.11 (unique constructor normal form w.r.t. ω-reductions, [Luc01])
Let $\mathcal{R}$ be a (finitarily) confluent TRS, $s \in \mathcal{T}^{\infty}(\Sigma,\mathcal{V})$, and $t,t' \in \mathcal{T}^{\infty}(\mathcal{C},\mathcal{V})$. If $s \hookrightarrow^{\leq\omega} t$ and $s \hookrightarrow^{\leq\omega} t'$, then $t = t'$.

With the additional restriction to left-linear systems, this can be generalised to arbitrarily long weakly convergent reduction sequences.

Theorem 5.3.12 (unique constructor normal form w.r.t. transfinite reductions, [Luc01])
Let $\mathcal{R}$ be a left-linear, (finitarily) confluent TRS, $s \in \mathcal{T}^{\infty}(\Sigma,\mathcal{V})$, and $t,t' \in \mathcal{T}^{\infty}(\mathcal{C},\mathcal{V})$. If $s \hookrightarrow t$ and $s \hookrightarrow t'$, then $t = t'$.

It is obvious that both Theorem 5.3.11 and Theorem 5.3.12 fail (even for finite reductions) if the restriction to confluent systems is omitted. Additionally, the following example shows that both theorems do not hold if arbitrary normal forms are considered:

Example 5.3.13 ([Luc01])
Consider the following (ground) TRS $\mathcal{R}$:

$$f(a) \to a \qquad f(a) \to f(f(a))$$

Since every ground term can be reduced to a, the above system $\mathcal{R}$ is ground confluent. The fact that $\mathcal{R}$ is ground implies that it is confluent in general. Yet, we have the two reductions $f(a) \hookrightarrow^{\omega} f^{\omega}$ and $f(a) \to a$. Both reducts are in normal form but only a is a constructor term.

From Theorem 5.3.12, we obtain the following corollary

Corollary 5.3.14 (unique constructor normal form w.r.t. transfinite reductions)
Let $\mathcal{R}$ be an weakly orthogonal TRS, $s \in \mathcal{T}^{\infty}(\Sigma,\mathcal{V})$, and $t,t' \in \mathcal{T}^{\infty}(\mathcal{C},\mathcal{V})$. If $s \hookrightarrow t$ and $s \hookrightarrow t'$, then $t = t'$.

Proof. Immediate consequence of Theorem 5.3.12 and Theorem 2.3.31. □

5.3.3 Connection to Strongly Convergent Reductions

As we will see in Section 5.4, there are far more results for strongly convergent than for weakly convergent reductions. Hence, by having criteria which ensure the strong convergence of weakly convergent reduction, or which at least assert strongly convergent reductions with the same start and end term, we are able to transfer some of the results known from strong convergent reductions to the setting of weakly convergent reductions.

The following proposition is often implied in the literature (e.g. in [Sim06]), but to the best of our knowledge no proof for this was explicitly given up to now.

Proposition 5.3.15 (strong convergence in top-terminating systems)
Let $\mathcal{R}$ be a left-linear and top-terminating ITRS. Then $s \hookrightarrow t$ implies $s \twoheadrightarrow t$.

Proof. Let $S\colon s \hookrightarrow t$ be a weakly convergent reduction sequence. By Theorem 5.3.1, there is a reduction sequence $T\colon s \hookrightarrow^{\leq\omega} t$. Let $T = (t_i \to_{\pi_i} t_{i+1})_{i<\alpha}$. If $T\colon s \hookrightarrow^{<\omega} t$, then $T\colon s \twoheadrightarrow t$ is obvious. Suppose that $T\colon s \hookrightarrow^{\omega} t$ is not strongly convergent. Then there is some depth $d \in \mathbb{N}$ such that infinitely many reduction steps in T occur at depth d. Let d^* be the minimal such depth. That is, there is some $n < \omega$ such that all reduction steps in $T|_{[n,\omega)}$ are at depth at least d^*, i.e. $|\pi_i| \geq d^*$ holds for all $n \leq i < \omega$. Of course, also $T|_{[n,\omega)}$ contains infinitely many steps at depth d^*. As all reduction steps in $T|_{[n,\omega)}$ take place at depth d^* or below, $t_i|d^* = t_j|d^*$ holds for all $n \leq i,j < \omega$. That is, all terms in $T|_{[n,\omega)}$ have the same set of positions of length d^*. Let $P^* = \{\pi \in \mathcal{P}(t_n) \mid |\pi| = d^*\}$ be this set. Since there are infinitely many steps in $T|_{[n,\omega)}$ taking place at a position in P^*, yet, P^* is finite, there has to be some position $\pi^* \in P^*$ at which infinitely many steps in $T|_{[n,\omega)}$ occur. Let T' be the reduction sequence that can be obtained from $T|_{[n,\omega)}$ by removing all reduction steps which occur at a position disjoint from π^*. Let $T' = (s_i \to_{\widehat{\pi}_i} s_{i+1})_{i<\omega}$. By construction, we have $\widehat{\pi}_i = \pi^* \cdot \pi'_i$ for some π'_i for each $i < \omega$. Hence, each s_i can be written as $s_0[s'_i]_{\pi^*}$ for an appropriate term s'_i. This gives rise to a reduction sequence $S' = (s'_i \to_{\pi'_i} s'_{i+1})_{i<\omega}$. Because infinitely many steps in T' are at position π^*, we can conclude that infinitely many steps in S' are at root position. This contradicts the assumption that $\mathcal{R}$ is top-terminating. Hence, $T\colon s \twoheadrightarrow t$. □

The requirement of top-termination is, of course, essential as Example 5.1.1 illustrates. Also left-linearity is crucial in order to ensure strong convergence. The following example illustrates this:

Example 5.3.16
Consider the TRS $\mathcal{R}$ containing the following rules:

$$\begin{aligned} \rho_1\colon &\quad a \to g(a) \\ \rho_2\colon &\quad b \to g(b) \\ \rho_3\colon &\quad f(x,x,y) \to f(a,b,g(y)) \end{aligned}$$

$\mathcal{R}$ is top-terminating. In $\mathcal{R}$ we can construct weakly convergent ω-reduction sequences of the form:

$$\begin{aligned} f(g^\omega,g^\omega,g^n(a)) \to f(a,b,g^{n+1}(a)) \to^2 f(g(a),g(b),g^{n+1}(a)) \to^2 f(g^2(a),g^2(b),g^{n+1}(a)) \\ \to^2 f(g^3(a),g^3(b),g^{n+1}(a)) \to \ \dots\ f(g^\omega,g^\omega,g^{n+1}(a)) \end{aligned}$$

By combining these reductions $f(g^\omega,g^\omega,g^n(a)) \hookrightarrow^\omega f(g^\omega,g^\omega,g^{n+1}(a))$, we obtain the weakly convergent ω^2-reduction sequence

$$f(g^\omega,g^\omega,a) \hookrightarrow^\omega f(g^\omega,g^\omega,g(a)) \hookrightarrow^\omega f(g^\omega,g^\omega,g^2(a)) \hookrightarrow^\omega \ \dots\ f(g^\omega,g^\omega,g^\omega)$$

This sequence is clearly not strongly convergent as infinitely many times rule ρ_3 is applied at the root. It is also clear that any weakly convergent reduction sequence from $f(g^\omega,g^\omega,a)$ to $f(g^\omega,g^\omega,g^\omega)$ has to apply rule ρ_3 infinitely often at the root. Hence, there is no strongly convergent reduction $f(g^\omega,g^\omega,a) \twoheadrightarrow f(g^\omega,g^\omega,g^\omega)$.

The restriction of the above proposition to top-termination is rather strict, also because of the fact that top-termination is in general undecidable. However, if we restrict ourselves to normalising reductions, we can achieve a similar behaviour:

Theorem 5.3.17 (normalising reductions, [KKSdV95a])
Let $\mathcal{R}$ be an orthogonal ITRS for which there is an upper bound on the depth of the left-hand sides of its rules. If $s \hookrightarrow t$, with t a normal form, then $s \twoheadrightarrow t$.

Both orthogonality and the boundedness of the depth of the rules' left-hand sides are vital for the above theorem:

Example 5.3.18

(i) (from [KKSdV95a]) Let $\mathcal{R}$ be the TRS consisting of the rules

$$f(g^n(c)) \to f(g^{n+1}(c)) \quad \text{for all } n \in \mathbb{N}$$

Note that $\mathcal{R}$ is orthogonal, but the depth of the left-hand sides of its rules are not bounded. The term $f(c)$ reduces to the normal form $f(g^\omega)$ by a weakly convergent reduction. However, there is no strongly convergent reduction $f(c) \twoheadrightarrow f(g^\omega)$ in $\mathcal{R}$.

(ii) Let $\mathcal{R}$ be the TRS over signature $\{f, g, a, b\}$ consisting of the rules

$$\begin{aligned} &g(x,x) \to c \\ &g(x,y) \to g(f(x), f(y)) \end{aligned}$$

$\mathcal{R}$ has only finitely many rules but is not orthogonal. From the term $g(a,b)$, we have the following weakly convergent reduction sequence to the normal form c:

$$g(a,b) \to g(f(a), f(b)) \to g(f^2(a), f^2(b)) \to g(f^3(a), f^3(b)) \to \ \dots\ g(f^\omega, f^\omega) \to c$$

Yet, there is no strongly convergent reduction sequence $g(a,b) \twoheadrightarrow c$ in $\mathcal{R}$.

When we consider constructor terms instead of general normal forms, the requirement of orthogonality and boundedness of left-hand sides can be dropped:

Theorem 5.3.19 (reductions to constructor terms, [Luc01])
Let $\mathcal{R}$ be an ITRS. If S: $s \hookrightarrow^{\leq\omega} t$, with $t \in \mathcal{T}^\infty(\mathcal{C}, \mathcal{V})$, then S: $s \twoheadrightarrow^{\leq\omega} t$, i.e. weakly convergent ω-reduction sequences ending in a constructor term are strongly convergent.

Here, the restriction to ω-reduction sequences is important as Example 5.3.18 (ii) shows. Also note that this cannot be generalised to arbitrary normal forms. Example 5.3.18 (i) illustrates this.

5.4 Strongly Convergent MRS Reductions

The discussion of weakly convergent reductions in Section 5.3 has shown that it is hard to establish useful properties in that setting. This was one of the motivations for considering strongly convergent reductions (cf. [KKSdV95a]). The purpose of this section is to present the most important results for strongly convergent reductions and also to argue why most of these results cannot be generalised to weakly convergent reductions in a satisfying way. We have already seen an example of this in the abstract setting of MRSs in Section 3.1: Strongly continuous reductions can only have countable ordinal length whereas weakly continuous reductions can be of arbitrary length. After the paradigm of strongly convergent reductions had been introduced and shown to have more advantageous properties (mostly by Kennaway et al. [KKSdV91, KKSdV95a]), most investigations of infinitary rewriting were focused on this paradigm. Therefore, the theory of strongly convergent reductions of ITRSs has become a comprehensive research topic.

In Section 5.4.1, we present criteria which allow to reduce the length of reduction sequences to at most ω. The sections 5.4.2 to 5.4.5 are concerned with confluence properties. Finally, in Section 5.4.6 and Section 5.4.7 we briefly discuss strategies and termination properties, respectively.

But before we begin with the in-depth discussion of strongly convergent reductions and their properties, let us have a look at the fundamental difference between strongly and weakly convergent reductions. By definition, strong convergence additionally requires that during the reduction the rewrite rules are applied at increasingly deeper positions. The following lemma shows that this is equivalent to the condition that there is an upper bound on the steps that occur at a particular position:

Lemma 5.4.1 (strong convergence)
Let $\mathcal{R}$ be an ITRS, and $S = (t_\iota \to_{\pi_\iota} t_{\iota+1})_{\iota<\alpha}$ a strongly continuous open reduction sequence in $\mathcal{R}$. S is strongly convergent iff, for each position π, there is an ordinal $\beta < \alpha$ such that $\pi_\iota \neq \pi$ for all $\beta \leq \iota < \alpha$.

Proof. The "only if" direction is easy: If S is strongly convergent, then $|\pi_\iota|$ has to tend to infinity as ι approaches α. Hence, the positions π_ι have to become longer and longer. In particular, this implies that, for each position π, there is an upper bound for the indices of the steps that occur at π.

For the converse direction, suppose that, for each position π, there is an upper bound for the indices of the steps in S that occur at π. By Proposition 3.1.14, we have to show that $|\pi_\iota|$ tends to infinity as ι approaches α. Assume that this is not the case. That is, there is some depth $d \in \mathbb{N}$ such that there is no upper bound on the indices of reduction steps taking place at depth d. Let d^* be the minimal such depth. That is, there is some $\beta < \alpha$ such that all reduction steps in $S|_{[\beta,\alpha)}$ are at depth at least d^*, i.e. $|\pi_\iota| \geq d^*$ holds for all $\beta \leq \iota < \alpha$. Of course, also in $S|_{[\beta,\alpha)}$ the indices of steps at depth d^* are not bounded from above. As all reduction steps in $S|_{[\beta,\alpha)}$ take place at depth d^* or below, $t_\iota|d^* = t_{\iota'}|d^*$ holds for all $\beta \leq \iota, \iota' < \alpha$. That is, all terms in $S|_{[\beta,\alpha)}$ have the same set of positions of length d^*. Let $P^* = \{\pi \in \mathcal{P}(t_n) \mid |\pi| = d^*\}$ be this set. Since there is no upper bound on the idices of steps in $S|_{[\beta,\alpha)}$ taking place at a position in P^*, yet, P^* is finite, there has to be some position $\pi^* \in P^*$ for which there is also no such upper bound. This contradicts the assumption that there is always such an upper bound. ☐

From this, we easily obtain the following proposition:

Proposition 5.4.2 (strong continuity and convergence)
Let $\mathcal{R}$ be an ITRS, and $S = (t_\iota \to_{\pi_\iota} t_{\iota+1})_{\iota<\alpha}$ a weakly continuous reduction sequence in $\mathcal{R}$. Then the following holds:

(i) S is strongly continuous iff, for each limit ordinal $\lambda < \alpha$ and for each position π, there is an ordinal $\beta < \lambda$ such that $\pi_\iota \neq \pi$ for all $\beta \leq \iota < \lambda$.

(ii) S is strongly convergent iff, for each limit ordinal $\lambda \leq \alpha$ and for each position π, there is an ordinal $\beta < \lambda$ such that $\pi_\iota \neq \pi$ for all $\beta \leq \iota < \lambda$.

Proof. This follows immediately from Proposition 3.1.17 and Lemma 5.4.1. ☐

5.4.1 Compression and Approximation

We have seen that under certain circumstances weakly convergent reduction sequences can be "compressed" to length at most ω, and that the final term of a transfinite reduction can be approximated arbitrarily well by a finite reduction. These results can be strengthened for strongly convergent reductions.

Unlike for weakly convergent reductions, the Compression Lemma does not require top-termination in the setting of strong convergence:

Theorem 5.4.3 (Compression Lemma, [KKSdV95a])
For a left-linear ITRS, $S\colon s \twoheadrightarrow t$ *implies* $T\colon s \twoheadrightarrow^{\leq\omega} t$.

With this we can also generalise Theorem 5.3.7 which was restricted to TRSs. It is not know whether this restriction is essential. However, for strongly convergent reductions, it is not:

Theorem 5.4.4 (finitary approximation)
Let $\mathcal{R}$ *be a left-linear ITRS and* $s \twoheadrightarrow t$. *Then, for each depth* $d \in \mathbb{N}$, *there is a finite reduction* $s \to^{*} t'$ *such that* t *and* t' *coincide up to depth* d, *i.e.* $\mathsf{sim}(t,t') > d$.

Proof. By Theorem 5.4.3, there is a reduction $S\colon s \twoheadrightarrow^{\leq\omega} t$. If S is of finite length, then we are done. If $S\colon s \twoheadrightarrow^{\omega} t$, then, by strong convergence, there is some $n < \omega$ such that all reductions steps in S after n take place at a depth greater than d. Consider $S|_{[0,n)}\colon s \to^{*} t'$. It is clear that t and t' coincide up to depth d. □

Again, for finite terms, this means that they can be computed by a finite reduction:

Corollary 5.4.5 (finite reductions to finite terms)
Let $\mathcal{R}$ *be a left-linear ITRS,* $s \in \mathcal{T}^{\infty}(\Sigma,\mathcal{V})$, *and* $t \in \mathcal{T}(\Sigma,\mathcal{V})$. *If* $s \twoheadrightarrow t$, *then* $s \to^{*} t$.

The assumption of left linearity is crucial for both Theorem 5.3.7 and Corollary 5.3.8. The following example illustrates this:

Example 5.4.6
Consider the following TRS [DKP91]:

$$a \to g(a) \qquad b \to g(b) \qquad f(x,x) \to c$$

Then we have the transfinite reduction $f(a,b) \twoheadrightarrow^{\omega} f(g^{\omega},g^{\omega}) \to c$. Yet, we do not have $f(a,b) \to^{*} c$.

5.4.2 Complete Developments

There are several methods to prove (finitary) confluence for orthogonal TRSs. One of them employs the concept of *complete developments*. This is a classic notion from the studies of λ-calculus and combinatory logic first used by Church and Rosser [CR36]. Intuitively speaking, complete developments are reduction sequences which contract an entire set of redexes in a term. In order to formalise this idea, one needs a formal means to track redexes in a reduction. The following definition of descendants serves this purpose:

Definition 5.4.7 (residuals/descendants, [KKSdV95a])
Let $S\colon t_0 \twoheadrightarrow^{\alpha} t_\alpha$ and U a set of occurrences in t_0. The *descendants* of U by S, denoted $U/\!/S$, is the set of occurrences in t_α inductively defined as follows:

(a) If $\alpha = 0$, then $U/\!/S = U$.

(b) If $\alpha = 1$, let $S\colon t_0 \to_{\pi,\rho} t_1$, where $\rho\colon l \to r$. Take any $u \in U$ and define the set R_u as follows: If $\pi \nleq u$, then $R_u = \{u\}$. If u is in the pattern of the redex of ρ, then $R_u = \emptyset$. Otherwise, i.e. if $u = \pi \cdot w \cdot x$, with $l|_w \in \mathcal{V}$, then $R_u = \{\pi \cdot w' \cdot x \mid r|_{w'} = l|_w\}$. Define $U/\!/S = \bigcup_{u \in U} R_u$.

(c) If $\alpha = \alpha' + 1$, then $U/\!/S = (U/\!/S|_{[0,\alpha')})/\!/S|_{[\alpha',\alpha)}$.

(d) If α is a limit ordinal, then $U/\!/S = \liminf_{\iota \to \alpha} U/\!/S|_{[0,\iota)}$.
That is, $\quad u \in U/\!/S \quad$ iff $\quad \exists\beta < \alpha \forall \beta < \iota < \alpha\colon u \in U/\!/S|_{[0,\iota)}$

If, in particular, U is a set of redex occurrences in t_0, then $U/\!/S$ is also called the set of *residuals* of U by S. Moreover, by abuse of notation, we write $u/\!/S$ instead of $\{u\}/\!/S$.

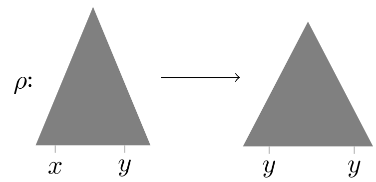

(a) A rewrite rule.

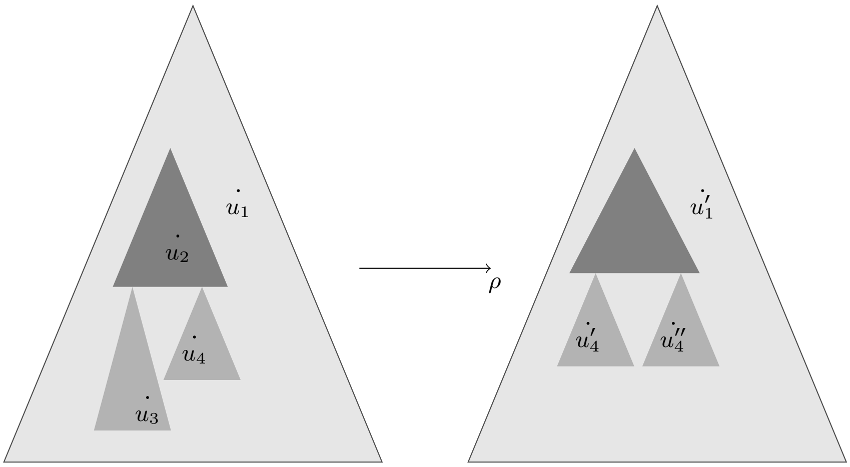

(b) Application of a rewrite rule.

Figure 5.2: Descendants by a single step.

The definition of descendants given above is a straightforward generalisation of the corresponding concept known from finite reductions. Item (a) needs no explanation. Item (b) is a simple case distinction of what can happen to occurrences during a single reduction step. The schematic example in Figure 5.2 illustrates its intuition. Figure 5.2a depicts a left-linear rewrite rule indicating all variable occurrences of its left- and right-hand side. Figure 5.2b shows the result of applying the rewrite rule to a term. In the initial term, four occurrences are singled out. One “outside” the redex, viz. u_1, and three “inside” the redex, viz. u_2, u_3 and u_4. The descendants of these occurrences are indicated in the resulting term as u'_1, u'_4 and u''_4. The occurrence outside the redex corresponds to the case that $\pi \nleq u_1$. The occurrence is preserved unaltered as u'_1, i.e. $u_1 = u'_1$. The occurrence in the pattern of the redex, viz. u_2, vanishes, i.e. it has no descendants, as demanded by the definition. Lastly, the positions u_3 and u_4 are in the variable part of the redex. Their propagation as descendants depends on how often and where the corresponding variable occurs on the right-hand side of the rewrite rule. As the variable x does not occur on the right-hand side of the rule, the occurrence u_3 has no descendants. On the other hand, y even occurs twice on the right-hand side. Hence, u_4 has two descendants, viz. u'_4 and u''_4. Item (c) is also straightforward: It simply states that descendants are transitively propagated by each step.

This is all well-known from finite reductions. The interesting part of the definition is (d): It asserts for the limit case that an occurrence is a descendant iff it becomes a stable descendant by all prefixes, i.e. it is a descendant by all prefixes from one point onwards. This also corresponds to the notion of weak convergence which essentially states that an open reduction sequence converges to a term whose every symbol at some position becomes stable at that position in the reductions before.

The concept of descendants is defined for strongly convergent reductions. One can easily see, however, that it is applicable to weakly convergent reductions, too. Yet, as we will see later, descendants are only meaningful for strongly convergent reductions. Along the way we will argue why this notion of descendants is not appropriate for weakly convergent reductions. As a matter of fact, the situation seems to be even worse than that: One can argue, as, for example, Simonsen [Sim04] does, that also any other notion of descendants either fails to be useful as a tool for proving important properties like the Infinitary Strip Lemma (cf. Proposition 5.4.23) or has peculiar properties itself which are counterintuitive to the notion of descendants.

With the concept of residuals it is possible to formalise the intuitive notion of complete developments. A complete development of a set of redexes is a reduction that contracts only redex occurrences which are residuals of the original set of redexes and stops in a term in which no residuals are left.

Definition 5.4.8 ((complete) development, [KKSdV95a])
Let $\mathcal{R}$ be an ITRS, s a term in $\mathcal{R}$ and U a set of pairwise non-conflicting redex occurrences in s.

(i) A *development* of U in s is a strongly convergent reduction $S\colon s \twoheadrightarrow^{\alpha}_{\mathcal{R}} t$ in which each reduction step $\varphi_\iota\colon t_\iota \to_{\pi_\iota} t_{\iota+1}$ contracts a redex at $\pi_\iota \in U//S|_{[0,\iota)}$ for $\iota < \alpha$.

(ii) A development $S\colon s \twoheadrightarrow t$ of U in s is called *complete*, denoted $S\colon s \twoheadrightarrow_U t$, if $U//S = \emptyset$.

An essential property that is needed to make the above definition meaningful is that the considered redex occurrences are independent from each other. The restriction to non-conflicting redex occurrences guarantees this. Moreover, in a left-linear system, the descendants of a set of non-conflicting redex occurrences is again a set of non-conflicting redex occurrences:

Fact 5.4.9 (non-conflicting residuals)
Let $\mathcal{R}$ be a left-linear ITRS, s a term in $\mathcal{R}$, U a set of pairwise non-conflicting redex occurrences in s, and $S\colon s \twoheadrightarrow_U t$ a development of U in s. Then also $U//S$ is a set of pairwise non-conflicting redex occurrences.

Usually, (almost) orthogonal systems are considered in the setting of complete developments. Recall, that (almost) orthogonal systems do not allow conflicting redex occurrences. Moreover, for technical reasons, also disjoint redex occurrences play an important rôle:

Proposition 5.4.10 ((disjoint) residuals, [KKSdV95a])
Let $\mathcal{R}$ be an almost orthogonal ITRS, $S\colon s \twoheadrightarrow_{\mathcal{R}} t$, and U a set of redex occurrences in s. Then the following holds:

(i) $U//S$ is a set of redex occurrences in t.

(ii) If the occurrences in U are pairwise disjoint, then so are the occurrences in $U//S$.

The following two propositions reveal more of the intuition of descendants.

Proposition 5.4.11 (pointwise definition of descendants)
Let $\mathcal{R}$ be an ITRS, $S\colon s \twoheadrightarrow_{\mathcal{R}} t$ and $U \subseteq \mathcal{P}(s)$. Then it holds that $U//S = \bigcup_{u\in U} u//S$.

Proof. Straightforward induction on the length of S. □

Proposition 5.4.12 (uniqueness of descendants)
Let $\mathcal{R}$ be a left-linear ITRS, $S\colon s \twoheadrightarrow_{\mathcal{R}} t$ and $U, V \subseteq \mathcal{P}(s)$. If $U \cap V = \emptyset$, then $U//S \cap V//S = \emptyset$.

Proof. Straightforward induction on the length of S. □

Remark 5.4.13. Particularly, the two propositions above imply that in left-linear systems each descendant $u' \in U//S$ of a set of occurrences is the descendant of a uniquely determined occurrence $u \in U$, i.e. $u' \in u//S$ for exactly one $u \in U$. This occurrence u is also called the *ancestor* of u' by S.

However, in general, this does only hold for strongly convergent reductions as the following example illustrates:

Example 5.4.14 ([Sim04])
Consider the TRS $\mathcal{R}$ with the single rule $f(x,y) \to f(y,x)$ and the term $f(a,a)$. Take the set $U = \{0,1\}$ of occurrences in $f(a,a)$ and the weakly convergent reduction:

$$S\colon f(\overline{a},\widehat{a}) \to f(\widehat{a},\overline{a}) \to f(\overline{a},\widehat{a}) \to \;\dots\; f(a,a)$$

The residuals of 0 and 1 are labelled with $^{-}$ and $\frown$, respectively, in order to make the occurring phenomenon clearly visible. One can see that $U /\!/ S = U$. On the other hand, we have both $0/\!/S = \varnothing$ and $1/\!/S = \varnothing$. This violates Proposition 5.4.11. The problem is that, individually, the two occurrences do not become stable during the reduction. However, if they are put into the same set, and are thus made indistinguishable, they seem to be stable. As a consequence, neither of the two occurrences 0 and 1 in the final term $f(a,a)$ has an ancestor.

The key property for the finitary version of complete developments is the so called Finite Developments Theorem (cf. [Ter03]) which states that all developments are finite. In the finitary setting, this means in particular that complete developments do always exists. In the infinitary setting, complete developments may be of infinite length, of course, but they are required to be strongly convergent. This is not always the case as the following example shows:

Example 5.4.15
Let $\mathcal{R}$ be an orthogonal TRS with the single rule $f(x) \to x$. Consider the term f^ω and the set $U = \{\varepsilon, 0, 0\cdot 0, \dots\}$ of all redex occurrences in f^ω. There is no complete development of U in f^ω as, for all strongly convergent reduction sequences S starting in f^ω, we have $S\colon f^\omega \twoheadrightarrow f^\omega$ and, therefore, $U/\!/S = U$.

This means that, unlike in the finitary setting, complete developments are not guaranteed to exist for infinite terms. The characteristic issue that occurs in the above example is that the set of redex occurrences U contains an infinite collapsing tower. Recall that a collapsing tower is a sequence of nested collapsing redex occurrences where each redex occurrence (except the topmost) is located at the collapsing position of the overlying collapsing redex. We say that a set of occurrences U *contains an infinite collapsing tower*, if there is an infinite collapsing tower consisting of elements in U only. If such infinite collapsing towers are excluded, complete developments do exist:

Proposition 5.4.16 (complete development, [KKSdV95a])
In an orthogonal ITRS, every set of redex occurrences not containing an infinite collapsing tower has a complete development.

This restriction of the existence of complete developments will turn out to have consequences for the infinitary confluence result that can be established for strongly convergent reductions.

Another important property of complete developments on the road to confluence is that the result of a complete development is uniquely determined by the set of redex occurrences under consideration.

Proposition 5.4.17 (descendants of complete developments, [KKSdV95a, Ter03])
Let $\mathcal{R}$ be an orthogonal ITRS, t a term in $\mathcal{R}$ and U a set of redex occurrences in t not containing an infinite collapsing tower. Then the following holds:

(i) Each complete development of U in t ends in the same term.

(ii) For each set $V \subseteq \mathcal{P}(t)$ and two complete developments S and T of U in t, it holds that $V/\!/S = V/\!/T$.

$$\begin{tikzcd}
& t \arrow[dl, two heads, "U"'] \arrow[dr, two heads, "V"] & \\
t_1 \arrow[dr, dashed, two heads, "V/\!/U"'] & & t_2 \arrow[dl, dashed, two heads, "U/\!/V"] \\
& t_3 &
\end{tikzcd}$$

Figure 5.3: Diamond property of complete developments according to Corollary 5.4.19.

Notation 5.4.18. Let $\mathcal{R}$ be an orthogonal ITRS, t a term in $\mathcal{R}$, U a set of occurrences in t, and V a set of redex occurrences in t not containing an infinite collapsing tower. We write $U/\!/V$ for the descendants $U/\!/S$ of U by some complete development S of V. Proposition 5.4.17 shows that $U/\!/V$ is well-defined.

As an immediate corollary, we obtain that complete developments have the diamond property as depicted in Figure 5.3:

Corollary 5.4.19 (diamond property of complete developments, [KKSdV95a])
Let $\mathcal{R}$ be an orthogonal ITRS and $t \twoheadrightarrow_U t_1$ and $t \twoheadrightarrow_V t_2$ be two complete developments of U respectively V in t. If U and V do not contain an infinite collapsing tower, then t_1 and t_2 are joinable by complete developments $t_1 \twoheadrightarrow_{V/\!/U} t_3$ and $t_2 \twoheadrightarrow_{U/\!/V} t_3$.

Note that pairwise disjointness of redex occurrences is a sufficient condition for the absence of infinite collapsing tower. Hence, the properties that require the absence of infinite collapsing towers hold in particular for sets of disjoint redex occurrence. Disjointness, however, has the advantage that it is preserved by strongly convergent reductions (cf. Proposition 5.4.10). This is not true for the absence of infinite collapsing towers as the following example shows.

Example 5.4.20
Let $\mathcal{R}$ be the TRS consisting of the rules

$$\rho_1 \colon f(x) \to x, \quad \rho_2 \colon g(x) \to x.$$

Consider the term $s = f(g(f(g(\dots$ and the set $U = \{\varepsilon, 0^2, 0^4, \dots\}$ of all ρ_1-redex occurrences in s. Note that U does not contain an infinite collapsing tower. Next, let $S \colon s \twoheadrightarrow t$ be a reduction that contracts all ρ_2-redexes, i.e. S is a complete development of the set $U = \{0, 0^3, 0^5, \dots\}$ of all ρ_2-redex occurrences in s. Hence, we have $t = f^\omega$ and $U/\!/S = \{\varepsilon, 0, 0^2, \dots\}$. That is, $U/\!/S$ contains the infinite collapsing tower $\varepsilon, 0, 0^2, \dots$.

5.4.3 Tiling Diagrams and Projections

Having a generalised notion of complete developments and residuals for transfinite reductions, one can also generalise the notion of projection to this setting. However, in contrast to the finitary setting projections do not need to exists even for orthogonal systems. The following definition formalises the concept of projections:

Definition 5.4.21 (tiling diagram, projection, [Ter03])
A *tiling diagram* for two strongly convergent reduction sequences $V \colon t_{0,0} \twoheadrightarrow^{\alpha} t_{\alpha,0}$ and $H \colon t_{0,0} \twoheadrightarrow^{\beta} t_{0,\beta}$ consists of a rectangular arrangement of strongly convergent reduction sequences as shown in Figure 5.4, subject to the following conditions:

$$
\begin{tikzcd}
 & & & & H & & & & \\
 & t_{0,0} \arrow[r] \arrow[d] & t_{0,1} \arrow[d, two heads] & \cdots\cdots & t_{0,\delta} \arrow[r] \arrow[d, two heads] & t_{0,\delta+1} \arrow[d, two heads] & \cdots\cdots & t_{0,\beta} \arrow[d, two heads] & \\
 & t_{1,0} \arrow[r, two heads] & t_{1,1} & \cdots\cdots & t_{1,\delta} \arrow[r, two heads] & t_{1,\delta+1} & \cdots\cdots & t_{1,\beta} & \\
 & \vdots & \vdots & & \vdots & \vdots & & \vdots & \\
V & t_{\gamma,0} \arrow[r, two heads] \arrow[d] & t_{\gamma,1} \arrow[d, two heads] & \cdots\cdots & t_{\gamma,\delta} \arrow[r, two heads, "H_{\gamma,\delta}"] \arrow[d, two heads, "V_{\gamma,\delta}"'] & t_{\gamma,\delta+1} \arrow[d, two heads] & \cdots\cdots & t_{\gamma,\beta} \arrow[d, two heads] & V/H \\
 & t_{\gamma+1,0} \arrow[r, two heads] & t_{\gamma+1,1} & \cdots & t_{\gamma+1,\delta} \arrow[r, two heads] & t_{\gamma+1,\delta+1} & \cdots & t_{\gamma+1,\beta} & \\
 & \vdots & \vdots & & \vdots & \vdots & & \vdots & \\
 & t_{\alpha,0} \arrow[r, two heads] & t_{\alpha,1} & \cdots\cdots & t_{\alpha,\delta} \arrow[r, two heads] & t_{\alpha,\delta+1} & \cdots\cdots & t_{\alpha,\beta} & \\
 & & & & H/V & & & &
\end{tikzcd}
$$

Figure 5.4: Tiling diagram.

(1) Each component reduction $H_{\gamma,\delta}\colon t_{\gamma,\delta} \twoheadrightarrow t_{\gamma,\delta+1}$ is a complete development of the set $\mathcal{H}_{\gamma,\delta}$ of redex occurrences in $t_{\gamma,\delta}$. Analogously, each $V_{\gamma,\delta}\colon t_{\gamma,\delta} \twoheadrightarrow t_{\gamma+1,\delta}$ is a complete development of the set $\mathcal{V}_{\gamma,\delta}$ of redex occurrences in $t_{\gamma,\delta}$.

(2) $\mathcal{H}_{0,\delta}$ is the singleton set containing the redex occurrence contracted in the δ-th step of H, and $\mathcal{V}_{\gamma,0}$ is the singleton set containing the redex occurrence contracted in the γ-th step of V.

(3) Let $H_{\gamma,[\delta,\delta')} = \prod_{\delta\leq\iota<\delta'} H_{\gamma,\iota}$ and $V_{[\gamma,\gamma'),\delta} = \prod_{\gamma\leq\iota<\gamma'} V_{\iota,\delta}$. $H_{\gamma,[\delta,\delta')}$ and $V_{[\gamma,\gamma'),\delta}$ have to be strongly convergent reduction sequences, and it has to hold that $\mathcal{H}_{\gamma,\delta} = \mathcal{H}_{0,\delta} /\!/ V_{[0,\gamma),\delta}$ and $\mathcal{V}_{\gamma,\delta} = \mathcal{V}_{\gamma,0} /\!/ H_{\gamma,[0,\delta)}$.

If H and V have a tiling diagram, then the projection of H by V, denoted H/V, and the projection of V by H, denoted V/H, are defined to be the strongly convergent reduction sequences along the bottom and right edges of the diagram, respectively.

The reason that a tiling diagram and, thus, the corresponding projections might not exists is the limit construction. The concatenation of infinitely many reduction steps might not yield a strongly convergent reduction. On the other hand, a single tile in the diagram can always be constructed due to Corollary 5.4.19.

The general existence of projections would, of course, imply infinitary confluence. Unfortunately, this property is not enjoyed by all orthogonal systems:

Example 5.4.22 ([KKSdV95a])
Let $\mathcal{R}$ be the TRS given by the rules

$$f(x) \to x, \quad g(x) \to x, \quad c \to f(g(x)).$$

$\mathcal{R}$ admits the strongly convergent reduction sequences

$$V\colon c \to f(g(c)) \to f(c) \to f(f(g(c))) \to f(f(c)) \to \ \dots\ f^\omega$$

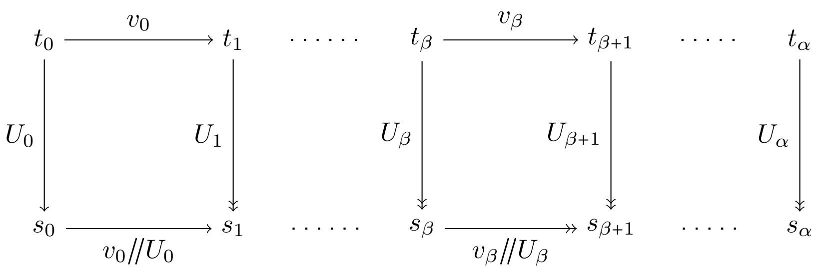

Figure 5.5: The Infinitary Strip Lemma.

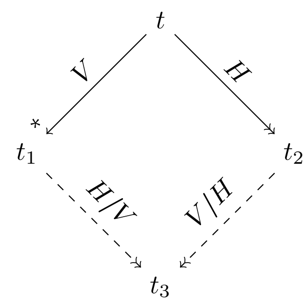

Figure 5.6: Semi-infinitary confluence according to Corollary 5.4.24.

and

$$V\colon c \to f(g(c)) \to g(c) \to g(f(g(c))) \to g(g(c)) \to \ \dots \ g^\omega$$

However, both f^ω and g^ω are only reducible to themselves and are, thus, not joinable. In particular, H and V do not have a tiling diagram.

Yet, for some restricted cases, a tiling diagram can be constructed. If at most one of the two reductions is infinite, then a tiling diagram can be constructed. Even if at least one reduction is a complete development of a set of disjoint redex occurrences, something similar to a tiling diagram can be constructed:

Proposition 5.4.23 (Infinitary Strip Lemma, [KKSdV95a])
Let $\mathcal{R}$ be an orthogonal ITRS, $S\colon t_0 \twoheadrightarrow^\alpha t_\alpha$ a strongly convergent reduction, and $t_0 \twoheadrightarrow_U s_0$ a complete development of a set U of disjoint redex occurrences. Then there is a complete development $t_\alpha \twoheadrightarrow_{U/\!/S} s_\alpha$ and a reduction $s_0 \twoheadrightarrow s_\alpha$.

The constellation that Proposition 5.4.23 describes is illustrated in Figure 5.5, in which U_ι denotes the residuals $U/\!/T|_{[0,\iota)}$. The diagram resembles a tiling diagram. The only difference is, that the reduction along the left edge is not a single step but a complete development.

Yet, if U is a singleton set, this is, in fact, a tiling diagram for a single step reduction V. Of course, by iterating this, we obtain a tiling diagram for each finite reduction V:

Corollary 5.4.24 (semi-infinitary confluence)
Let $\mathcal{R}$ be an orthogonal ITRS, $H\colon t \twoheadrightarrow t_2$, and $V\colon t \to^* t_1$. Then the projections $V/H\colon t_2 \twoheadrightarrow t_3$ and $H/V\colon t_1 \twoheadrightarrow t_3$ exist.

Proof. This can be shown by an induction on the length of V. If V is empty, then the statement trivially holds. The induction step follows from Proposition 5.4.23. $\square$

Figure 5.6 summarises the statement of the corollary. It establishes an asymmetric variant of the infinitary confluence property.

It should be pointed out that both Proposition 5.4.23 and Corollary 5.4.24 do not hold in general if weakly convergent reductions are considered:

Example 5.4.25 ([Sim04])
Let $\mathcal{R}$ be the non-collapsing orthogonal TRS given by the rules:

$$\begin{array}{lll}\rho_1\colon & a \to b & \\ \rho_2\colon & f(g^k(c), x, y) \to f(g^{k+1}(c), y, y) & \text{for all even } k \in \mathbb{N} \\ \rho_3\colon & f(g^k(c), x, y) \to f(g^{k+1}(c), a, y) & \text{for all odd } k \in \mathbb{N}\end{array}$$

We obtain the reduction $f(c,a,a) \to f(c,a,b)$ by applying rule ρ_1, and the reduction $f(c,a,a) \hookrightarrow^\omega f(g^\omega,a,a)$ by alternately applying rules ρ_2 and ρ_3 at the root. Yet, $f(c,a,b)$ and $f(g^\omega,a,a)$ are not joinable by weakly convergent reductions. A common reduct of the two terms must be of the form $f(g^\omega,s,t)$. As argued in [Sim04], there is no weakly convergent reduction from $f(c,a,b)$ to such a term. Intuitively, the problem is that such a reduction has to alternately apply rules ρ_2 and ρ_3 at the root which would cause the second argument of the f symbol to "flicker" between a and b. Hence, it cannot be weakly convergent.

5.4.4 Confluence

We have seen in Example 5.4.22 that not all orthogonal systems enjoy the infinitary confluence property. The characteristic trait of the counterexample is the availability of collapsing rules. Example 5.4.22 employs two collapsing rules. It is, however, possible to conceive a counterexample with only a single collapsing rule:

Example 5.4.26 ([KKSdV95a])
Let $\mathcal{R}$ be the orthogonal TRS consisting of the rules

$$f(x,y) \to y, \quad c \to f(a, f(b,c)).$$

In $\mathcal{R}$ the term c is reducible to both $f(a,c)$ and $f(b,c)$ by the reductions

$$\underline{c} \to f(a, \underline{f(b,c)}) \to f(a,c) \qquad \text{resp.} \qquad \underline{c} \to \underline{f(a, f(b,c))} \to f(b,c)$$

The underlinings indicate the redex that is contracted in each step. Hence, $\mathcal{R}$ admits the strongly convergent reduction sequences

$$c \to^2 f(a,c) \to^2 f(a,f(a,c)) \to^2 f(a,f(a,f(a,c))) \to \dots f(a,f(a,f(a,\dots)))$$

and

$$c \to^2 f(b,c) \to^2 f(b,f(b,c)) \to^2 f(b,f(b,f(b,c))) \to \dots f(b,f(b,f(b,\dots)))$$

Yet, there are no strongly convergent reductions which join the terms $f(a,f(a,f(a,\dots)))$ and $f(b,f(b,f(b,\dots)))$.

Basically, the rule $f(x,y) \to y$ simulates the two collapsing rules $f(x) \to x$ and $g(x) \to x$ from Example 5.4.22. Applied to a term of the form $f(a,f(b,t))$ it allows to either delete the constant a or the constant b. Similarly, the two collapsing rules of Example 5.4.22 allowed to either delete the function symbol f or the function symbol g in a term of the form $f(g(t))$.

As a matter of fact, Example 5.4.22 and Example 5.4.26 are characteristic for infinitarily non-confluent behaviour of orthogonal systems. If none of these characteristics are present in an orthogonal system, then infinitary confluence holds. The following notion of almost non-collapsingness summarises this criterion.

$$\begin{tikzcd}
 & t \arrow[dl, two heads] \arrow[dr, two heads] & \\
t_1 \arrow[dr, dashed, two heads] & & t_2 \arrow[dl, dashed, two heads] \\
 & t_3 \equiv_{hc} t'_3 &
\end{tikzcd}$$

Figure 5.7: Infinitary confluence modulo hypercollapsing terms.

Definition 5.4.27 (almost non-collapsing, [KKSdV95a])
An ITRS $\mathcal{R}$ is called *almost non-collapsing* if it contains at most one collapsing rule, which, if present, is only allowed to contain a single variable.

One can easily see that the two counterexamples for infinitary confluence are not almost non-collapsing. The TRS from Example 5.4.22 contains two collapsing rules. The system from Example 5.4.26 has only a single collapsing rule, but this collapsing rule contains two variables, viz. x and y.

Theorem 5.4.28 (CR^∞ for almost non-collapsing systems, [KKSdV95a])
An orthogonal ITRS is infinitarily confluent iff it is almost non-collapsing.

This theorem cannot be generalised to weakly convergent reductions. Example 5.4.25 provides a counterexample.

Nevertheless, also for orthogonal systems, which are not almost non-collapsing, there is a weaker variant of infinitary confluence that these systems satisfy. For this variant, a certain set of terms is identified. As we will see Section 5.4.5, this set can be interpreted as a (subset of a) set of meaningless terms.

Definition 5.4.29 (hyper-collapsing term, [KKSdV95a])
Let $\mathcal{R}$ be an ITRS and $t, t' \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$.

(i) t is called *hypercollapsing* if every reduct of t can be reduced to a collapsing redex. That is, for each reduction $t \to^\star t'$, there is a reduction $t' \to^\star s$ to a collapsing redex s. The set of all hypercollapsing terms of $\mathcal{R}$ is denoted $\mathcal{HC}_\mathcal{R}$ or simply $\mathcal{HC}$ if $\mathcal{R}$ is clear from the context.

(ii) t and t' are called *hc-equivalent*, denoted $t \equiv_{hc} t'$ if t' can be obtained from t by replacing some pairwise disjoint hypercollapsing subterm occurrences in t with other hypercollapsing terms.

(iii) $\mathcal{R}$ is called *infinitarily confluent modulo hypercollapsing terms* (CR^∞_{hc}) if, for each two reductions $t \twoheadrightarrow t_1$ and $t \twoheadrightarrow t_2$, there are reductions $t_1 \twoheadrightarrow t_3$ and $t_2 \twoheadrightarrow t'_3$ such that $t_3 \equiv_{hc} t'_3$.

Intuitively speaking the hc-equivalence $t \equiv_{hc} t'$ means that t and t' are equal provided all hypercollapsing terms are identified. The structure of the property CR^∞_{hc} is illustrated in Figure 5.7. This weaker variant is indeed enjoyed by all orthogonal ITRSs:

Theorem 5.4.30 (CR^∞_{hc} for orthogonal systems, [KKSdV95a])
Every orthogonal ITRS is infinitary confluent modulo hypercollapsing terms.

Also this theorem cannot be generalised to weakly convergent reductions. The counterexample presented in Example 5.4.25 does apply to this setting as well. The TRS that is given there is non-collapsing. Hence, there are no hypercollapsing terms, i.e. in this case CR^∞_{hc} is equivalent to CR^∞.

Note that CR^∞_{hc} also implies the property $\mathsf{UN}^\infty_{\rightarrow}$. But, as a matter of fact, orthogonal systems even satisfy the stronger property NF^∞:

Theorem 5.4.31 (NF^∞ for orthogonal systems, [KKSdV95a])
Every orthogonal ITRS has the infinitary normal form property NF^∞.

Proof. In Theorem 7.15 from [KKSdV95a], this is proven for a different definition of NF^∞. However, as argued in Remark 3.3.4 this variant is equivalent to our definition of NF^∞. □

Recall that, according to Proposition 3.3.5, also in the infinitary setting, NF^∞ implies both UN^∞ and $\mathsf{UN}^\infty_{\rightarrow}$.

However, the above theorem does not hold for weakly convergent reductions. The TRS from Example 5.4.25 serves as a counterexample: It admits the finite reduction $f(c,a,a) \to f(c,a,b)$ and the weakly convergent reduction

$$f(c,a,a) \hookrightarrow^\omega f(g^\omega,a,a) \to^2 f(g^\omega,b,b)$$

to the normal form $f(g^\omega,b,b)$. Yet, as argued in Example 5.4.25, there is no weakly convergent reduction from $f(c,a,b)$ to $f(g^\omega,b,b)$.

5.4.5 Meaningless Terms and Böhm Trees

In Section 5.4.4, we have seen that, unlike in the finitary setting, orthogonal systems might fail to be infinitarily confluent. As presented in Theorem 5.4.30, orthogonal systems can be regarded as infinitary confluent provided that hypercollapsing terms are identified. This is reasonable as one can argue that hypercollapsing terms can be considered to be meaningless and, thus, contain the same information. The theory of meaningless terms [AKK+94, KvOdV99] in rewrite systems generalises this idea by studying axioms that describe a set of terms that may intuitively be seen as meaningless or undefined. This idea stems from the λ-calculus in which several notions of undefinedness exist (cf. [Bar92]). Intuitively, one can think of meaningless or undefined terms as terms that cannot be distinguished from one another and which cannot contribute any information to any computation.

Eventually, the notion of meaningless terms can be used to associate to each term a unique "normal form" which is constructed by performing a (transfinite) reduction and replacing undefined subterms by the special symbol $\bot$ along the way. The corresponding reduction systems are called *Böhm reductions*. Their normal forms, also called *Böhm trees* or *Böhm-like trees*, provide a syntax-based denotational semantics for rewrite systems parametrised by the choice of the set of terms to be considered as meaningless. This idea was first pursued in λ-calculus [Lév78, Bar84, Lév76, Lon83] using different notions of meaninglessness. Initially, Böhm trees were defined using *partial functions* [Bar84] and *direct approximants* [Lév78]. However, these construction can also be carried out using Böhm reductions [Ber96, KKSdV97]. We do not intend to go into all the details of Böhm trees. Instead we refer to the PhD thesis of Ketema [Ket06] which provides a comprehensive overview of the topic.

This section presents the axiomatisation of meaningless terms introduced in [KvOdV99] and the resulting theory of Böhm reductions and Böhm trees using the infinitary normal form approach. This will become particularly useful when the PRS model of transfinite term rewriting is studied in Section 5.5. It will turn out that term rewriting within the PRS model coincides with a particular Böhm reduction – the least one.

The least set of meaningless terms consists of the root-active terms which generalises the notion of hypercollapsing terms:

Definition 5.4.32 (root-active)
Let $\mathcal{R}$ be an ITRS and $t \in \mathcal{T}^\infty(\Sigma,\mathcal{V})$. t is called *root-active* if every reduct of t can be

reduced to a redex. That is, for each reduction $t \to^\star t'$, there is a reduction $t' \to^\star s$ to a redex s. The set of all root-active terms of $\mathcal{R}$ is denoted $\mathcal{RA}_\mathcal{R}$ or simply $\mathcal{RA}$ if $\mathcal{R}$ is clear from the context.

In contrast to hypercollapsing terms, we only require root-active terms to have all their reducts being reducible to a redex instead of a collapsing redex. Hence, $\mathcal{HC} \subseteq \mathcal{RA}$.

In order to define the axioms of meaningless terms, we need a means to identify certain subterm occurrences of two terms:

Definition 5.4.33 ($\mathcal{U}$-equivalence, [KvOdV99])
Let $\mathcal{U} \subseteq \mathcal{T}^\infty(\Sigma, \mathcal{V})$ and $s, t \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$. We say that s and t are *$\mathcal{U}$-equivalent*, denoted $s \equiv_\mathcal{U} t$, if t can be obtained from s by replacing some pairwise disjoint subterms of s in $\mathcal{U}$ with terms in $\mathcal{U}$.

Intuitively speaking, the $\mathcal{U}$-equivalence $s \equiv_\mathcal{U} t$ means that the terms s and t are equal modulo the terms in $\mathcal{U}$. This definition is a generalisation of hc-equivalence: It holds that $s \equiv_{hc} t$ iff $s \equiv_{\mathcal{HC}} t$.

Before we present the axiomatisation of meaningless terms we need a particular notion of overlapping:

Definition 5.4.34 (overlap)
Let t be a redex, i.e. an instance $l\sigma$ of the left-hand side l of some rewrite rule. The redex t is said to *overlap* its subterm at position π if π is a non-empty pattern position, i.e. $\pi \in \mathcal{P}_\Sigma(l) \setminus \{\varepsilon\}$.

Note that this concept of overlapping differs from the usual concept of overlapping of two rewrite rules. Here, only a single rule is concerned, whose redex overlaps an *arbitrary proper subterm* instead of another redex.

The following definition lists the axioms that we require in order to formally call the terms in a set of terms *meaningless*.

Definition 5.4.35 (set of meaningless terms, [KvOdV99])
Let $\mathcal{R}$ be an ITRS over Σ. A set $\mathcal{U} \subseteq \mathcal{T}^\infty(\Sigma, \mathcal{V})$ is said to be a *set of meaningless terms* of $\mathcal{R}$ if it satisfies the following axioms

(i) $\mathcal{U}$ is closed under finite reductions in $\mathcal{R}$. (descendants)

(ii) If an $\mathcal{R}$-redex t overlaps a subterm in $\mathcal{U}$, then $t \in \mathcal{U}$. (overlap)

(iii) $\mathcal{U}$ contains all root-active terms, i.e. $\mathcal{RA}_\mathcal{R} \subseteq \mathcal{U}$. (root-activeness)

(iv) If $s \equiv_\mathcal{U} t$, then $s \in \mathcal{U}$ iff $t \in \mathcal{U}$. (indiscernibility)

The root-activeness axiom illustrates what was meant by saying that $\mathcal{RA}$ is the least set of meaningless term. However, for this to make sense, it is, of course, necessary that $\mathcal{RA}$ is itself a set of meaningless terms. The following proposition confirms this:

Proposition 5.4.36 (root-active terms are meaningless, [KvOdV99])
For every orthogonal ITRS, the set of root-active terms is a set of meaningless terms.

The heart of the theory of meaningless terms is that all terms in such a set $\mathcal{U}$ of meaningless terms are considered to be identical, viz. identically undefined. This is the raison d'être of the notion of $\mathcal{U}$-equivalence $\equiv_\mathcal{U}$. Another way of enforcing the equivalence of all terms in $\mathcal{U}$ is to enrich the ITRS under consideration by additional rules which allow a meaningless term to be immediately contracted to $\bot$, a fresh constant symbol that we assume to be available for this purpose. However, one has to be careful when formalising this idea:

Definition 5.4.37 (Böhm reduction, [KvOdV99])
Let $\mathcal{R}$ be an ITRS over Σ, and $\mathcal{U} \subseteq \mathcal{T}^\infty(\Sigma, \mathcal{V})$.

(i) A term $t \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$ is called a *$\bot$,$\mathcal{U}$-instance* of a term $s \in \mathcal{T}^\infty(\Sigma_\bot, \mathcal{V})$ if t can be obtained from s by replacing each occurrence of $\bot$ in s with some term in $\mathcal{U}$. (Different occurrences of $\bot$ may be replaced with different terms.)

(ii) $\mathcal{U}_\bot$ is the set of terms in $\mathcal{T}^\infty(\Sigma_\bot, \mathcal{V})$ that have a $\bot$,$\mathcal{U}$-instance in $\mathcal{U}$.

(iii) The *Böhm reduction* of $\mathcal{R}$ w.r.t. $\mathcal{U}$ is the ITRS $\mathcal{B}_{\mathcal{R},\mathcal{U}} = (\Sigma_\bot, R \cup B)$, where

$$B = \{t \to \bot \mid t \in \mathcal{U}_\bot \setminus \{\bot\}\}$$

and we restrict the application of the rules in B to be allowed w.r.t. the identity substitution only. We write $\to_{\mathcal{U},\bot}$ for a reduction by a rule in B. That is, $s \to_{\mathcal{U},\bot} t$ iff $s = C[l]$ and $t = C[\bot]$ for $l \in \mathcal{U}_\bot \setminus \{\bot\}$ and some context $C[]$. If $\mathcal{R}$ and $\mathcal{U}$ are clear from the context, we simply write $\mathcal{B}$ and $\to_\bot$ instead of $\mathcal{B}_{\mathcal{R},\mathcal{U}}$ and $\to_{\mathcal{U},\bot}$, respectively.

By abuse of language, we will call a reduction sequence in a Böhm reduction a *Böhm reduction sequence* or simply *Böhm reduction*.

Note that the additional rules in the Böhm reduction are only allowed to be applied uninstantiated. The reason for this is the intuition of Böhm reductions to identify the terms in $\mathcal{U}$ (rather than instances of terms in $\mathcal{U}$). Additionally, this choice is motivated in the axiomatic character of the theory: It is simply not necessary to also allow instantiation of these rules.[1] Practically, however, this is not a restriction since all sets $\mathcal{U}$ of meaningless terms which are usually considered (cf. [KvOdV99]) are, in fact, closed under substitutions and, hence, so are the corresponding sets $\mathcal{U}_\bot$.

It is necessary to use the set $\mathcal{U}_\bot$ for the definition of Böhm reductions instead of $\mathcal{U}$ as one has to take into account that a subterm that is in $\mathcal{U}$ might have been already contracted to $\bot$. Hence, for each term $t \in \mathcal{U}$, the set $\mathcal{U}_\bot$ provides a variant of t in which some subterms (that are in $\mathcal{U}$ themselves) have been replaced with $\bot$.

Note that, formally, the above definition of Böhm reductions does not require the set $\mathcal{U}$ to be a set of meaningless terms. For some properties, none or only a subset of the axioms of meaningless terms is necessary. This will become crucial in Section 5.5. Additionally, we will also need some technical lemmas in order to work with Böhm reductions:

Lemma 5.4.38 (postponement of $\to_\bot$-steps, [KvOdV99])
Let $\mathcal{R}$ be a left-linear ITRS over Σ and $\mathcal{U} \subseteq \mathcal{T}^\infty(\Sigma, \mathcal{V})$. Then $s \twoheadrightarrow_{\mathcal{B}} t$ implies that there is some term $s' \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$ such that $s \twoheadrightarrow_{\mathcal{R}} s' \twoheadrightarrow_\bot t$[2].

Lemma 5.4.39 ($\bot$,$\mathcal{U}$-instances, [KvOdV99])
Let $\mathcal{U} \subseteq \mathcal{T}^\infty(\Sigma, \mathcal{V})$ be a set of terms satisfying the indiscernibility property and $t \in \mathcal{T}^\infty(\Sigma_\bot, \mathcal{V})$. If some $\bot$,$\mathcal{U}$-instance of t is in $\mathcal{U}$, then every $\bot$,$\mathcal{U}$-instance of t is.

Lemma 5.4.40 (compression of $\to_\bot$-steps)
$\mathcal{U} \subseteq \mathcal{T}^\infty(\Sigma, \mathcal{V})$ be a set of terms satisfying the indiscernibility property, $t_0 \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$, $t_\alpha \in \mathcal{T}^\infty(\Sigma_\bot, \mathcal{V})$, and $S\colon t_0 \twoheadrightarrow_\bot^\alpha t_\alpha$. Then there is a reduction sequence $T\colon t_0 \twoheadrightarrow_\bot^{\le\omega} t_\alpha$ that is a complete development of a set of disjoint occurrences of terms $u \in \mathcal{U}$ in t_0.

Proof. The argument is essentially the same as in Lemma 7.2.4 from [Ket06]. Let $S = (t_\iota \to_{\pi_\iota} t_{\iota+1})_{\iota<\alpha}$ be the mentioned reduction sequence strongly converging to t_α, and let π be a position at which some reduction step in S takes place. That is, there is some β such that $\pi_\beta = \pi$. We will prove by induction on β that $t_0|_\pi \in \mathcal{U}$.

Consider the term $t_\beta|_\pi$. Since a $\bot$-rule is applied here, we have that $t_\beta|_\pi \in \mathcal{U}_\bot$. Let $V = \mathcal{P}_\bot(t_\beta|_\pi)$. Hence, for each $v \in V$, there is some $\gamma < \beta$ such that $\pi_\gamma = \pi \cdot v$. Therefore, we can apply the induction hypothesis and get that $t_0|_{\pi\cdot v} \in \mathcal{U}$ for all $v \in V$. It is clear that we can

[1] unless one considers higher-order reduction systems such as λ-calculus (cf. [KvOdV99])

[2] Strictly speaking, if s is not a total term, i.e. it contains $\bot$, then we have to consider the system that is obtained from $\mathcal{R}$ by extending its signature to $\Sigma_\bot$.

obtain $t_0|_\pi$ from $t_\beta|_\pi$ by replacing each $\bot$-occurrence at $v \in V$ with the corresponding term $t_0|_{\pi \cdot v}$. That is, $t_0|_\pi$ is a $\bot,\mathcal{U}$-instance of $t_\beta|_\pi$. Because $t_\beta|_\pi \in \mathcal{U}_\bot$, there is some $\bot,\mathcal{U}$-instance of $t_\beta|_\pi$ in $\mathcal{U}$. Thus, by Lemma 5.4.39, also $t_0|_\pi$ is in $\mathcal{U}$. This closes the proof of the claim.

Now let $V = \mathcal{P}_\bot(t_\alpha)$. Clearly, all positions in V are pairwise disjoint. Moreover, for each $v \in V$, there is a step in S that takes place at v. Hence, by the claim shown above, V is a set of occurrences in t_0 of terms in $\mathcal{U}$. A complete development of V in t_0 leads to t_α and can be performed in at most ω. □

The following two theorems summarise the most important properties for Böhm reductions for orthogonal systems:

Theorem 5.4.41 (infinitary normalisation of Böhm reductions, [KvOdV99])
Let $\mathcal{R}$ be an orthogonal ITRS and $\mathcal{U}$ a set of terms satisfying the root-activeness axiom. Then the Böhm reduction $\mathcal{B}_{\mathcal{R},\mathcal{U}}$ is infinitarily normalising.

Theorem 5.4.42 (infinitary confluence of Böhm reductions, [KvOdV99])
Let $\mathcal{R}$ be an orthogonal ITRS and $\mathcal{U}$ a set of meaningless terms in $\mathcal{R}$. Then the Böhm reduction $\mathcal{B}_{\mathcal{R},\mathcal{U}}$ is infinitarily confluent.

That is, the Böhm reduction of an orthogonal ITRS w.r.t. a set of meaningless terms provides a unique infinitary normal form for each term. This normal form is called the *Böhm tree* of the term.

5.4.6 Reduction Strategies

This section is concerned with infinitarily normalising reduction strategies, i.e. strategies which generate possibly transfinite reductions always leading to a normal form, provided the starting term has a normal form. In particular, we are interested in reduction strategies which find a normal form within at most ω steps, as only such strategies are of practical impact. They can be considered as a realisation of Theorem 5.4.4 and Corollary 5.4.5 which assert that it is possible to approximate the final term of a possibly transfinite reduction arbitrarily well by a finite reduction resp. to reach a finite term by a finite reduction. In a similar way, these strategies realise the Compression Lemma.

For our purposes, the following quite abstract definition of a reduction strategy will suffice:

Definition 5.4.43 (reduction strategy)
Let $\mathcal{R}$ be an ITRS. A *reduction strategy* for $\mathcal{R}$ is a set $\mathcal{S}$ of strongly convergent reductions in $\mathcal{R}$. For a reduction strategy $\mathcal{S}$, we also write $\mathcal{S}(t)$ to denote the set $\{T \in \mathcal{S} \,|\, T\colon t \twoheadrightarrow t'\} \cup \left\{t \to^0 t\right\}$ of reduction sequences in $\mathcal{S}$ starting in t plus the empty reduction sequence.

The above definition deviates from the notion of a reduction strategy usually found in the literature. Typically, a reduction strategy consist of a mapping that assigns to each term, which is not in normal form, a single reduction step or a non-empty finite reduction sequence that starts in this term. Here, a reduction strategy $\mathcal{S}$ is simply a set that lists all reduction sequences that it “generates”. For investigating strategies in detail, this definition might not be appropriate. However, we are only interested in presenting some strategies and their properties. On the other hand, the usual notion of a reduction strategy is, in fact, not adequate for our purposes since we want to discuss reduction strategies with a fairness constraint.

Definition 5.4.44 (infinitarily normalising, ω-normalising)
Let $\mathcal{R}$ be an ITRS, T be a strongly continuous reduction in $\mathcal{R}$, and $\mathcal{S}$ a reduction strategy for $\mathcal{R}$.

(i) T is called *infinitarily normalising* if it strongly converges to a normal form.

(ii) T is called *ω-normalising* if it is of length at most ω and strongly converges to a normal form.

(iii) $\mathcal{S}$ is called *strongly convergent* if, whenever arbitrarily long proper prefixes of an open strongly continuous reduction sequence R are contained in $\mathcal{S}$, then also R itself is in $\mathcal{S}$.

(iv) $\mathcal{S}$ is called *infinitarily normalising* resp. *ω-normalising* if $\mathcal{S}$ is strongly convergent and, for each term t having a normal form, all maximal reduction sequences in $\mathcal{S}(t)$ are infinitarily normalising resp. ω-normalising reduction sequences.

Including strong convergence in the notions of ω-normalisation and infinitary normalisation is necessary in order to guarantee that maximal reduction sequences always exist. Otherwise, every reduction strategy that only permits finite reductions would be ω-normalising. This is not desired. For example, consider the system with the rules $a \to a$ and $a \to b$. a has the normal form b. Without the requirement of strong convergence, the reduction strategy

$$\mathcal{S} = \{T \,|\, T\colon a \to^{\star} a\}$$

would be normalising. This is due to the fact that $\mathcal{S}(a) = \{T \,|\, T\colon a \to^{\star} a\}$ has no maximal element. However, as $\mathcal{S}$ has to be strongly convergent, it also has to contain the ω-reduction sequence $T\colon a \to a \to \dots$ because $\mathcal{S}(a)$ contains arbitrarily long proper prefixes of T. But note that in this particular case including T into a reduction strategy is not allowed as T is not strongly convergent.

Next we present the needed reduction strategy [HL91a, HL91b]. The idea of needed reduction is that only those redexes are contracted which cannot be avoided during a reduction to normal form (and, therefore, are needed).

Definition 5.4.45 (needed redex, needed reduction, [KKSdV95a])
Let $\mathcal{R}$ be an ITRS, and t a term in $\mathcal{R}$. A redex occurrence u in t is called *needed* if, in every reduction starting in t that strongly converges to a normal form, some residual of u is contracted. A strongly convergent reduction sequence in $\mathcal{R}$ is called *needed* if all its reduction steps contract a needed redex occurrence.

This merely defines a property on reduction sequences. However, any property on reduction sequences induces a reduction strategy:

Notation 5.4.46. Every property P on reduction sequences gives rise to a reduction strategy $\mathcal{S}_P$, called *P-reduction*, which is the set of strongly convergent reduction sequences satisfying P. For example, *needed reduction* is the reduction strategy that only contracts needed redexes.

For infinitary normalisation, it is necessary that the reduction strategy only "halts" at normal forms. For needed reductions, this means that needed redexes must always exists in terms having a normal form, but which are not normal forms themselves.

Theorem 5.4.47 (existence of needed redexes, [KKSdV95a])
For orthogonal ITRSs, every term having a normal form but not being a normal form contains a needed redex occurrence.

With this theorem the only thing that remains, in order to obtain infinitary normalisation, is that strongly continuous needed reduction sequences are always strongly convergent. This was confirmed in [KKSdV95a] and thus yields the following theorem:

Theorem 5.4.48 (needed reduction is infinitarily normalising, [KKSdV95a])
For orthogonal ITRSs, needed reduction is infinitary normalising.

(a) Full binary tree. (b) Degenerated binary tree.

Figure 5.8: Binary trees.

This is a generalisation of a corresponding result for finitary term rewriting [HL91a, HL91b].

Unfortunately, needed reduction is not ω-normalising. The problem that might occur is that needed redex occurrences are infinitely often ignored:

Example 5.4.49
Let $\mathcal{R}$ be the TRS consisting of the single rule $a \to f(a,a)$. In every term each occurrence of the redex a is needed, i.e. every reduction sequence is needed. The term a has the unique normal form depicted in Figure 5.8a. However, the needed ω-reduction sequence

$$f(a,a) \to f(f(a,a),a) \to f(f(f(a,a),a),a) \to \ \dots$$

strongly converges to the non-normal form depicted in Figure 5.8b.

An issue similar to the one illustrated in Example 5.4.49 can also occur if a term contains infinitely many needed redexes. In order to obtain an ω-normalising reduction strategy, a notion of fairness has to be introduced:

Definition 5.4.50 ((needed-)fairness, [KKSdV95a])
Let $\mathcal{R}$ be an ITRS.

(i) Let R be a function that maps each term t to a set of redex occurrences in t. A reduction sequence $(t_i \to_{\pi_i} t_{i+1})_{i<\alpha}$ of length at most ω strongly converging to t_α in $\mathcal{R}$ is called *R-fair* if, for each $i < \alpha$ and $u \in R(t_i)$, there is some $i \leq j < \alpha$ such that $\pi_j \leq u$.

(ii) A reduction sequence is called *needed-fair* if it is R-fair, where R is the function that maps each term to the set of all its needed redex occurrences.

(iii) A reduction sequence is called *fair* if it is R-fair, where R is the function that maps each term to the set of all its redex occurrences.

That is, in order to obtain (R-)fairness, there can only be finitely many steps between the creation of a redex occurrence (according to R) and the contraction of that redex occurrence or a redex occurrence above.

Note that each fair reduction sequence is also R-fair for any R, i.e. particularly, each fair reduction is also needed-fair.

Theorem 5.4.51 ((needed-)fair reduction is ω-normalising, [KKSdV95a])
For orthogonal ITRSs, needed-fair and fair reduction are ω-normalising.

Also the parallel-outermost reduction strategy is well-known from finitary term rewriting. It can be straightforwardly generalised to the infinitary setting:

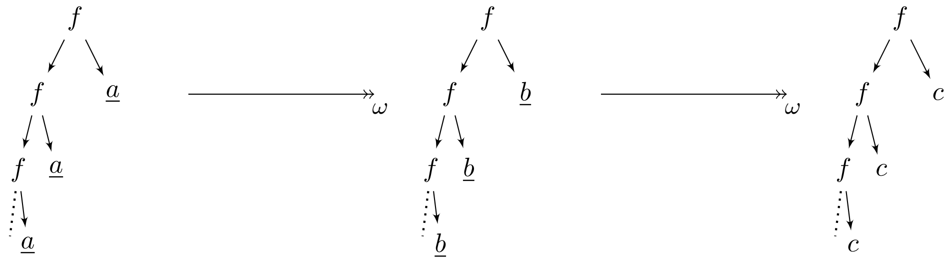

Figure 5.9: A parallel-outermost reduction.

Definition 5.4.52 (parallel-outermost reduction)
Let $\mathcal{R}$ be an ITRS. A strongly convergent reduction sequence S in $\mathcal{R}$ is called *parallel-outermost* if there is a sequence $(S_\iota)_{\iota<\alpha}$ of complete developments $S_\iota\colon s_\iota \twoheadrightarrow_{U_\iota} t_\iota$, where U_ι is the set of all outermost redex occurrences in s_ι, such that $S = \prod_{\iota<\alpha} S_\iota$.

Theorem 5.4.53 (parallel-outermost reduction is inf. normalising, [KKSdV95a])
For orthogonal ITRSs, parallel-outermost reduction is infinitarily normalising.

If we restrict ourselves to TRS, i.e. finite right-hand sides, and finite starting terms, then parallel-outermost reduction is even ω-normalising.

Theorem 5.4.54 (parallel-outermost reduction is ω-normalising, [Mid97])
For almost orthogonal TRSs, parallel-outermost reduction is ω-normalising for finite terms.

It is rather obvious that parallel-outermost reductions are in general not ω-normalising for infinite terms. The problem is that an infinite term could contain infinitely many outermost redex occurrences. Hence, a parallel step itself might take ω steps:

Example 5.4.55
Let $\mathcal{R}$ be the orthogonal TRS given by the rules:

$$a \to b, \quad b \to c$$

$\mathcal{R}$ admits the parallel-outermost reduction depicted in Figure 5.9. The outermost redex occurrences of the terms are indicated by underlinings. Each of the two ω-reduction sequences is a parallel-outermost step.

Another reduction strategy that is ω-normalising is the so-called depth-increasing reduction:

Definition 5.4.56 (depth-increasing reduction, [KKSdV95a])
Let $\mathcal{R}$ be an orthogonal ITRS. A strongly convergent reduction sequence S in $\mathcal{R}$ is called *depth-increasing* if there is a sequence $(S_i)_{i<\omega}$ of complete developments $S_i\colon s_i \twoheadrightarrow_{U_i} t_i$, where U_i is the set of all redex occurrences in s_i at depth $\leq i$, such that $S = \prod_{i<\omega} S_i$.

Theorem 5.4.57 (depth-increasing reduction is inf. normalising, [KKSdV95a])
For orthogonal ITRSs, depth-increasing reduction is ω-normalising.

As mentioned in [Mid97], the reduction strategy $\mathcal{S}_\omega$ defined in [AM96] is ω-normalising as well:

Theorem 5.4.58 ($\mathcal{S}_\omega$ is ω-normalising, [Mid97])
For almost orthogonal TRSs, $\mathcal{S}_\omega$ is ω-normalising for finite terms.

5.4.7 Termination

This section contains some brief remarks about infinitary termination properties. We have seen that for transfinite reductions, similarly to the finitary case, infinitary termination implies infinitary normalisation. Astonishingly, when considering orthogonal ITRSs, infinitary normalisation also implies infinitary termination – at least globally:

Theorem 5.4.59 (equivalence of WN^∞ and SN^∞, [KdV05])
Each orthogonal ITRS is SN^∞ *iff it is* WN^∞*.*

This, however, does not hold for individual terms as the following example reveals:

Example 5.4.60 ([KdV05])
Let $\mathcal{R}$ be the orthogonal TRS

$$\rho_1\colon c \to c, \quad \rho_2\colon f(x) \to f(f(x))$$

The term $f(c)$ is WN^∞ as it has the normal form f^ω reachable by the strongly convergent ω-reduction sequence

$$f(c) \to_{\rho_2} f^2(c) \to_{\rho_2} f^3(c) \to_{\rho_2} \dots\ f^\omega$$

Yet, $f(c)$ is not SN^∞ as it starts a strongly divergent ω-reduction sequence

$$f(c) \to_{\rho_1} f(c) \to_{\rho_1} f(c) \to_{\rho_1} \dots$$

Also other variants of termination properties such as infinitary weak resp. strong head normalisation (cf. [KdV05]) are equivalent to SN^∞ at the global level.

The requirement of orthogonality is essential for Theorem 5.4.59:

Example 5.4.61
Let $\mathcal{R}$ be the TRS:

$$a \to a, \quad a \to b$$

$\mathcal{R}$ is WN^∞ as both a and b have the normal form b. Yet, $\mathcal{R}$ is not SN^∞ because of the strongly divergent reduction

$$a \to a \to a \to \dots$$

5.5 Strongly Convergent PRS Reductions

In this section we want to study the properties of strongly convergent reductions in ITRSs w.r.t. its PRS semantics that was introduced in this thesis in Section 3.2. According to Corollary 5.2.3, the PRS model is a conservative extension of the MRS model. Therefore, some results known from the MRS world can be carried over to the PRS world. This includes, for example, results involving the infinitary normalisation property WN^∞.

We restrict our study of strongly convergent PRS reductions chiefly to confluence properties. For this purpose, we develop a theory of descendants and complete developments similar to that of MRS reductions. This is done in Section 5.5.1 and Section 5.5.2, respectively. Subsequently, in Section 5.5.3, the equivalence of transfinite Böhm reductions in the MRS model and transfinite reductions in the PRS model is established. Afterwards, this equivalence result is used in order to show confluence, normalisation and compression properties. As in the previous section, the analysis of PRS reductions in this section is mostly restricted to orthogonal systems.

An analysis of strongly convergent reductions in a partial order model comparable to ours was done by Blom [Blo04] for the λ-calculus. In his paper Blom was able to show the equivalence of partial order reductions and Böhm reductions in the metric model. Due to the characteristics of infinitary λ-calculi, for which usually several different metrics are considered, the investigated partial order model is different from ours: It is rather involved, specifically devised for the needs of the λ-calculus, and does not allow a notion of weak convergence.

Also Corradini [Cor93] has considered a partial order model of transfinite rewriting. However, it is only used in order to formalise parallel reductions of an infinite set of redexes in left-linear term rewriting systems. To achieve this, he employs a non-standard semantics of term rewriting allowing the partial matching of rules. With the theory of complete developments that we will develop in this chapter, we are able to formalise infinite parallel reductions in a more straightforward way. Nevertheless, the connection to cyclic term graph rewriting that Corradini has established is intriguing and suggests that future research on PRS reductions in this direction is worthwhile.

Before we begin with the investigation of the topics promised above, we want to establish some fundamental properties of infinitary term rewriting in the PRS model. These will not only serve as useful lemmas for later proofs, but will also give further insight into the behaviour of transfinite reductions in this model.

The foremost difference between the MRS and the PRS model is that the latter can spawn $\bot$ symbols. The following lemmas provide several characterisations for this behaviour and its absence, respectively.

Lemma 5.5.1 (non-$\bot$ symbols in open reductions)
Let $\mathcal{R}$ be an ITRS and $S\colon s \twoheadrightarrow^{\lambda}_{\mathcal{R}} t$ an open reduction sequence with $S = (t_\iota \to_{\pi_\iota, c_\iota} t_{\iota+1})_{\iota<\lambda}$. Then the following statements are equivalent for all positions π:

(a) $\pi \in \mathcal{P}_{\not\bot}(t)$,

(b) there is some $\alpha < \lambda$ such that $c_\iota(\pi) = t(\pi) \neq \bot$ for all $\alpha \leq \iota < \lambda$, and

(c) there is some $\alpha < \lambda$ such that $t_\alpha(\pi) = t(\pi) \neq \bot$ and $\pi_\iota \nleq \pi$ for all $\alpha \leq \iota < \lambda$.

(d) there is some $\alpha < \lambda$ such that $\pi \in \mathcal{P}_{\not\bot}(t_\alpha)$ and $\pi_\iota \nleq \pi$ for all $\alpha \leq \iota < \lambda$.

Proof. At first consider the implication from (a) to (b). To this end, let $\pi \in \mathcal{P}_{\not\bot}(t)$, i.e. we have that $t(\pi) \neq \bot$. Define $s_\gamma = \bigsqcap^{\bot}{}_{\gamma\leq\iota<\lambda} c_\iota$ for each $\gamma < \lambda$. Note that then $t = \bigsqcup^{\bot}{}_{\gamma<\lambda} s_\gamma$. Applying Corollary 4.4.16 yields that there is some $\alpha < \lambda$ such that $s_\alpha(\pi) = t(\pi)$. Hence, the fact that $s_\alpha \leq_\bot s_\iota$ for all $\alpha \leq \iota < \lambda$ implies, by Corollary 4.4.14, that $s_\iota(\pi) = t(\pi)$ for all $\alpha \leq \iota < \lambda$. Since $s_\iota \leq_\bot c_\iota$ for all $\iota < \lambda$, we can again employ Corollary 4.4.14 to obtain that $c_\iota(\pi) = t(\pi)$ for all $\alpha \leq \iota < \lambda$.

Next consider the implication from (b) to (c). Let $\alpha < \lambda$ be such that $c_\iota(\pi) = t(\pi) \neq \bot$ for all $\alpha \leq \iota < \lambda$. Recall that $c_\iota = t_\iota[\bot]_{\pi_\iota}$ for all $\iota < \lambda$. Hence, the fact that $\pi \in \mathcal{P}_{\not\bot}(c_\iota)$ for all $\alpha \leq \iota < \lambda$ implies that $t_\alpha(\pi) = c_\alpha(\pi)$ and that $\pi_\iota \nleq \pi$ for all $\alpha \leq \iota < \lambda$. Since $c_\alpha(\pi) = t(\pi) \neq \bot$, we also have $t_\alpha(\pi) = t(\pi) \neq \bot$.

The implication from (c) to (d) is trivial.

Finally, consider the implication from (d) to (a). For this purpose, let $\alpha < \lambda$ be such that (1) $\pi \in \mathcal{P}_{\not\bot}(t_\alpha)$ and (2) $\pi_\iota \nleq \pi$ for all $\alpha \leq \iota < \lambda$. Consider the set P consisting of all positions in t_α that are prefixes of π. P is obviously closed under prefixes and, because of (2), all terms in the set $T = \{c_\iota \mid \alpha \leq \iota < \lambda\}$ coincide in all positions in P. According to Lemma 4.4.23, also $s_\alpha = \bigsqcap^{\bot} T$ coincides with all terms in T in all positions in P. Since $\pi \in P$ and $c_\alpha \in T$, we thereby obtain that $c_\alpha(\pi) = s_\alpha(\pi)$. As we also have $t_\alpha(\pi) = c_\alpha(\pi)$, due to (2), and $\pi \in \mathcal{P}_{\not\bot}(t_\alpha)$ we can infer that $\pi \in \mathcal{P}_{\not\bot}(s_\alpha)$. Applying Corollary 4.4.14 eventually yields that $\pi \in \mathcal{P}_{\not\bot}(t)$ because $s_\alpha \leq_\bot t$. □

In a nutshell, the above lemma states that the symbol at a certain position π of the final term of an open reduction is not $\bot$ iff there is an upper bound α on the indexes of the reduction steps that take place at π or above. And, additionally, the symbol at that position of the final term coincides with the symbol at the same position in all terms and contexts after the α-th step.

Next we consider the phenomenon that causes $\bot$ symbols to appear. The concept of so-called volatile positions will turn out to be crucial to this effect:

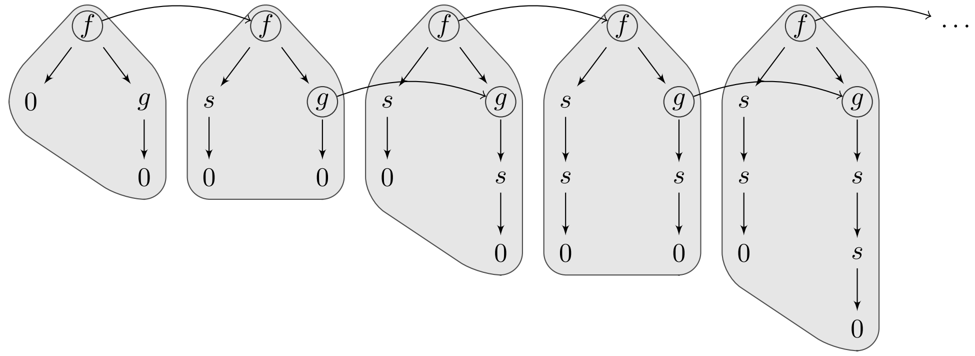

Figure 5.10: Reduction sequence with volatile positions.

Definition 5.5.2 (volatile position)
Let $\mathcal{R}$ be an ITRS, $S = (t_\iota \to_{\pi_\iota} t_{\iota+1})_{\iota<\lambda}$ a open strongly continuous reduction sequence in $\mathcal{R}$, and π a position. π is said to be *volatile* in S if, for each ordinal $\beta < \lambda$, there is some $\beta \leq \gamma < \lambda$ such that $\pi_\gamma = \pi$. If π is volatile in S and no proper prefix of π is volatile in S, then π is called *outermost-volatile* .

The notion of volatile positions in some sense negates the characterisation of the creation non-$\bot$ symbols: A position π is volatile if the indexes of reduction steps performed at π are *not* bounded by some ordinal $\alpha < \lambda$.

Example 5.5.3
Consider the TRS $\mathcal{R}$ consisting of the rules

$$\rho_1\colon f(x,y) \to f(s(x),y), \quad \rho_2\colon g(x) \to g(s(x))$$

$\mathcal{R}$ admits the ω-reduction sequence

$$\begin{aligned} S\colon f(0,g(0)) &\to_{\rho_1} f(s(0),g(0)) \to_{\rho_2} f(s(0),g(s(0))) \\ &\to_{\rho_1} f(s^2(0),g(s(0))) \to_{\rho_2} f(s^2(0),g(s^2(0))) \to_{\rho_1} \dots \end{aligned}$$

which weakly converges to $f(s^\omega, g(s^\omega))$ and strongly converges to $\bot$. The reductions sequence S is also illustrated in Figure 5.10. The positions at which the reduction steps are performed are indicated by circles and reduction arrows. One can see from the picture that both $\pi = \varepsilon$ and $\pi' = 1$ are volatile positions in S. Again and again reductions take place at π and π'. Since these are the only volatile positions in S and it holds that $\pi \leq \pi'$, we have that π is outermost-volatile.

The following lemma shows that $\bot$ symbols are produced in open reduction sequence at outermost-volatile positions and that this is the only way in which $\bot$ symbols can appear.

Lemma 5.5.4 ($\bot$ symbols in open reductions)
Let $\mathcal{R}$ be an ITRS and $S = (t_\iota \to_{\pi_\iota} t_{\iota+1})_{\iota<\alpha}$ an open reduction sequence in $\mathcal{R}$ strongly converging to t_α. Then, for every position π, we have

(i) If π is volatile in S, then $\pi \notin \mathcal{P}_{\not\bot}(t_\alpha)$.

(ii) $t_\alpha(\pi) = \bot$ iff

(a) π is outermost-volatile in S, or

(b) there is some $\beta < \alpha$ such that $t_\beta(\pi) = \bot$ and $\pi_\iota \nleq \pi$ for all $\beta \leq \iota < \alpha$.

(iii) Let t_ι be total for all $\iota < \alpha$. Then $t_\alpha(\pi) = \bot$ iff π is outermost-volatile in S.

Proof. (i) follows from Lemma 5.5.1, in particular the equivalence of (a) and (c).

(ii) At first consider the "only if" direction. To this end, suppose that $t_\alpha(\pi) = \bot$. In order to show that then (a) or (b) holds, we will prove that (b) must hold true whenever (a) does not hold. For this purpose, we assume that π is not outermost-volatile in S. Note that no proper prefix π' of π can be volatile in S as this would imply, according to clause (i), that $\pi' \notin \mathcal{P}_{\setminus\bot}(t_\alpha)$ and, therefore, $\pi \notin \mathcal{P}(t_\alpha)$. Hence, π is also not volatile in S. In sum, no prefix of π is volatile in S. Consequently, there is an upper bound $\beta < \alpha$ on the indices of reduction steps taking place at π or above. But then $t_\beta(\pi) = \bot$ since otherwise Lemma 5.5.1 would imply that $t_\alpha(\pi) \neq \bot$. This shows that (b) holds.

For the converse direction, we will show that both (a) and (b) imply that $t_\alpha(\pi) = \bot$.

(a) Let π be outermost-volatile in S. By clause (i), this implies $\pi \notin \mathcal{P}_{\setminus\bot}(t_\alpha)$. Hence, it remains to be shown that $\pi \in \mathcal{P}(t_\alpha)$. If $\pi = \varepsilon$, then this is trivial. Otherwise, π is of the form $\pi' \cdot i$. Since all proper prefixes of π are not volatile, there is some $\beta < \alpha$ such that $\pi_\beta = \pi$ and $\pi_\iota \nleq \pi'$ for all $\beta \leq \iota < \alpha$. This implies that $\pi \in \mathcal{P}(t_\beta)$. Hence, $t_\beta(\pi') = f$ is a symbol having an arity of at least $i + 1$. Consequently, according to Lemma 5.5.1, also $t_\alpha(\pi') = f$. Since f's arity is at least $i + 1$, also $\pi = \pi' \cdot i \in \mathcal{P}(t_\alpha)$.

(b) Let $\beta < \alpha$ such that $t_\beta(\pi) = \bot$ and $\pi_\iota \nleq \pi$ for all $\beta \leq \iota < \alpha$. According to Proposition 2.1.32, we have that $t_\alpha = \bigsqcup^\bot{}_{\beta \leq \gamma < \alpha} \bigsqcap^\bot{}_{\gamma \leq \iota < \alpha} c_\iota$. Define $s_\gamma = \bigsqcap^\bot{}_{\gamma \leq \iota < \alpha} c_\iota$ for each $\gamma < \alpha$. Since from β onwards no reduction takes place at π or above, it holds that all c_ι, for $\beta \leq \iota < \alpha$, coincide in all prefixes of π. By Lemma 4.4.23, this also holds for all s_ι and c_ι with $\beta \leq \iota < \alpha$. Since $c_\beta(\pi) = t_\beta(\pi) = \bot$, this means that $s_\iota(\pi) = \bot$ for all $\beta \leq \iota < \alpha$. Recall that $t_\alpha = \bigsqcup^\bot{}_{\beta \leq \gamma < \alpha} s_\gamma$. Hence, according to Corollary 4.4.16, we can conclude that $t_\alpha(\pi) = \bot$.

(iii) is a special case of (ii): If each t_ι, $\iota < \alpha$, is total, then (b) cannot be true. □

We can apply this lemma to Example 5.5.3: As we have seen, the position $\pi = \varepsilon$ is outermost-volatile in the reduction sequence S mentioned in the example. Hence, S strongly converges to a term that has, according to Lemma 5.5.4, the symbol $\bot$ at position π. In other words: S strongly converges to $\bot$.

From the lemma above, one can see that the absence of volatile positions in strongly convergent reductions is equivalent to the totality of strongly convergent reductions.

Lemma 5.5.5 (total reductions)
Let $\mathcal{R}$ be an ITRS, s a total term in $\mathcal{R}$, and $S\colon s \twoheadrightarrow t$ a strongly convergent reduction sequence in $\mathcal{R}$. $S\colon s \twoheadrightarrow t$ is total iff no prefix of S has a volatile position.

Proof. The "only if" direction follows straightforwardly from Lemma 5.5.4(iii).

We prove the "if" direction by induction on the length of S. If $|S| = 0$, then the totality of S follows by the assumption of s being total. If $|S|$ is a successor ordinal, then the totality of S follows from the induction hypothesis since single step reductions preserve totality. If $|S|$ is a limit ordinal, then the totality of S follows from the induction hypothesis using Lemma 5.5.4(iii). □

5.5.1 Descendants

In this section we introduce the notion of descendants to the setting of strongly convergent PRS reductions and analyse its properties. Descendants were already covered in the setting of strongly convergent MRS reductions in Section 5.4.2. The definition of this concept in the present setting is very similar. For technical reasons, however, it is restricted to non-$\bot$-occurrences:

Definition 5.5.6 (descendants, residuals)
Let $\mathcal{R}$ be a ITRS, $S\colon t_0 \twoheadrightarrow^\alpha_{\mathcal{R}} t_\alpha$, and $U \subseteq \mathcal{P}_{\setminus\bot}(t_0)$. The *descendants* of U by S, denoted $U /\!/ S$, is the set of occurrences in t_α inductively defined as follows:

(a) If $\alpha = 0$, then $U/\!/S = U$.

(b) If $\alpha = 1$, let $S\colon t_0 \to_{\pi,\rho} t_1$ with $\rho\colon l \to r$. Take any $u \in U$ and define the set R_u as follows: If $\pi \nleq u$, then $R_u = \{u\}$. If u is in the pattern of the redex of ρ, then $R_u = \emptyset$. Otherwise, i.e. if $u = \pi \cdot w \cdot x$, with $l|_w \in \mathcal{V}$, then $R_u = \{\pi \cdot w' \cdot x \mid r|_{w'} = l|_w\}$. Define $U/\!/S = \bigcup_{u\in U} R_u$.

(c) If $\alpha = \alpha' + 1$, then $U/\!/S = (U/\!/S|_{[0,\alpha')})/\!/S|_{[\alpha',\alpha)}$.

(d) If α is a limit ordinal, then $\quad U/\!/S = \mathcal{P}_{\setminus\bot}(t_\alpha) \cap \liminf_{\iota\to\alpha} U/\!/S|_{[0,\iota)}$
That is, $\quad u \in U/\!/S \quad$ iff $\quad u \in \mathcal{P}_{\setminus\bot}(t_\alpha)$ and $\exists\beta < \alpha \forall\beta \le \iota < \alpha\colon u \in U/\!/S|_{[0,\iota)}$

If, in particular, U is a set of redex occurrences, then $U/\!/S$ is also called the set of *residuals* of U by S. Moreover, by abuse of notation, we write $u/\!/S$ instead of $\{u\}/\!/S$.

Note that the above definition can also, just as the corresponding definition for MRS reductions, be applied to weakly convergent reductions. This, however, would not yield a desirable formalisation of the concept of descendants as we will see later.

The definition of descendants by reductions in the PRS model is almost verbatimly the same as the one for the MRS model. The only difference is the explicit restriction of descendants to non-$\bot$-occurrences for the limit ordinal case. As this restriction is satisfied for MRS reductions anyway, the above definition is equivalent to the MRS version if restricted to MRS reductions.

Remark 5.5.7. One can easily see that the descendants of a set of non-$\bot$-occurrences is again a set of non-$\bot$-occurrences. The restriction to non-$\bot$-occurrences has to be made explicit for the case of open reductions. In fact, without this explicit restriction the definition would yield descendants which might not even be occurrences in the final term t_α of the reduction. For example, consider the system with the single rule $f(x) \to x$ and the term f^ω and let S be the strongly convergent reduction sequence

$$S\colon f^\omega \to f^\omega \to \ \ldots\ \bot$$

that contracts the redex at the root of f^ω in each step. Consider $U = \{\varepsilon, 0, 0^2, 0^3, \ldots\}$, the set of all occurrences in t^ω. Without the abovementioned restriction, the descendants of U by S would be U itself as the descendants of U by each proper prefix of S is also U. However, none of the occurrences $0, 0^2, 0^3, \ldots \in U$ is even an occurrence in the final term $\bot$. The $\varepsilon \in U$ is an occurrence in $\bot$, but only a $\bot$-occurrence. With the restriction to non-$\bot$-occurrences we indeed get the expected result $U/\!/S = \emptyset$.

On the other hand, the exclusion of $\bot$-occurrences does not affect the definition of residuals. By definition, the root symbol of a redex cannot be $\bot$. Additionally, the concept of descendants defined on PRS reductions is compatible with the corresponding concept on MRS reductions, i.e. whenever a reduction sequence is strongly convergent in both the MRS and the PRS model both notions of descendants coincide.

The following proposition confirms a property of descendants that one expects intuitively: The descendants of descendants are again descendants. That is, the concept of descendants is composable.

Proposition 5.5.8 (descendants of sequential reductions)
Let $\mathcal{R}$ be a ITRS, $S\colon t_0 \twoheadrightarrow^\alpha_\mathcal{R} t_1$, $T\colon t_1 \twoheadrightarrow^\beta_\mathcal{R} t_2$, and $U \subseteq \mathcal{P}_{\setminus\bot}(t_0)$. Then $U/\!/S\cdot T = (U/\!/S)/\!/T$.

Proof. We conduct the proof by induction on β. If $\beta = 0$, then the statement is trivially true. Suppose that β is a successor ordinal $\beta' + 1$. Let T_1 and T_2 denote $T|_{[0,\beta')}$ and $T|_{[\beta',\beta)}$, respectively. Then we have the following equations:

$$U/\!/S\cdot T = U/\!/S\cdot T_1\cdot T_2 = (U/\!/S\cdot T_1)/\!/T_2 \overset{IH}{=} ((U/\!/S)/\!/T_1)/\!/T_2 = (U/\!/S)/\!/T_1\cdot T_2 = (U/\!/S)/\!/T$$

Suppose β is a limit ordinal. Note that, by Lemma 2.1.19, also $\alpha + \beta$, the length of $S \cdot T$, is a limit ordinal. Therefore, we can reason as follows:

$$\begin{array}{rll} u \in U /\!/ S \cdot T & \text{iff} & u \in \mathcal{P}_{\smallsetminus\perp}(t_2) \text{ and } \exists \gamma < \alpha + \beta \forall \gamma \leq \iota < \alpha + \beta \;\; u \in U /\!/ (S \cdot T)|_{[0,\iota)} \\ & \text{iff} & u \in \mathcal{P}_{\smallsetminus\perp}(t_2) \text{ and } \exists \gamma < \beta \forall \gamma \leq \iota < \beta \;\; u \in U /\!/ S \cdot (T|_{[0,\iota)}) \\ & \text{iff} & u \in \mathcal{P}_{\smallsetminus\perp}(t_2) \text{ and } \exists \gamma < \beta \forall \gamma \leq \iota < \beta \;\; u \in (U /\!/ S) /\!/ T|_{[0,\iota)} \qquad \text{(ind. hyp.)} \\ & \text{iff} & u \in (U /\!/ S) /\!/ T \end{array}$$

□

Also the next lemma verifies an intuitive property of descendants.

Lemma 5.5.9 (monotonicity of descendants)
Let $\mathcal{R}$ be an ITRS, $S\colon s \twoheadrightarrow_{\mathcal{R}} t$ and $U, V \subseteq \mathcal{P}_{\smallsetminus\perp}(s)$. If $U \subseteq V$, then $U /\!/ S \subseteq V /\!/ S$.

Proof. Straightforward induction on the length of S. □

The definition of descendants of open reduction sequences is quite subtle which makes its use in proofs fairly cumbersome. The lemma below establishes an alternative characterisation which will turn out to be useful in later proofs.

Lemma 5.5.10 (descendants of open reductions)
Let $\mathcal{R}$ be an ITRS, $S\colon s \twoheadrightarrow^{\lambda}_{\mathcal{R}} t$ and $U \subseteq \mathcal{P}_{\smallsetminus\perp}(s)$, with λ a limit ordinal and $S = (t_\iota \to_{\pi_\iota, c_\iota} t_{\iota+1})_{\iota<\lambda}$. Then it holds that

$$\pi \in U /\!/ S \quad \textit{iff} \quad \textit{there is some } \beta < \lambda \textit{ with } \pi \in U /\!/ S|_{[0,\beta)} \textit{ and } \forall \beta \leq \iota < \lambda \;\; \pi_\iota \nleq \pi.$$

Proof. We first prove the "only if" direction. To this end, assume that $\pi \in U /\!/ S$. Hence, it holds that

$$\pi \in \mathcal{P}_{\smallsetminus\perp}(t) \text{ and there is some } \gamma_1 < \lambda \text{ such that } \pi \in U /\!/ S|_{[0,\iota)} \text{ for all } \gamma_1 \leq \iota < \lambda \tag{1}$$

Particularly, we have that $t(\pi) \neq \perp$. Applying Lemma 5.5.1 then yields that

$$\text{there is some } \gamma_2 < \lambda \text{ such that } \pi_\iota \nleq \pi \text{ for all } \gamma_2 \leq \iota < \lambda \tag{2}$$

Now take $\beta = \max\{\gamma_1, \gamma_2\}$. Then it holds that $\pi \in U /\!/ S|_{[0,\beta)}$ and that $\pi_\iota \nleq \pi$ for all $\beta \leq \iota < \lambda$ due to (1) and (2), respectively.

Next consider the converse direction of the statement: Let $\beta < \lambda$ be such that $\pi \in U /\!/ S|_{[0,\beta)}$ and $\pi_\iota \nleq \pi$ for all $\beta \leq \iota < \lambda$. We will show that $\pi \in U /\!/ S$ by proving the stronger statement that $\pi \in U /\!/ S|_{[0,\gamma)}$ for all $\beta \leq \gamma \leq \lambda$. We do this by induction on γ.

For $\gamma = \beta$, this is trivial. Let $\gamma = \gamma' + 1 > \beta$. Note that, by definition, $U /\!/ S|_{[0,\gamma)} = \big(U /\!/ S|_{[0,\gamma')}\big) /\!/ S|_{[\gamma',\gamma)}$. Hence, since for the γ'-th step we have, by assumption, $\pi_{\gamma'} \nleq \pi$ and for the preceding reduction we have, by induction hypothesis, that $\pi \in U /\!/ S|_{[0,\gamma')}$, we can conclude that $\pi \in U /\!/ S|_{[0,\gamma)}$.

Let $\gamma > \beta$ be a limit ordinal. By induction hypothesis, we have that $\pi \in U /\!/ S|_{[0,\iota)}$ for each $\beta \leq \iota < \gamma$. Particularly, this implies that $\pi \in \mathcal{P}_{\smallsetminus\perp}(t_\beta)$. Together with the assumption that $\pi_\iota \nleq \pi$ for all $\beta \leq \iota < \gamma$, this yields that $\pi \in \mathcal{P}_{\smallsetminus\perp}(t_\gamma)$ according to Lemma 5.5.1. Hence, $\pi \in U /\!/ S|_{[0,\gamma)}$. □

Similarly to Proposition 5.5.8, the proposition below confirms a property that we intuitively expect from a formalisation of the concept of descendants. It is also a much stronger variant of Lemma 5.5.9 and is the correspondent to Proposition 5.4.11.

Proposition 5.5.11 (pointwise definition of descendants)
Let $\mathcal{R}$ be an ITRS, $S\colon s \twoheadrightarrow_{\mathcal{R}} t$ and $U, V \subseteq \mathcal{P}_{\smallsetminus\perp}(s)$. Then it holds that $U /\!/ S = \bigcup_{u \in U} u /\!/ S$.

Proof. Let $S = (t_\iota \to_{\pi_\iota, c_\iota} t_{\iota+1})_{\iota<\alpha}$. For $\alpha = 0$ and $\alpha = 1$, the statement is trivially true. If $\alpha = \alpha' + 1 > 1$, then abbreviate $S|_{[0,\alpha')}$ and $S|_{[\alpha',\alpha)}$ by S_1 and S_2, respectively, and reason as follows:

$$\begin{aligned} U/\!/S = (U/\!/S_1)/\!/S_2 &\overset{IH}{=} \underbrace{(\bigcup_{u\in U} \overbrace{u/\!/S_1}^{V_u})}_{V}/\!/S_2 \overset{IH}{=} \bigcup_{u\in V} u/\!/S_2 \\ &= \bigcup_{u\in U} \bigcup_{v\in V_u} v/\!/S_2 \overset{IH}{=} \bigcup_{u\in U} V_u/\!/S_2 = \bigcup_{u\in U} (u/\!/S_1)/\!/S_2 = \bigcup_{u\in U} u/\!/S \end{aligned}$$

Let α be a limit ordinal. The "$\supseteq$" direction of the equation follows from Lemma 5.5.9. For the converse direction, assume that $\pi \in U/\!/S$. By Lemma 5.5.10, there is some $\beta < \alpha$ such that $\pi_\iota \nleq \pi$ for all $\beta \le \iota < \alpha$ and $\pi \in U/\!/S|_{[0,\beta)}$. Applying the induction hypothesis yields that $\pi \in \bigcup_{u\in U} u/\!/S|_{[0,\beta)}$, i.e. there is some $u^* \in U$ such that $\pi \in u^*/\!/S|_{[0,\beta)}$. By employing Lemma 5.5.10 again, we can conclude that $\pi \in u^*/\!/S$ and, therefore, that $\pi \in \bigcup_{u\in U} u/\!/S$. □

As we have mentioned, one can reasonably argue that the above proposition confirms a behaviour that one intuitively expects from descendants. In this light, the present formalisation of descendants is not desirable for weakly convergent reductions as it would violate Proposition 5.5.11. The reason is essentially the same as for MRS reductions (cf. Example 5.4.14). Another counterexample is the following one:

Example 5.5.12
Consider the system with the single rule $f(x) \to x$ and the weakly convergent reduction sequence

$$S\colon f^\omega \to f^\omega \to \dots f^\omega$$

which in each step contracts the redex at the root of f^ω. Consider the set $U = \{\varepsilon, 0, 0^2, 0^3, \dots\}$ of all occurrences in f^ω. Then the set of descendants of U by S is again U. Proposition 5.5.11 fails as, for each $\pi \in U$, the set of residuals of the singleton set $\{\pi\}$ by S is the empty set.

The following proposition corresponds to Proposition 5.4.12 from the MRS world.

Proposition 5.5.13 (uniqueness of descendants)
Let $\mathcal{R}$ be a left-linear ITRS, $S\colon s \twoheadrightarrow_{\mathcal{R}} t$ and $U, V \subseteq \mathcal{P}_{\setminus\bot}(s)$. If $U \cap V = \emptyset$, then $U/\!/S \cap V/\!/S = \emptyset$.

Proof. We will prove the contraposition of the statement. To this end, suppose that there is some occurrence $w \in U/\!/S \cap V/\!/S$. By Proposition 5.5.11, there are occurrences $u \in U$ and $v \in V$ such that $w \in u/\!/S \cap v/\!/S$. We will show by induction on the length of S that then $u = v$ and, therefore, $U \cap V \neq \emptyset$. If S is empty, then this is trivial. If S is of successor ordinal length, then this follows straightforwardly from the induction hypothesis. If S is open, then $u = v$ follows from the induction hypothesis using Lemma 5.5.10. □

As in the MRS model, Proposition 5.5.11 and Proposition 5.5.13 give rise to the notion of *ancestors* (cf. Remark 5.4.13).

The next proposition corresponds to Proposition 5.4.10.

Proposition 5.5.14 ((disjoint) residuals)
Let $\mathcal{R}$ be an almost orthogonal ITRS, $S\colon s \twoheadrightarrow_{\mathcal{R}} t$ and U a set of redex occurrences in s. Then the following holds:

(i) $U/\!/S$ is a set of redex occurrences in t.

(ii) If the occurrences in U are pairwise disjoint, then so are the occurrences in $U/\!/S$.

Proof. Let $S = (t_\iota \to_{\pi_\iota, c_\iota} t_{\iota+1})_{\iota<\alpha}$.

We will prove (i) by an induction on α. For α being 0, the statement is trivial, and, for α a successor ordinal, the statement follows straightforwardly from the induction hypothesis.

So assume that α is a limit ordinal and that $\pi \in U/\!/S$. From Lemma 5.5.10 we obtain that

$$\text{there is some } \beta < \alpha \text{ with } \pi \in U/\!/S|_{[0,\beta)} \text{ and } \pi_\iota \nleq \pi \text{ for all } \beta \le \iota < \alpha. \tag{1}$$

By applying the induction hypothesis, we get that π is a redex occurrence in t_β. Hence, there is some rule $l \to r \in R$ such that $t_\beta|_\pi$ is an instance of l.

We conclude this proof by proving the following stronger claim that implies (i):

$$\text{for all } \beta \le \gamma \le \alpha \qquad t_\gamma|_\pi \text{ is an instance of } l, \text{ and} \tag{2}$$

$$c_\iota|_\pi \text{ is an instance of } l \text{ for all } \beta \le \iota < \gamma \tag{3}$$

Then (i) is simply (2) for the case $\gamma = \alpha$.

We proceed by an induction on γ. For $\gamma = \beta$, part (2) of the claim has already been shown and (3) is vacuously true. Let $\gamma = \gamma' + 1 > \beta$. According to the induction hypothesis, (2) and (3) hold for γ'. Hence, it remains to be shown that both $t_\gamma|_\pi$ and $c_{\gamma'}|_\pi$ are instances of l. At first consider $c_{\gamma'}|_\pi$. Recall that $c_{\gamma'} = t_{\gamma'}[\bot]_{\pi_{\gamma'}}$. Assume that π and π'_γ are disjoint. Then $c_{\gamma'}|_\pi = t_{\gamma'}|_\pi$. Since, by induction hypothesis, $t_{\gamma'}|_\pi$ is an instance of l, so is $c_{\gamma'}|_\pi$. Next assume that π and $\pi_{\gamma'}$ are not disjoint. Because of (1), we then have that $\pi < \pi_{\gamma'}$, i.e. there is some non-empty π' with $\pi_{\gamma'} = \pi \cdot \pi'$. Since $\mathcal{R}$ is an overlay system, π' cannot be a position in the pattern of the redex $t_{\gamma'}|_\pi$ w.r.t. l. Therefore, also $c_{\gamma'}|_\pi$ is an instance of l. So in either case $c_{\gamma'}|_\pi$ is an instance of l. Since $c_{\gamma'} \le_\bot t_\gamma$, we can apply Corollary 4.4.14 to obtain that also $t_\gamma|_\pi$ is an instance of l.

Let $\gamma > \beta$ be a limit ordinal. Using Lemma 2.1.15, part (3) of the claim follows immediately from the induction hypothesis. Hence, $c_\iota|_\pi$ is an instance of l for all $\beta \le \iota < \gamma$. This and (1) implies that all terms in the set $T = \{c_\iota \mid \beta \le \iota < \gamma\}$ coincide in all occurrences in the set

$$P = \{\pi' \mid \pi' \le \pi\} \cup \{\pi \cdot \pi' \mid \pi \in \mathcal{P}_\Sigma(l)\}$$

P is obviously closed under prefixes. Therefore, we can apply Lemma 4.4.23 in order to obtain that $\bigsqcap^\bot T$ coincides with all terms in T in all occurrences in P. Since $\bigsqcap^\bot T \le_\bot t_\gamma$, this property carries over to t_γ. Consequently, also $t_\gamma|_\pi$ is an instance of l.

We also prove (ii) by induction on α. For α being 0, the statement is trivial, and, for α being a successor ordinal, the statement follows straightforwardly from the induction hypothesis. Let α be limit ordinal and suppose that (ii) does not hold. That is, there are two positions $u, v \in U/\!/S$ which are not disjoint. By definition, there are ordinals $\beta_1, \beta_2 < \alpha$ such that $u \in U/\!/S|_{[0,\iota)}$ for all $\beta_1 \le \iota < \alpha$, and $v \in U/\!/S|_{[0,\iota)}$ for all $\beta_2 \le \iota < \alpha$. Let $\beta = \max\{\beta_1, \beta_2\}$. Then we have that $u, v \in U/\!/S|_{[0,\beta)}$. This, however, contradicts the induction hypothesis which, in particular, states that $U/\!/S|_{[0,\beta)}$ is a set of pairwise disjoint redex occurrences. □

The first part of the proposition just proven is essential for the concept of complete developments that will be covered in Section 5.5.2. Similarly to the case of MRS reductions, sets of disjoint redex occurrences are of great importance mainly because they are preserved by strongly convergent reductions.

Next we want to establish an alternative characterisation of descendants based on labellings. This is a well-known technique that keeps track of descendants by labelling, in the initial term, the symbols at the positions that are to be tracked. In order to formalise this idea, we need to extend a given ITRS such that it can also deal with terms that contain labelled symbols:

Definition 5.5.15 (labelled ITRSs/terms)
Let $\mathcal{R} = (\Sigma, R)$ be a ITRS over Σ.

(i) The *labelled signature* Σ^l is defined as

$$\Sigma^l = \Sigma \cup \left\{ f^l \,\middle|\, f \in \Sigma \right\}$$

The arity of the function symbol f^l is the same as that of f. The symbols f^l are called *labelled*; the symbols $f \in \Sigma$ are called *unlabelled*. Terms over Σ^l are called *labelled terms*.

(ii) The labelled terms can be projected back to the original unlabelled ones by removing the labels. We, therefore, define the projection function $\|\cdot\| \colon \mathcal{T}^\infty(\Sigma^l_\perp, \mathcal{V}) \to \mathcal{T}^\infty(\Sigma_\perp, \mathcal{V})$ to this effect by setting

$$\begin{aligned} \|x\| &= x && \text{for all } x \in \mathcal{V}, \text{ and} \\ \|\perp\| &= \perp \\ \left\|f^l(t_1, \dots, t_k)\right\| &= \|f(t_1, \dots, t_k)\| = f(\|t_1\|, \dots, \|t_k\|) && \text{for all } f \in \Sigma^{(k)} \end{aligned}$$

(iii) The *labelled ITRS* $\mathcal{R}^l$ is defined as (Σ^l, R^l), where

$$R^l = \{ l \to r \mid \|l\| \to r \in R \}$$

(iv) For each rule $l \to r \in R^l$, we define its unlabelled original $\|l \to r\| \in R$ by setting $\|l \to r\| = \|l\| \to r$.

(v) Let $t \in \mathcal{T}^\infty(\Sigma_\perp, \mathcal{V})$ and $U \subseteq \mathcal{P}_{\not\perp}(t)$. The term $t^{(U)} \in \mathcal{T}^\infty(\Sigma^l_\perp, \mathcal{V})$ is defined by

$$t^{(U)}(\pi) = \begin{cases} t(\pi) & \text{if } \pi \notin U \\ t(\pi)^l & \text{if } \pi \in U \end{cases}$$

That is, $\left\|t^{(U)}\right\| = t$ and the labelled symbols in $t^{(U)}$ are exactly those at positions in U.

The key property which is needed in order to make the labelling approach work is that any reduction in a left-linear ITRS that starts in some term t can be lifted for any labelling t' of t to a unique equivalent reduction in the labelled ITRS that starts in t':

Proposition 5.5.16 (lifting of reduction sequences to labelled ITRS)
Let $\mathcal{R} = (\Sigma, R)$ be a left-linear ITRS, $S = (s_\iota \to_{\rho_\iota, \pi_\iota} s_{\iota+1})_{\iota<\alpha}$ a reduction sequence in $\mathcal{R}$ strongly converging to s_α, and $t_0 \in \mathcal{T}^\infty(\Sigma^l_\perp, \mathcal{V})$ a labelled term with $\|t_0\| = s_0$. Then there is a unique strongly continuous reduction sequence $T = (t_\iota \to_{\rho'_\iota, \pi_\iota} t_{\iota+1})_{\iota<\alpha}$ in $\mathcal{R}^l$ converging to t_α such that

(a) $\|t_\iota\| = s_\iota$, $\|\rho'_\iota\| = \rho_\iota$, for all $\iota < \alpha$, and

(b) $\|t_\alpha\| = s_\alpha$.

Proof. We prove this by an induction on α. For the case of α being zero, the statement is trivially true. For the case of α being a successor ordinal, the statement follows straightforwardly from the induction hypothesis (the argument is the same as for finite reduction sequences; e.g. consult [Ter03]).

Let α be a limit ordinal. By induction hypothesis, each proper prefix $S|_{[0,\gamma)}$ of S has a uniquely defined strongly convergent reduction sequence T_γ in $\mathcal{R}^l$ satisfying (a) and (b). Note that, for each two $\gamma \le \gamma' < \alpha$, it holds that $S|_{[0,\gamma)} \le S|_{[0,\gamma')}$. Consequently, also for the lifted reduction sequences, it holds that $T_\gamma \le T_{\gamma'}$. Define $T = \bigsqcup_{\iota<\alpha} T_\iota$. This is well-defined as the set $\{T_\iota \mid \iota < \alpha\}$ is directed. Therefore, $T_\gamma \le T$ holds for each $\gamma < \alpha$, and we can use the induction hypothesis to obtain part (a) of the proposition.

Next, we show that $\|t_\alpha\| \le_\perp s_\alpha$. To this end, let $\pi \in \mathcal{P}_{\not\perp}(\|t_\alpha\|)$. According to Corollary 4.4.14, we have to show that $\|t_\alpha\|(\pi) = s_\alpha(\pi)$. Let $\|t_\alpha\|(\pi) = f \in \Sigma \cup \mathcal{V}$. That is, either

$t_\alpha(\pi) = f$ or $t_\alpha(\pi) = f^l$. In either case, we can employ Lemma 5.5.1 to obtain some $\beta < \alpha$ such that $t_\beta(\pi) = f$ resp. $t_\beta(\pi) = f^l$ and $\pi_\iota \nleq \pi$ for all $\beta \le \iota < \alpha$. Since, by (a), $s_\beta = \|t_\beta\|$, we have in both cases that $s_\beta(\pi) = f$. By applying Lemma 5.5.1 again, we get that $s_\alpha(\pi) = f$, too.

Lastly, we show the converse inequality $s_\alpha \le_\bot \|t_\alpha\|$. For this purpose, let $\pi \in \mathcal{P}_{\smallsetminus\bot}(s_\alpha)$ and $s_\alpha(\pi) = f \in \Sigma \cup \mathcal{V}$. By Lemma 5.5.1, there is some $\beta < \alpha$ such that $s_\beta(\pi) = f$ and $\pi_\iota \nleq \pi$ for all $\beta \le \iota < \alpha$. Since, by (a), $s_\beta = \|t_\beta\|$, we have that $t_\beta(\pi) \in \{f, f^l\}$. Applying Lemma 5.5.1 again then yields that $t_\alpha(\pi) \in \{f, f^l\}$ and, therefore, $\|t_\alpha\|(\pi) = f$.

By combining both inequalities, we obtain that $s_\alpha = \|t_\alpha\|$ due to the antisymmetry of $\le_\bot$. This proves part (b) of the proposition. □

Having this, we can establish an alternative characterisation of descendants using labellings:

Proposition 5.5.17 (alternative characterisation of descendants)
Let $\mathcal{R}$ be a left-linear ITRS over Σ, $S\colon s_0 \twoheadrightarrow_{\mathcal{R}} s_\alpha$, and $U \subseteq \mathcal{P}_{\smallsetminus\bot}(s_0)$. Following Proposition 5.5.16, let $T\colon t_0 \twoheadrightarrow_{\mathcal{R}} t_\alpha$ be the unique lifting of S to $\mathcal{R}^l$ starting with the term $t_0 = s_0^{(U)}$. Then it holds that $t_\alpha = s_\alpha^{(U/\!/S)}$. That is, for all $\pi \in \mathcal{P}_{\smallsetminus\bot}(s_\alpha)$, it holds that $t_\alpha(\pi)$ is labelled iff $\pi \in U/\!/S$.

Proof. Let $S = (s_\iota \to_{\pi_\iota} s_{\iota+1})_{\iota<\alpha}$ and $T = (t_\iota \to_{\pi_\iota} t_{\iota+1})_{\iota<\alpha}$ We show the statement by an induction on the length α of S. If $\alpha = 0$, then the statement is trivially true. If α is a successor ordinal, then a straightforward argument shows that the statement follows from the induction hypothesis. Here the restriction to left-linear systems is vital.

Let α be a limit ordinal and let $\pi \in \mathcal{P}_{\smallsetminus\bot}(s_\alpha)$. We can then reason as follows:

$$\begin{aligned}
t_\alpha(\pi) \text{ is labelled} \quad &\text{iff} \quad \exists\beta < \alpha\colon\ t_\beta(\pi) \text{ is labelled and } \forall\beta \le \iota < \alpha\colon\ \pi_\iota \nleq \pi && \text{(Lem. 5.5.1)}\\
&\text{iff} \quad \pi \in U/\!/S|_{[0,\beta)} \text{ and } \forall\beta \le \iota < \alpha\colon\ \pi_\iota \nleq \pi && \text{(ind. hyp.)}\\
&\text{iff} \quad \pi \in U/\!/S && \text{(Lem. 5.5.10)}
\end{aligned}$$

□

5.5.2 Complete Developments

We have already come across complete developments in Section 5.4.2 where they were presented for strongly convergent MRS reductions. The intuition of complete development is that of a strongly convergent reduction that simulates the contraction of an entire set of redex occurrences. That was the motivation for considering descendants and, in particular, residuals in the previous section.

We have seen in Section 5.4.2 that complete developments exists for a wide, yet still restricted, range of sets of redex occurrences. In the present setting of PRS reductions, orthogonal systems do always admit complete developments – for any set of redex occurrences. Additionally, we can establish the same properties that we already have in the MRS case: The final terms of complete developments are uniquely determined, descendants by compete developments are well-defined, and the Infinitary Strip Lemma holds.

The definition of complete developments can be copied verbatimly from the MRS setting (cf. Definition 5.4.8):

Definition 5.5.18 ((complete) development)
Let $\mathcal{R}$ be an ITRS, s a partial term in $\mathcal{R}$, and U a set of pairwise non-conflicting redex occurrences in s.

(i) A *development* of U in s is a strongly convergent reduction $S\colon s \twoheadrightarrow^\alpha_{\mathcal{R}} t$ in which each reduction step $\varphi_\iota\colon t_\iota \to_{\pi_\iota} t_{\iota+1}$ contracts a redex at $\pi_\iota \in U/\!/S|_{[0,\iota)}$ for $\iota < \alpha$.

(ii) A development $S\colon s \twoheadrightarrow t$ of U in s is called *complete*, denoted $S\colon s \twoheadrightarrow_U t$, if $U/\!/S = \emptyset$.

(iii) A (complete) development $S\colon s \twoheadrightarrow t$ is called *total* if it is a total strongly convergent reduction.

Again, the restriction to non-conflicting redex occurrences is essential in order guarantee that the redex occurrences are independent from each other. Moreover, in conjunction with left-linearity this ensures that the descendants of a set of non-conflicting redex occurrences is again a set of non-conflicting redex occurrences:

Fact 5.5.19 (non-conflicting residuals)
Let $\mathcal{R}$ be a left-linear ITRS, s a partial term in $\mathcal{R}$, U a set of pairwise non-conflicting redex occurrences in s, and $S\colon s \twoheadrightarrow_U t$ a development of U in s. Then also $U/\!/S$ is a set of pairwise non-conflicting redex occurrences.

It is relatively easy to show that complete developments do always exists in the PRS setting – no matter which set of redex occurrences is considered. The reason for this is that the PRS induced by an ITRS is always complete according to Proposition 3.2.4 which implies, by Fact 3.2.7, that every strongly continuous reduction is also strongly convergent. This means that as long as there are (residuals of) redex occurrences left after an incomplete development, one can extend this development arbitrarily by contracting some of the remaining redex occurrences. The only thing that remains to be shown is that one can devise a reduction strategy which eventually contracts (all residuals of) all redexes. The proposition below shows that a parallel-outermost reduction strategy will always yield a complete development in an orthogonal system.

Proposition 5.5.20 (existence of complete developments)
Let $\mathcal{R}$ be an orthogonal ITRS, t a partial term in $\mathcal{R}$, and U a set of redex occurrences in t. Then U has a complete development in t.

Proof. Let $t_0 = t$, $U_0 = U$ and V_0 the set of outermost occurrences in U_0, i.e. the minimal elements w.r.t. the prefix order. Furthermore, let $S_0\colon t_0 \twoheadrightarrow_{V_0} t_1$ be some complete development of V_0 in t_0. S_0 can be constructed by contracting the redex occurrences in V_0 in a left-to-right order. This can be continued for each $i < \omega$ by taking $U_{i+1} = U_i/\!/S_i$, V_{i+1} the outermost redex occurrences in U_{i+1} and $S_{i+1}\colon t_{i+1} \twoheadrightarrow_{V_{i+1}} t_{i+2}$ some complete development of V_{i+1} in t_{i+1}.

Note that then, by iterating Proposition 5.5.8, it holds that

$$U/\!/S_0 \cdot \ldots \cdot S_{n-1} = U_n \qquad \text{for all } n < \omega \tag{1}$$

If there is some $n < \omega$ for which $U_n = \emptyset$, then $S_0 \cdot \ldots \cdot S_{n-1}$ is a complete development of U according to (1).

If this is not the case, consider the reduction sequence $S = \prod_{i<\omega} S_i$, i.e. the concatenation of all 'S_i's. We claim that S is a complete development of U. Suppose that this is not the case, i.e. $U/\!/S \neq \emptyset$. Hence, there is some $u \in U/\!/S$. Since all 'U_i's are non-empty, so are the 'V_i's. Consequently, all 'S_i's are non-empty reduction sequences which implies that S is a reduction sequence of limit ordinal length, say λ. Therefore, we can apply Lemma 5.5.10 to infer from $u \in U/\!/S$ that there is some $\alpha < \lambda$ such that $u \in U/\!/S|_{[0,\alpha)}$ and all reduction steps beyond α do not take place at u or above. This is not possible due to the parallel-outermost reduction strategy that S adheres. □

There are several techniques to show that in the MRS model of transfinite reductions the final terms of complete developments are uniquely determined by the starting term and the set of redex occurrences. One of these approaches, introduced in [KdV03], uses so-called *paths*. Paths are constructed such that they, starting from the root, run through the initial term t of the complete development, and whenever a redex occurrence of the development is hit, the path jumps to the root of the right-hand side of the corresponding rule and jumps back to the term t when it reaches a variable in the right-hand side. Figure 5.11a illustrates

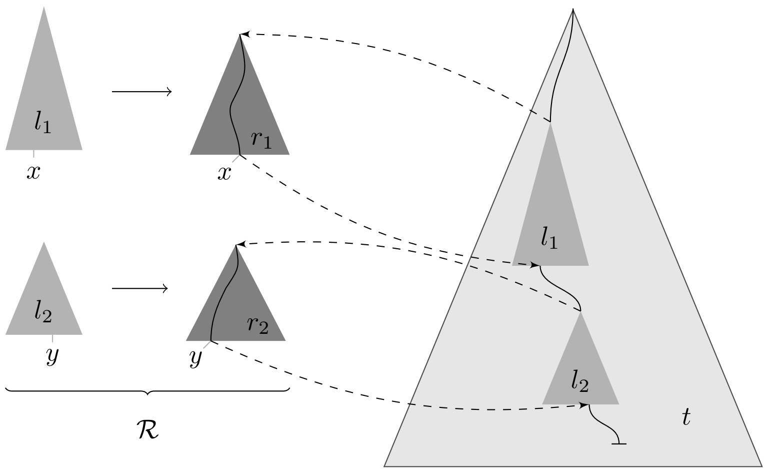

(a) Constructing a path in a term.

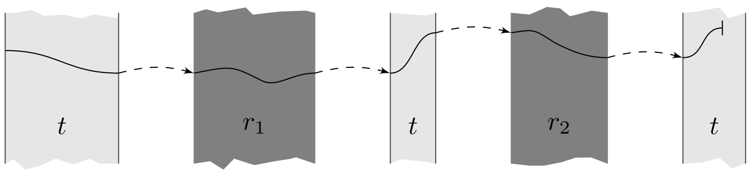

(b) The constructed path.

Figure 5.11: A path.

the this idea. It shows a path of a term t that encounters two redex occurrences of the complete development. As soon as such a redex occurrence is encountered, the path jumps to the right-hand side of the corresponding rule as indicated by the dashed arrows. Then the path runs through the right-hand side. When a variable is encountered, the path jumps back to the position of the term t that matches the variable. This jump is again indicated by a dashed arrow. The path that is obtained by this construction is shown in Figure 5.11b. With the collection of the thus obtained paths one can then construct the final term of the complete development. This technique can also be applied in the present setting.

At first we need to formalise the concept of a path that we have intuitively described above.

Definition 5.5.21 (path, [Ter03])
Let $\mathcal{R}$ be an orthogonal ITRS, t a partial term in $\mathcal{R}$, and U a set of redex occurrence in t. A *path* of t w.r.t. U and $\mathcal{R}$ is a sequence of length at most ω containing so-called *nodes* and *edges* in an alternating manner like this:

$$n_0 \cdot e_0 \cdot n_1 \cdot e_1 \cdot n_2 \cdot e_2 \cdot \cdots$$

where the 'n_i's are nodes and the 'e_i's are edges. A node is either a pair of the form (t, π) with $\pi \in \mathcal{P}(t)$ or a triple of the form (r, π, u) with r the right-hand side of a rule in $\mathcal{R}$, $\pi \in \mathcal{P}(r)$, and $u \in U$. Edges are denoted by arrows $\to$. Both edges and nodes might be labelled by elements in $\Sigma_\perp \cup \mathcal{V}$ and $\mathbb{N}$, respectively. We write paths as the one sketched above as

$$n_0 \to n_1 \to n_2 \to \cdots$$

or, when explicitly indicating labels, as

$$n_0{}^{l_0} \xrightarrow{l_1} n_1{}^{l_2} \xrightarrow{l_3} n_2{}^{l_4} \xrightarrow{l_5} \cdots$$

where empty labels are explicitly given by the symbol $\varnothing$. If a path has a segment of the form $n \to n'$, then we say there is an edge from n to n' or that n has an outgoing edge to n'.

Every path starts with the node (t, ε) and is either infinitely long or ends with a node. For each node n having an outgoing edge to a node n', the following must hold:

(1) If n is of the form (t, π), then

  (a) $n' = (t, \pi \cdot i)$ and the edge is labelled by i, with $0 \leq i < \mathsf{ar}_t(\pi)$, whenever $\pi \notin U$, and

  (b) $n' = (r, \varepsilon, \pi)$ and the edge is unlabelled, with $l \to r \in R$ the rule for the redex $t|_\pi$, whenever $\pi \in U$.

(2) If n is of the form (r, π, u), then

  (a) $n' = (r, \pi \cdot i, u)$ and the edge is labelled by i, with $0 \leq i < \mathsf{ar}_r(\pi)$, whenever $r|_\pi$ is not a variable, and

  (b) $n' = (t, u \cdot \pi')$ and the edge is unlabelled, with $l \to r \in R$ the rule for the redex $t|_u$ and π' the unique occurrence of $r|_\pi$ in l, whenever $r|_\pi$ is a variable.

Additionally, the nodes of a path are supposed to be labelled in the following way:

(3) A node of the form (t, π) is unlabelled if $\pi \in U$ and is labelled by $t(\pi)$ otherwise.

(4) A node of the form (r, π, u) is unlabelled if $r|_\pi$ is a variable and labelled by $r(\pi)$ otherwise.

Remark 5.5.22. The above definition is actually a coinductive one. This is necessary to also define paths of infinite length. Also in [KdV03] paths are considered to be possibly infinite, although they are defined inductively and are, therefore, finite.

The purpose of nodes of the form (t, π) and (r, π, u), respectively, is that they encode that the path is currently at position π in the term t resp. r. The additional component u provides the information that the path jumped to the right-hand side r from the redex $t|_u$. The labellings of the nodes represent the symbols at the current location of the path, unless it is a redex occurrence in the main term or a variable occurrence in a right-hand side. The labellings of the edges provide information on how the path moves through the terms: A labelling i represents a move along the i-th edge in the term tree from the current location whereas an empty labelling indicates a jump from or to a right-hand side.

Returning to the schematic example illustrated in Figure 5.11, we can observe how the construction of a path is carried out: The path starts with a segment in the term t. This segment is entirely regulated by the rule (1a); all its edges and nodes are labelled according to (1a) and (3). The jump to the right-hand side r_1 following that initial segment is justified by rule (1b). This jump consists of a node (t, u_1), unlabelled according to (3), corresponding to the redex occurrence u_1, and an unlabelled edge to the node (r_1, ε, u_1), corresponding to the root of the right-hand side r_1. The segment of the path that runs through the right-hand side r_1 is subject to rule (2a); again all its nodes and edges are labelled, now, according to (2a) and (4). As soon as a variable is reached in the right-hand side, in the schematic example it is the variable x, a jump to the main term t is performed as required by rule (2b). This jump consists of a node (r_1, π, u_1), unlabelled according to (4), where π is the current position in r_1, i.e. the variable occurrence, and an unlabelled edge to the node $(t, u_1 \cdot \pi')$. The position π' is the occurrence of the variable x in the left-hand side. As we only consider orthogonal systems, this occurrence is unique. Afterwards, the same behaviour is repeated: A segment in t is followed by a jump to a segment in the right-hand side r_2 which is in turn followed by a jump back to a final segment in t.

Note that paths do not need to be maximal. As indicated in the schematic example, the path ends somewhere within the main term, i.e. not necessarily at a constant symbol or a variable. What the example does not show, but which is obvious from the definition, is that paths can also terminate within a right-hand side. A jump back to the main term is only required if variable is encountered.

The purpose of the concept of paths is to simulate the contraction of all redexes of the complete development in a locally restricted manner, i.e. only along some branch of the term tree. This locality will keep the proofs more concise and makes them easier to understand once we have grasped the idea behind paths. The strategy to prove our conjecture of uniquely determined final terms is to show that paths can be used to define a term and that a contraction of a redex of the complete development preserves a property of the collection of all paths which ensures that the induced term remains invariant. Then we only have to observe that the induced term of paths in a term with no redexes (in U) is the term itself.

The following fact is obvious from the definition of a path.

Fact 5.5.23
Let $\mathcal{R}$ be an orthogonal ITRS, t a partial term in $\mathcal{R}$, and U a set of redex occurrences in t.

(i) An edge in a path of t w.r.t. U and $\mathcal{R}$ is unlabelled iff the preceeding node is unlabelled.

(ii) Any prefix of a path of t w.r.t. U and $\mathcal{R}$ that ends in a node is also a path of t w.r.t. U and $\mathcal{R}$.

The part of the information encoded in paths that is necessary in order to define the final term of the complete development is contained solely in the labels of their nodes and edges. The inner structure of nodes is only necessary for the definition of paths. The following notion of traces defines projections to the labels of paths:

Definition 5.5.24 (trace)
Let $\mathcal{R}$ be an orthogonal ITRS, t a partial term in $\mathcal{R}$, and U a set of redex occurrences in t.

(i) Let Π be a path of t w.r.t. U and $\mathcal{R}$. The *trace* of Π, denoted $\mathsf{tr}(\Pi)$, is the projection of Π to the labelling of its nodes and edges ignoring empty labels and the node label $\bot$.

(ii) $\mathcal{P}(t,U,\mathcal{R})$ is used to denote the set of all paths of t w.r.t. U and $\mathcal{R}$ that end in a labelled node, or are infinite but have a finite trace. The set of traces of paths in $\mathcal{P}(t,U,\mathcal{R})$ is denoted by $\mathcal{T}r(t,U,\mathcal{R})$.

(iii) U is said to have *finite jumps* if, for each infinite path $\Pi \in \mathcal{P}(t,U,\mathcal{R})$, also $\mathsf{tr}(\Pi)$ is infinite.

Remark 5.5.25. By Fact 5.5.23, the trace of a path is a sequence alternating between elements in $\Sigma \cup \mathcal{V}$ and $\mathbb{N}$, which, if non-empty, starts with an element in $\Sigma \cup \mathcal{V}$. Moreover, by definition, $\mathcal{T}r(t,U,\mathcal{R})$ is a set of finite traces of paths of t w.r.t. U and $\mathcal{R}$.

Example 5.5.26
Consider the term $t = g(f(g(h(\bot))))$ and the TRS $\mathcal{R}$ consisting of the two rules

$$f(x) \to h(x), \qquad h(x) \to x.$$

Furthermore, let U be the set of all redex occurrences in t, viz. $U = \{0, 0^3\}$. The following path Π is a path of t w.r.t. U and $\mathcal{R}$:

$$(t,\varepsilon)^g \xrightarrow{0} (t,0)^\varnothing \xrightarrow{\varnothing} (r_1,\varepsilon,0)^h \xrightarrow{0} (r_1,0,0)^\varnothing \xrightarrow{\varnothing} (t,0^2)^g \xrightarrow{0} (t,0^3)^\varnothing \xrightarrow{\varnothing} (r_2,\varepsilon,0^3)^\varnothing \xrightarrow{\varnothing} (t,0^4)^\bot$$

As a matter of fact, Π is the greatest path of t. Hence, according to Fact 5.5.23, the set of all prefixes of Π ending in a node is the set of all paths of t. Note that since Π itself ends in a labelled node, it is contained in $\mathcal{P}(t, U, \mathcal{R})$. The trace $\mathsf{tr}(\Pi)$ of Π is the sequence

$$g \cdot 0 \cdot h \cdot 0 \cdot g \cdot 0$$

Now consider the term $t' = g(f(g(h^\omega)))$ and the set U' of all its redexes, viz. $U' = \{0\} \cup \{0^3, 0^4, \dots\}$. Then the following path Π' is a path of t' w.r.t. U' and $\mathcal{R}$:

$$\begin{aligned}(t', \varepsilon)^g \overset{0}{\to} (t', 0)^\varnothing \overset{\varnothing}{\to} (r, \varepsilon, 0)^h \overset{0}{\to} (r, 0, 0)^\varnothing \overset{\varnothing}{\to} (t', 0^2)^g \overset{0}{\to} (t', 0^3)^\varnothing \overset{\varnothing}{\to} (r, \varepsilon, 0^3)^\varnothing \overset{\varnothing}{\to} (t', 0^4)^\varnothing& \\ \overset{\varnothing}{\to} (r, \varepsilon, 0^4)^\varnothing \overset{\varnothing}{\to} (t', 0^5)^\varnothing& \\ \overset{\varnothing}{\to} \dots&\end{aligned}$$

Π' is the greatest path of t'. The trace $\mathsf{tr}(\Pi')$ of Π' is the sequence

$$g \cdot 0 \cdot h \cdot 0 \cdot g \cdot 0$$

Since Π' is infinitely long but has a finite trace, it is contained in $\mathcal{P}(t', U, \mathcal{R})$.

The lemma below shows that there is a one-to-one correspondence between paths in $\mathcal{P}(t, U, \mathcal{R})$ and their traces in $\mathcal{T}r(t, U, \mathcal{R})$.

Lemma 5.5.27 ($\mathsf{tr}(\cdot)$ is a bijection)
Let $\mathcal{R}$ be an orthogonal ITRS, t a partial term in $\mathcal{R}$, and U a set of redex occurrences in t. $\mathsf{tr}(\cdot)$ is a bijection from $\mathcal{P}(t, U, \mathcal{R})$ to $\mathcal{T}r(t, U, \mathcal{R})$.

Proof. By definition, $\mathsf{tr}(\cdot)$ is surjective. Let Π_1, Π_2 be two paths having the same trace. We will show that then $\Pi_1 = \Pi_2$ by an induction on the length of the common trace.

Let $\mathsf{tr}(\Pi_1) = \varepsilon$. Following Fact 5.5.23, there are two different cases: The first case is that $\Pi_1 = \Pi \cdot (t, \pi)^\perp$, where the prefix Π corresponds to a finite maximal collapsing tower $(u_i)_{i \le \alpha}$ starting at the root of t or Π is empty if such a collapsing tower does not exists. If the collapsing tower exists, then

$$\Pi = (t, u_0)^\varnothing \overset{\varnothing}{\to} (r_0, \varepsilon, u_0)^\varnothing \overset{\varnothing}{\to} (t, u_1)^\varnothing \overset{\varnothing}{\to} (r_1, \varepsilon, u_1)^\varnothing \overset{\varnothing}{\to} \dots \overset{\varnothing}{\to} (t, u_\alpha)^\varnothing \overset{\varnothing}{\to}$$

But then also Π_2 starts with the prefix $\Pi \cdot (t, \pi)$ due to the uniqueness of the collapsing tower and the involved rules. Then $\Pi_1 = \Pi_2$ follows easily. If Π is empty, $\Pi_1 = \Pi_2$ follows immediately.

The second case is that Π_1 is infinite. Then there is an infinite collapsing tower $(u_i)_{i<\omega}$ starting at the root of t. Hence,

$$\Pi_1 = (t, u_0)^\varnothing \overset{\varnothing}{\to} (r_0, \varepsilon, u_0)^\varnothing \overset{\varnothing}{\to} (t, u_1)^\varnothing \overset{\varnothing}{\to} (r_1, \varepsilon, u_1)^\varnothing \overset{\varnothing}{\to} \dots$$

$\Pi_1 = \Pi_2$ follows from the uniqueness of the infinite collapsing tower.

At first glance one might additionally find a third case where $\Pi_1 = \Pi \cdot (t, \pi)^\varnothing \overset{\varnothing}{\to} (r, \varepsilon, \pi)^\perp$ with Π a prefix corresponding to a collapsing tower as in the first case. However, this is not possible as it would require the occurrence of a $\perp$ in a rule.

Let $\mathsf{tr}(\Pi_1) = f$. Then there are two cases: Either $\Pi_1 = \Pi \cdot (t, \pi)^f$ or $\Pi_1 = \Pi \cdot (t, \pi)^\varnothing \overset{\varnothing}{\to} (r, \varepsilon, \pi)^f$, where the prefix Π corresponds to a finite maximal collapsing tower $(u_i)_{i \le \alpha}$ starting at the root of t or Π is empty if such a collapsing tower does not exists. The argument is analogous to the argument employed for the first case of the induction base above.

Finally, we consider the induction step. Hence, there are the two cases: Either $\mathsf{tr}(\Pi_1) = T \cdot f \cdot i$ or $\mathsf{tr}(\Pi_1) = T \cdot f \cdot i \cdot g$. For both cases, the induction hypothesis can be invoked by taking two prefixes Π_1' and Π_2' of Π_1 and Π_2, respectively, which both have the trace T and, therefore, are equal according to the induction hypothesis. The argument that the remaining suffixes of Π_1 and Π_2 are equal is then analogous to the argument for two base cases. ∎

As mentioned above, the traces of paths contain all information necessary to define a term which we will later identify to be the final term of the corresponding complete development. The following definition explains how such a term, called a *matching term*, is determined:

Definition 5.5.28 (matching term)
Let $\mathcal{R}$ be an orthogonal ITRS, t a partial term in $\mathcal{R}$, and U a set of redex occurrences in t.

(i) The *position* of a trace $T \in \mathcal{T}r(t, U, \mathcal{R})$, denoted $\mathsf{pos}(T)$, is the subsequence of T containing only the edge labels. The set of all positions of traces in $\mathcal{T}r(t, U, \mathcal{R})$ is denoted $\mathcal{PT}r(t, U, \mathcal{R})$.

(ii) The *symbol* of a trace $T \in \mathcal{T}r(t, U, \mathcal{R})$, denoted $\mathsf{sym}(T)$, is f if T ends in a node label f, and is $\perp$ otherwise, i.e. whenever T is infinite, empty, or ends in an edge label.

(iii) A term t' is said to *match* $\mathcal{T}r(t, U, \mathcal{R})$ if, for all traces $T \in \mathcal{T}r(t, U, \mathcal{R})$, it holds that $t'(\mathsf{pos}(T)) = \mathsf{sym}(T)$.

Returning to the definition of paths, one can see that the label of a node is the symbol of the "current" position in a term. Similarly, the label of an edge says which edge in the term tree was taken at that point in the construction of the path. Hence, by projecting to the edge labels, we obtain the "history" of the path or in other words its position. In the same way we obtain the symbol of that node by taking the label of the last node of the path, provided the corresponding path ended in a non-$\perp$-labelled node. In the other case that the trace does not end in a node label, the corresponding path ended in a node labelled $\perp$ or was infinite. As we will see, infinite paths with finite traces correspond to infinite collapsing towers, which in turn yield volatile positions within the complete development. Eventually, these volatile positions will also give rise to $\perp$ subterms.

The following lemma shows that there is also a one-to-one correspondence between the traces in $\mathcal{T}r(t, U, \mathcal{R})$ an their positions in $\mathcal{PT}r(t, U, \mathcal{R})$:

Lemma 5.5.29 ($\mathsf{pos}(\cdot)$ is a bijection)
Let $\mathcal{R}$ be an orthogonal ITRS, t a partial term in $\mathcal{R}$ and U a set of redex occurrences in t. $\mathsf{pos}(\cdot)$ is a bijection from $\mathcal{T}r(t, U, \mathcal{R})$ to $\mathcal{PT}r(t, U, \mathcal{R})$.

Proof. An argument similar to the one for Lemma 5.5.27 can be given in order to show that the composition $\mathsf{pos}(\cdot) \circ \mathsf{tr}(\cdot)$ is a bijection. Together with the bijectivity of $\mathsf{tr}(\cdot)$, according to Lemma 5.5.27, this yields the bijectivity of $\mathsf{pos}(\cdot)$. □

Having this lemma, the following proposition is an easy consequence of the definition of matching terms. It shows that matching terms are uniquely determined:

Proposition 5.5.30 (unique matching term)
Let $\mathcal{R}$ be an orthogonal ITRS, t a partial term in $\mathcal{R}$, and U a set of redex occurrences in t. Then there is a unique term, denoted $\mathcal{F}(t, U, \mathcal{R})$, that matches $\mathcal{T}r(t, U, \mathcal{R})$.

Proof. Define the mapping $\varphi\colon \mathcal{PT}r(t, U, \mathcal{R}) \to \Sigma_\perp \cup \mathcal{V}$ by setting $\varphi(\mathsf{pos}(T)) = \mathsf{sym}(T)$ for each trace $T \in \mathcal{T}r(t, U, \mathcal{R})$. By Lemma 5.5.29, φ is well-defined. Moreover, it is easy to see from the definition of paths, that $\mathcal{PT}r(t, U, \mathcal{R})$ is closed under prefixes and that φ respects the arity of the symbols. Following Remark 2.3.17, φ uniquely determines a term s. By construction, s matches $\mathcal{T}r(t, U, \mathcal{R})$. Moreover, any other term s' matching $\mathcal{T}r(t, U, \mathcal{R})$ must satisfy $s'(\pi) = \varphi(\pi)$ for all $\pi \in \mathcal{PT}r(t, U, \mathcal{R})$, which implies that $s' = s$. □

It is also obvious that the matching term of a term t w.r.t. an empty set of redex occurrences is the term t itself.

Lemma 5.5.31 (matching term w.r.t. empty redex set)
For any ITRS $\mathcal{R}$ and any partial term t in $\mathcal{R}$, it holds that $\mathcal{F}(t, \emptyset, \mathcal{R}) = t$.

Proof. Straightforward. □

Remark 5.5.32. Now it only remains to be shown that the matching term stays invariant during a development, i.e. that, for each development $S\colon t \twoheadrightarrow t'$ of U, the matching terms $\mathcal{F}(t,U,\mathcal{R})$ and $\mathcal{F}(t',U/\!\!/S,\mathcal{R})$ coincide. Since the matching term $\mathcal{F}(t,U,\mathcal{R})$ only depends on the set $\mathcal{T}r(t,U,\mathcal{R})$ of traces, it is sufficient to show that $\mathcal{T}r(t,U,\mathcal{R})$ and $\mathcal{T}r(t',U/\!\!/S,\mathcal{R})$ coincide. The key observation is that in each step $s \to s'$ in a development the paths of s' differ from the paths of s only in that they might omit some jumps. This can be seen in Figure 5.11a: In a step $s \to s'$ of a development, (some residual of) some redex occurrence in U is contracted. In the picture this corresponds to removing the pattern, say l_1, of the redex and replacing it by the corresponding right-hand side of the rule r_1. One can see that, except for the jump to and from the right-hand side r_1 the path remains the same.

In order to establish the above observation formally, we need a means to simulate reduction steps in a development directly as an operation on paths. The following definition provides a tool for this.

Definition 5.5.33 (position and prefix of a path)
Let $\mathcal{R}$ be an orthogonal ITRS, t a partial term in $\mathcal{R}$, U a set of redex occurrences in t, and $\Pi \in \mathcal{P}(t,U,\mathcal{R})$.

(i) Π is said to contain an occurrence $\pi \in \mathcal{P}(t)$ if it contains the node (t,π).

(ii) Let $u \in U$. The prefix of Π by u, denoted $\Pi^{(u)}$, is defined as Π whenever Π does not contain u and otherwise as the unique prefix Π_1 such that $\Pi = \Pi_1 \cdot (t,u) \cdot \Pi_2$ for some Π_2.

Remark 5.5.34. It is obvious from the definition that each prefix $\Pi^{(u)}$ of a path $\Pi \in \mathcal{P}(t,U,\mathcal{R})$ by an occurrence u is the maximal prefix of Π, that does not contain positions that are extensions of u (including u itself). Equivalently, $\Pi^{(u)}$ is the maximal prefix of Π that only contains proper prefixes of u.

The following lemma is the key step towards proving the invariance of matching terms in developments. It formalises the observation described in Remark 5.5.32.

Lemma 5.5.35 (preservation of traces)
Let $\mathcal{R}$ be an orthogonal ITRS, t a partial term in $\mathcal{R}$, U a set of redex occurrences in t, and $S\colon t \twoheadrightarrow t'$ a development of U in t. There is a surjective mapping $\vartheta_S\colon \mathcal{P}(t,U,\mathcal{R}) \to \mathcal{P}(t',U/\!\!/S,\mathcal{R})$ such that $\mathsf{tr}(\Pi) = \mathsf{tr}(\vartheta_S(\Pi))$ for all $\Pi \in \mathcal{P}(t,U,\mathcal{R})$.

Proof. Let $S = (t_\iota \to_{\pi_\iota,c_\iota} t_{\iota+1})_{\iota<\alpha}$. We prove the statement by an induction on α.

If $\alpha = 0$, then the statement is trivially true.

Suppose that α is a successor ordinal $\beta+1$. Let $T\colon t_0 \twoheadrightarrow^\beta t_\beta$ be the prefix of S of length β and $\varphi_\beta\colon t_\beta \to_{\pi_\beta} t_\alpha$ the last step of S, i.e. $S = T \cdot \varphi_\beta$. By the induction hypothesis, there is a surjective mapping $\vartheta_T\colon \mathcal{P}(t,U,\mathcal{R}) \to \mathcal{P}(t_\beta,U',\mathcal{R})$, where $U' = U/\!\!/T$, satisfying $\mathsf{tr}(\Pi) = \mathsf{tr}(\vartheta_T(\Pi))$ for all $\Pi \in \mathcal{P}(t,U,\mathcal{R})$. By a careful case analysis (as done in [KS06]), one can show that there is a surjective mapping $\vartheta\colon \mathcal{P}(t_\beta,U',\mathcal{R}) \to \mathcal{P}(t_\alpha,U'',\mathcal{R})$, where $U'' = U'/\!\!/\varphi_\beta = U/\!\!/S$, satisfying $\mathsf{tr}(\Pi) = \mathsf{tr}(\vartheta(\Pi))$ for all $\Pi \in \mathcal{P}(t_\beta,U',\mathcal{R})$. Hence, the composition $\vartheta_S = \vartheta \circ \vartheta_T$ is a surjective mapping from $\mathcal{P}(t,U,\mathcal{R})$ to $\mathcal{P}(t',U/\!\!/S,\mathcal{R})$ and satisfies $\mathsf{tr}(\Pi) = \mathsf{tr}(\vartheta_S(\Pi))$ for all $\Pi \in \mathcal{P}(t,U,\mathcal{R})$.

Let α be a limit ordinal. By induction hypothesis, there is a surjective mapping ϑ_T for each proper prefix T of S that satisfies $\mathsf{tr}(\Pi) = \mathsf{tr}(\vartheta_T(\Pi))$ for all $\Pi \in \mathcal{P}(t,U,\mathcal{R})$. Let $\Pi \in \mathcal{P}(t,U,\mathcal{R})$ and $\Pi_\iota = \vartheta_{S|_{[0,\iota)}}(\Pi)$ for each $\iota < \alpha$. We define $\vartheta_S(\Pi)$ as follows:

$$\vartheta_S(\Pi) = \liminf_{\iota\to\alpha} \Pi_\iota^{(\pi_\iota)}$$

Note that the prefixes $\Pi_\iota^{(\pi_\iota)}$ are the respective maximal prefixes of Π_ι that do not contain any extension of π_ι. It is easy to see that ϑ_S is surjective and satisfies $\mathsf{tr}(\Pi) = \mathsf{tr}(\vartheta_S(\Pi))$ for all $\Pi \in \mathcal{P}(t, U, \mathcal{R})$. □

The above lemma effectively establishes the invariance of matching terms during a development. Together with Lemma 5.5.31 this implies the uniqueness of final terms of complete developments of the same redex occurrences. As a corollary from this, we obtain that descendants are also unique among all complete developments:

Proposition 5.5.36 (final term and descendants of complete developments) *Let $\mathcal{R}$ be an orthogonal ITRS, t a partial term in $\mathcal{R}$, and U a set of redex occurrences in t. Then the following holds:*

(i) Each complete development of U in t has the final term $\mathcal{F}(t, U, \mathcal{R})$.

(ii) For each set $V \subseteq \mathcal{P}_{\setminus\bot}(t)$ and two complete developments S and T of U in t, respectively, it holds that $V/\!/S = V/\!/T$.

Proof. (i) Let $S\colon t \twoheadrightarrow_U t'$ be a complete development of U in t strongly converging to t'. By Lemma 5.5.35, there is a surjective mapping $\vartheta\colon \mathcal{P}(t, U, \mathcal{R}) \to \mathcal{P}(t', U', \mathcal{R})$ with $\mathsf{tr}(\Pi) = \mathsf{tr}(\vartheta(\Pi))$ for all $\Pi \in \mathcal{P}(t, U, \mathcal{R})$, where $U' = U/\!/S$. Hence, it holds that $\mathcal{T}r(t, U, \mathcal{R}) = \mathcal{T}r(t', U', \mathcal{R})$ and, consequently, $\mathcal{F}(t, U, \mathcal{R}) = \mathcal{F}(t', U', \mathcal{R})$. Since S is a complete development of U in t, we have that $U' = \emptyset$ which implies, according to Lemma 5.5.31, that $\mathcal{F}(t', U', \mathcal{R}) = t'$. Therfore, $\mathcal{F}(t, U, \mathcal{R}) = t'$.

(ii) Let $t' = t^{(V)}$. By Proposition 5.5.17, both reduction sequences S and T can be uniquely lifted to reduction sequences S' and T' in $\mathcal{R}^l$, respectively, such that $V/\!/S$ and $V/\!/T$ are determined by the final term of S' and T', respectively. It is easy to see that also $\mathcal{R}^l$ is an orthogonal ITRS and that S' and T' are complete developments of U in t'. Hence, we can invoke clause (i) of this proposition to conclude that the final terms of S' and T' coincide and that, therefore, also $V/\!/S$ and $V/\!/T$ coincide. □

The above proposition corresponds to Proposition 5.4.17 which establishes the same properties for complete developments in the MRS model. Therefore, we can adopt the same notational abbreviation $U/\!/V$, mentioned in Notation 5.4.18, which denotes the descendants of a set U of non-$\bot$-occurrences in some term t by a complete development of a set of redex occurrences V in t. Furthermore, Proposition 5.5.36 yields the following corollary establishing the diamond property of complete developments as illustrated in Figure 5.3.

Corollary 5.5.37 (diamond property of complete developments) *Let $\mathcal{R}$ be an orthogonal ITRS and $t \twoheadrightarrow_U t_1$ and $t \twoheadrightarrow_V t_2$ be two complete developments of U respectively V in t. Then t_1 and t_2 are joinable by complete developments $t_1 \twoheadrightarrow_{V/\!/U} t'$ and $t_2 \twoheadrightarrow_{U/\!/V} t'$.*

Proof. By Proposition 5.5.11, it holds that

$$(U \cup V)/\!/U = U/\!/U \cup V/\!/U = V/\!/U.$$

Let $S\colon t \twoheadrightarrow_U t_1$, $T\colon t \twoheadrightarrow_V t_2$, $S'\colon t_1 \twoheadrightarrow_{V/\!/U} t'$ and $T'\colon t_2 \twoheadrightarrow_{U/\!/V} t''$. By the equation above and Proposition 5.5.8, we have that $S \cdot S'\colon t \twoheadrightarrow_U t_1 \twoheadrightarrow_{V/\!/U} t'$ is a complete development of $U \cup V$. Analogously, we obtain that $T \cdot T'\colon t \twoheadrightarrow_V t_2 \twoheadrightarrow_{U/\!/V} t''$ is a complete development of $U \cup V$, too. According to Proposition 5.5.36, this implies that both $S \cdot S'$ and $T \cdot T'$ end in the same term, i.e. $t' = t''$. □

With these properties of complete developments we can establish an Infinitary Strip Lemma for PRS reductions in the same way as for MRS reductions (cf. Proposition 5.4.23):

Proposition 5.5.38 (Infinitary Strip Lemma)
Let $\mathcal{R}$ be an orthogonal ITRS, $S\colon t_0 \twoheadrightarrow^{\alpha} t_\alpha$ a strongly convergent reduction, and $t_0 \twoheadrightarrow_U s_0$ a complete development of a set U of disjoint redex occurrences in t_0. Then t_α and s_0 are joinable by a reduction $s_0 \twoheadrightarrow s_\alpha$ and a complete development $t_\alpha \twoheadrightarrow_{U/\!/S} s_\alpha$.

Proof. We prove this statement by constructing the diagram shown in Figure 5.5. The 'U_ι's in the diagram are sets of redex occurrences: $U_\iota = U/\!/S|_{[0,\iota)}$ for all $0 \le \iota \le \alpha$. In particular, $U_0 = U$. All arrows in the diagram represent complete developments of the indicated sets of redex occurrences. Particularly, in each ι-th step of S the redex at v_ι is contracted. We will construct the diagram by an induction on α.

If $\alpha = 0$, then the diagram is trivial. If α is a successor ordinal $\beta + 1$, then we can take the diagram for the prefix $S|_{[0,\beta)}$, which exists by induction hypothesis, and extend it to a diagram for S. The existence of the additional square that completes the diagram for S is asserted by Corollary 5.5.37 since $U_{\beta+1} = U_\beta/\!/v_\beta$.

Let α be a limit ordinal. Moreover, let s'_α be the uniquely determined final term of a complete development of U_α in t_α. By induction hypothesis, the diagram exists for each proper prefix of S. Let $T_\iota\colon s_0 \twoheadrightarrow s_\iota$ denote the reduction sequence at the bottom of the diagram for the reduction sequence $S|_{[0,\iota)}$ for each $\iota < \alpha$. The set of all T_ι is directed. Hence, $T = \bigsqcup_{\iota<\alpha} T_\iota$ exists. Since $T_\iota < T$ for each $\iota < \alpha$, the diagram for S with $T\colon s_0 \twoheadrightarrow s_\alpha$ at the bottom satisfies almost all required properties. Only the equality of s_α and s'_α remains to be shown.

Note that, by Proposition 5.5.14, the redex occurrences in U_α are pairwise disjoint. Let $\pi \in U_\alpha$. By Lemma 5.5.10 and the definition of descendants, there is some $\beta < \alpha$ such that $\pi \in U_\iota$ and $v_\iota \nleq \pi$ for all $\beta \le \iota < \alpha$. Hence, for all $\pi' \in v_\iota/\!/U_\iota$ with $\beta \le \iota < \alpha$, we also have $\pi' \nleq \pi$. That is, in the remaining reductions $t_\beta \twoheadrightarrow t_\alpha$ and $t_\beta \twoheadrightarrow_{U_\beta} s_\beta \twoheadrightarrow s_\alpha$, no reduction takes place at a proper prefix of π. Hence, by Lemma 5.5.1, t_β coincides with t_α and s_α in all proper prefixes of π. Since in the reduction $t_\alpha \twoheadrightarrow_{U_\alpha} s'_\alpha$ also no reduction takes place at a proper prefix of π, we obtain that t_α and s'_α and, thus, also s_α and s'_α coincide in all proper prefixes of π.

Let $\rho\colon l \to r$ be the rule for the redex $t_\beta|_\pi$ and $C\langle,\dots,\rangle, D\langle,\dots,\rangle$ ground contexts such that $l = C\langle x_1,\dots,x_k\rangle$ and $r = D\langle x_{p(1)},\dots,x_{p(m)}\rangle$ for some pairwise distinct variables $x_1,\dots,x_k$ and an appropriate mapping $p\colon \{1,\dots,m\} \to \{1,\dots,k\}$. Moreover, let $t^\iota_1,\dots,t^\iota_k$ be terms such that $t_\iota = t_\iota[C\langle t^\iota_1,\dots,t^\iota_k\rangle]_\pi$ and $s_\iota = s_\iota[D\langle t^\iota_{p(1)},\dots,t^\iota_{p(m)}\rangle]_\pi$ for all $\beta \le \iota \le \alpha$. The argument in the previous paragraph justifies the assumption of these elements. From β onward, all horizontal reduction steps in the diagram take place within the contexts $t_\iota[\cdot]_\pi$ and $s_\iota[\cdot]_\pi$, respectively, or inside the terms t^ι_i, and all vertical reductions take place within the contexts $t_\iota[C\langle,\dots,\rangle]_\pi$ and $s_\iota[D\langle,\dots,\rangle]_\pi$, respectively. In particular, we have $t_\alpha = t_\alpha[C\langle t^\alpha_1,\dots,t^\alpha_k\rangle]_\pi$ and $s_\alpha = s_\alpha[D\langle t^\alpha_{p(1)},\dots,t^\alpha_{p(m)}\rangle]_\pi$. Let $t_\alpha \to_\pi t'_\alpha$. This reduction contracts the redex $C\langle t^\alpha_1,\dots,t^\alpha_k\rangle$ to the subterm $D\langle t^\alpha_{p(1)},\dots,t^\alpha_{p(m)}\rangle$ using rule ρ. Note that a complete development $t_\alpha \twoheadrightarrow_{U_\alpha} s'_\alpha$ contracts, besides π, only redex occurrences disjoint with π. Hence, t'_α and s'_α coincide in all extensions of π. Since $t'_\alpha = t_\alpha[D\langle t^\alpha_{p(1)},\dots,t^\alpha_{p(k)}\rangle]_\pi$ (and $s_\alpha = s_\alpha[D\langle t^\alpha_{p(1)},\dots,t^\alpha_{p(m)}\rangle]_\pi$), we can conclude that s_α and s'_α coincide in all extensions of π.

Since the residual $\pi \in U_\alpha$ was chosen arbitrarily, the above holds for all elements in U_α. That is, s_α and s'_α coincide in all prefixes and all extensions of elements in U_α. It remains to be shown, that they also coincide in positions that are disjoint to all positions in U_α. To this end, we only need to show that t_α and s_α coincide in these positions since the complete development $t_\alpha \twoheadrightarrow_{U_\alpha} s'_\alpha$ keeps positions disjoint with all positions in U_α unchanged. Let π be such a position.

Suppose $t_\alpha(\pi) = f \neq \bot$. By Lemma 5.5.1, there is some $\beta < \alpha$ such that $t_\beta(\pi) = f$ and $v_\iota \nleq \pi$ for all $\beta \le \iota < \alpha$. Note that no prefix π' of π is in U_β since otherwise $\pi' \in U_\alpha$, by Lemma 5.5.10, which contradicts the assumption that π is disjoint to all positions in U_α. Hence, $s_\beta(\pi) = f$ and $\pi' \nleq \pi$ for all $\pi' \in v_\iota/\!/U_\iota$ and $\beta \le \iota < \alpha$, which means that no reduction step in $s_\beta \twoheadrightarrow s_\alpha$ takes place at some prefix of π. Thus, we can conclude, according to Lemma 5.5.1, that $s_\alpha(\pi) = f$. Similarly, one can show that $s_\alpha(\pi) = f \neq \bot$ implies $t_\alpha(\pi) = f$.

Suppose $t_\alpha(\pi) = \bot$. Hence, according to Lemma 5.5.4, π is outermost-volatile in S or there is some $\beta < \alpha$ such that $t_\beta(\pi) = \bot$ and $v_\iota \nleq \pi$ for all $\beta \leq \iota < \alpha$. For the latter case, we can argue as in the case for $t_\alpha(\pi) \neq \bot$ above. In the former case, π is outermost-volatile in T as well. Thus, by applying Lemma 5.5.4, we obtain that $s_\alpha(\pi) = \bot$. A similar argument can be employed for the reverse direction. □

Note that the concepts of tiling diagrams and projections that we have seen for MRS reductions in Definition 5.4.21 can be carried over to our present setting of PRS reductions verbatimly. Then the above proposition states that tiling diagrams always exists for a pair of reductions consisting of an arbitrary strongly convergent reduction and a single reduction step. This can be used in analogy to Corollary 5.4.24 in order to obtain a semi-infinitary confluence property for orthogonal systems as illustrated in Figure 5.6.

Corollary 5.5.39 (semi-infinitary confluence)
Let $\mathcal{R}$ be an orthogonal ITRS, $H\colon t \twoheadrightarrow t_2$, and $V\colon t \to^* t_1$. Then the projections $V/H\colon t_2 \twoheadrightarrow t_3$ and $H/V\colon t_1 \twoheadrightarrow t_3$ exist.

Proof. This can be shown by an induction on the length of V. If V is empty, the statement trivially holds. The induction step follows from Proposition 5.5.38. □

5.5.3 Relation to Böhm Trees

This section is going to provide an interesting insight into the concepts of Böhm reductions and Böhm trees which were covered in Section 5.4.5. Recall that Böhm reductions are based on the intuition of meaningless terms. The Böhm reduction of a ITRS is obtained by adjoining additional rules to the system which allow to rewrite terms that are considered meaningless directly to the fresh constant symbol $\bot$. The least set of meaningless terms is $\mathcal{RA}$, the set of root-active terms. Below, the notion of *fragile terms* is introduced which has some similarity to root-active terms.

Definition 5.5.40 (destructive reductions, fragile terms)
Let $\mathcal{R}$ be an ITRS.

(i) A strongly convergent reduction sequence $S\colon t \twoheadrightarrow s$ is called *destructive* if ε is a volatile position in S.

(ii) A partial term t in $\mathcal{R}$ is called *fragile* if a destructive reduction sequence starts in t.

It is quite obvious that destructive reductions are open reductions which strongly converge to $\bot$ – hence the name. Consequently, fragile terms are terms which are reducible to $\bot$ by an open reduction.

Fact 5.5.41 (destructive reductions, fragile terms)
Let $\mathcal{R}$ be an ITRS.

(i) A strongly convergent reduction sequence $S\colon t \twoheadrightarrow^\lambda s$ is destructive iff λ is a limit ordinal and $s = \bot$

(ii) A partial term t in $\mathcal{R}$ is fragile iff there is an open strongly convergent reduction $t \twoheadrightarrow \bot$.

Proof. This follows immediately from Lemma 5.5.4. □

One has to keep in mind, however, that a closed reduction to $\bot$ is not destructive. Such a notion of destructiveness would include the empty reduction sequence from $\bot$ to $\bot$, and reductions that end with the contraction of a collapsing redex as, for example, in the single step reduction $f(\bot) \to \bot$ induced by the rule $f(x) \to x$. Such reductions do not "produce"

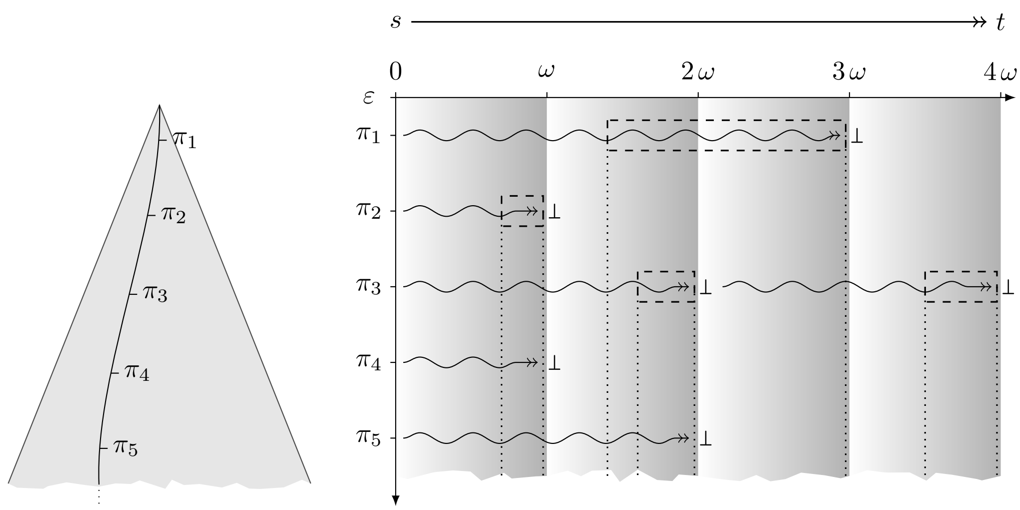

(a) Nested volatile positions. (b) Replacing nested destructive reductions with $\to_\perp$ steps.

Figure 5.12: Turning a PRS reduction into a Böhm reduction.

the $\perp$ term. They are merely capable of copying a subterm $\perp$ by a collapsing rule. In this sense, fragile terms are, according to Lemma 5.5.5, the only terms which can produce the $\perp$ term. This is the key observation for studying the relation between PRS reductions and Böhm reductions.

Fragile terms are, as already mentioned above, the only terms that can produce the $\perp$ term. The Böhm reduction w.r.t. the set of all fragile terms can contract fragile terms to $\perp$ immediately. Since the only difference between PRS reductions and MRS reduction is, by Corollary 5.2.3, the ability of PRS reductions to spawn $\perp$ symbols, the question arises whether the concept of Böhm reductions can bridge this gap. The objective of this section is to confirm this conjecture. More specifically, we will establish the equivalence of PRS reductions and the MRS reductions in the Böhm reduction w.r.t. the total fragile terms. This equivalence will then allow us to hijack some of the results that are already known for Böhm reductions (cf. Section 5.4.5).

One direction of the equivalence of PRS reductions and Böhm reductions is comparatively easy: Given a PRS reduction sequence one can construct a Böhm reduction sequence by removing reduction steps which cause the volatility of a position in some open prefix of the definition and replacing them by a $\to_\perp$-step. The intuition of this construction is illustrated in Figure 5.12. It shows a PRS reduction of length $4\,\omega$ from s to t. In order to maintain readability, we restrict the attention to a particular branch of the term (tree) as indicated in Figure 5.12a. The picture shows five positions which are volatile in some open prefix of the reduction. We assume that they are the only volatile positions at least in the considered branch. Note that the positions do not need to occur in all of the terms in the reduction sequence. They might disappear and reappear repeatedly. Each of them, however, appears in infinitely many terms in the reduction sequence, as, by definition of volatility, infinitely many steps take place at each of these positions. In Figure 5.12b, the prefixes of the reduction that contain a volatile position are indicated by a waved rewrite arrow pointing to a $\perp$. The level of an arrow indicates the position which is volatile. A prefix might have multiple volatile positions. For example, both π_2 and π_4 are volatile in the prefix of length ω. But a position might also be volatile for several prefixes. For instance, π_3 is volatile in the prefix of length $2\,\omega$ and the prefix of length $4\,\omega$.

By Lemma 5.5.4, outermost-volatile positions are responsible for the emergence of $\perp$ symbols. By their nature, at some point there are no reductions taking place above outermost-volatile positions. The suffix where this is the case is a nested destructive reduction sequence. The subterm where this suffix starts is, therefore, a fragile term and we can replace this suffix with a single $\to_\perp$-step. The segments which are replaced in this way are highlighted by dashed boxes in Figure 5.12b. As indicated by the dotted lines, this then also includes reduction steps which occur below the outermost-volatile positions. Therefore, also volatile positions which are not outermost are removed as well. Eventually, we obtain a reduction sequence without volatile positions, which is, by Lemma 5.5.5, a Böhm reduction.

The following proposition summarises the above observation and provides a more rigorous proof.

Proposition 5.5.42 (PRS reductions are Böhm reductions)
Let $\mathcal{R}$ be an ITRS, $\mathcal{U}$ the set of fragile terms in $\mathcal{T}^\infty(\Sigma, \mathcal{V})$, and $\mathcal{B}$ the Böhm reduction of $\mathcal{R}$ w.r.t. $\mathcal{U}$. Then, for each PRS reduction $s \twoheadrightarrow^p_{\mathcal{R}} t$ in $\mathcal{R}$, there is a MRS reduction $s \twoheadrightarrow^m_{\mathcal{B}} t$ in $\mathcal{B}$.

Proof. W.l.o.g. we can assume that s, the starting term of the reduction, is total. If s is not total, consider the signature $\Sigma' = \Sigma \uplus \{\perp\}$ and take $\perp'$ as the "new bottom" symbol, i.e. $\mathcal{T}^\infty(\Sigma'_{\perp'}, \mathcal{V})$ is the set of partial terms.

Assume that there is a strongly convergent PRS reduction sequence $S\colon (t_\iota \to_{\pi_\iota} t_{\iota+1})_{\iota<\alpha}$ in $\mathcal{R}$ that converges to t_α. We will construct a strongly convergent MRS reduction sequence $T\colon t_0 \twoheadrightarrow t_\alpha$ in $\mathcal{B}$ by removing reduction steps in S that take place at or below outermost-volatile positions of some prefix of S and replace them by $\to_\perp$-steps.

Let π be an outermost-volatile position of some prefix $S|_{[0,\lambda)}$. Then there is some ordinal $\beta < \lambda$ such that no reduction step between β and λ in S takes place strictly above π, i.e. $\pi_\iota \not< \pi$ for all $\beta \le \iota < \lambda$. Such an ordinal β must exist since otherwise π would not be an outermost-volatile position in $S|_{[0,\lambda)}$. Hence, we can construct a destructive reduction sequence $S'\colon t_\beta|_\pi \twoheadrightarrow^p \perp$ by taking the subsequence of $S|_{[\beta,\lambda)}$ containing the reduction steps at π or below. Note that $t_\beta|_\pi$ might still contain the symbol $\perp$. Since $\perp$ is not relevant for the applicability of rules in $\mathcal{R}$, each of the $\perp$ symbols in $t_\beta|_\pi$ can be safely replaced by arbitrary total terms, in particular by terms in $\mathcal{U}$. Let r be a term that is obtained in this way. Then there is a destructive reduction sequence $S''\colon r \twoheadrightarrow^p \perp$ that applies the same rules at the same positions as S'. Hence, $r \in \mathcal{U}$. By construction, r is a $\perp,\mathcal{U}$-instance of $t_\beta|_\pi$ which means that $t_\beta|_\pi \in \mathcal{U}_\perp$. Additionally, $t_\beta|_\pi \neq \perp$ since there is a non-empty reduction $S'\colon t_\beta|_\pi \twoheadrightarrow^p \perp$ starting in $t_\beta|_\pi$. Consequently, there is a rule $t_\beta|_\pi \to \perp$ in $\mathcal{B}$. Let T' be the reduction sequence that is obtained from $S|_{[0,\lambda)}$ by replacing the β-th step, which we can assume w.l.o.g. to take place at π, by a step with the rule $t_\beta|_\pi \to \perp$ at the same position π and removing all reduction steps φ_ι taking place at π or below for all $\beta < \iota < \lambda$. Let t' be the term that the reduction sequence T' converges to. t_λ and t' can only differ in positions π or below. However, by construction, we have $t'(\pi) = \perp$ and, by Lemma 5.5.4, $t_\lambda(\pi) = \perp$. Consequently, $t' = t_\lambda$.

This construction can be done for all outermost-volatile positions and all prefixes of S. Thereby, we obtain a PRS reduction sequence $T\colon t_0 \twoheadrightarrow^p_{\mathcal{B}} t_\alpha$ for which no prefix has a volatile position. By Lemma 5.5.5, T is a total reduction sequence. Note that $\mathcal{B}$ is a ITRS over the extended signature $\Sigma' = \Sigma \cup \{\perp\}$, i.e. terms containing $\perp$ are considered total (cf. the initial remark of this proof). Hence, by Corollary 5.2.3, $T\colon t_0 \twoheadrightarrow^m_{\mathcal{B}} t_\alpha$. □

With this half of the equivalence we can already establish a result similar to Theorem 5.4.4 which allows the approximation of an arbitrarily large finite part of a result term of a PRS reduction using a finite reduction:

Proposition 5.5.43 (finite approximation of non-$\perp$ occurrences)
Let $\mathcal{R}$ be a left-linear ITRS and $s \twoheadrightarrow t$. Then, for each finite set $P \subseteq \mathcal{P}_{\setminus\perp}(t)$, there is a finite reduction $s \to^* t'$ such that t and t' coincide in all positions in P.

Proof. Assume that $s \twoheadrightarrow^p_{\mathcal{R}} t$. Then, by Proposition 5.5.42, there is a reduction $s \twoheadrightarrow^m_{\mathcal{B}} t$, where $\mathcal{B}$ is the Böhm reduction of $\mathcal{R}$ w.r.t. the set of total fragile terms of $\mathcal{R}$. By Lemma 5.4.38, there is a reduction $s \twoheadrightarrow^m_\mathcal{R} s' \twoheadrightarrow^m_\bot t$. Then we have that s' and t coincide in all positions in $\mathcal{P}_{\setminus\bot}(t)$. Let $d = \max\{|\pi|\,|\,\pi \in P\}$. d is well defined as P is finite. By Theorem 5.4.4, there is a reduction $s \to^\star_\mathcal{R} t'$ such that t' and s' coincide up to depth d and, thus, in particular they coincide in all positions in P. Consequently, t and t' coincide in all positions in P, too. $\square$

The next step for our goal is to show that destructive reduction sequences can be compressed to a length of exactly ω. To achieve this, we need that the projection of a destructive reduction is again destructive:

Lemma 5.5.44 (preservation of destructive reductions by finite projections)
Let $\mathcal{R}$ be an orthogonal ITRS, $S\colon t_0 \twoheadrightarrow t_\alpha$ a destructive reduction, and $T\colon t_0 \to s_0$ a single reduction step. Then the projection $S/T\colon s_0 \twoheadrightarrow s_\alpha$ is also destructive.

Proof. We consider the situation depicted in Figure 5.5 on page 104. Since $S\colon t_0 \twoheadrightarrow t_\alpha$ is destructive, we have, for each $\beta < \alpha$, some $\beta \le \gamma < \alpha$ such that $v_\gamma = \varepsilon$. If $v_\gamma = \varepsilon$, then also $\varepsilon \in v_\gamma /\!/ U_\gamma$ unless $\varepsilon \in U_\gamma$. As by Proposition 5.5.14, U_γ is a set of pairwise disjoint positions, $\varepsilon \in U_\gamma$ implies $U_\gamma = \{\varepsilon\}$. This means that if $v_\gamma = \varepsilon$ and $\varepsilon \in U_\gamma$, then $U_\iota = \emptyset$ for all $\gamma < \iota < \alpha$. Thus, this can only happen at most once. Therefore, we have, for each $\beta < \alpha$, some $\beta \le \gamma < \alpha$ such that $\varepsilon \in v_\gamma /\!/ U_\gamma$. Hence, T is destructive. $\square$

As a consequence of this preservation of destructiveness by finite projections, we obtain that the set of fragile terms is closed under finite reductions:

Lemma 5.5.45 (closure of fragile terms under finite reductions)
In each orthogonal ITRS, the set of fragile terms is closed under finite reductions.

Proof. Let t be a fragile term and $T\colon t \to^\star t'$ a finite reduction sequence. We prove by an induction on the length of T that then also t' is fragile. If T is the empty reduction sequence, this is trivial. Let T be of the form $t \to^\star s \to t'$. By induction hypothesis, s is fragile. Hence, there is a destructive reduction S starting in s. By applying Proposition 5.5.38 and Lemma 5.5.44 to S and $s \to t'$, we obtain a destructive reduction sequence starting in t'. Hence, also t' is fragile. $\square$

Now we can show that destructiveness does not need more that ω steps in orthogonal systems. This property will become important when proving the compression property for PRS reductions.

Proposition 5.5.46 (compression of destructive reductions)
Let $\mathcal{R}$ be an orthogonal ITRS and t a partial term in $\mathcal{R}$. If there is a destructive reduction sequence starting in t, then there is a destructive reduction sequence of length ω starting in t.

Proof. Let $S\colon t_0 \twoheadrightarrow^\lambda \bot$ be a destructive reduction sequence starting in t_0. Hence, there is some $\alpha < \lambda$ such that $S|_{[0,\alpha)}\colon t_0 \twoheadrightarrow s_1$, where s_1 is a ρ-redex for some $\rho \in R$. Let P be the set of pattern positions of the ρ-redex s_1, i.e. $P = \mathcal{P}_\Sigma(l)$ for l the left-hand side of ρ. By Proposition 5.5.43, there is a finite reduction $t_0 \to^\star s_1'$ such that s_1 and s_1' coincide in all positions in P. Hence, because $\mathcal{R}$ is left-linear, also s_1' is a ρ-redex. Now consider the reduction sequence $T_0\colon t_0 \to^\star s_1' \to_{\rho,\varepsilon} t_1$. T_0 is of finite length and, by applying Lemma 5.5.45, we get that t_1 is fragile.

The above argument can be repeated arbitrarily often which yields for each $i < \omega$ a finite reduction sequence $T_i\colon t_i \to^\star t_{i+1}$ whose last step is a contraction at the root. Then the concatenation $T = \prod_{i<\omega} T_i$ of these reduction sequences is a destructive reduction sequence of length ω starting in t_0. $\square$

The above proposition bridges the gap between fragility and root-activeness. Whereas the former concept is defined in terms of transfinite reductions, the latter is defined in terms of finite reductions. By Proposition 5.5.46, however, a fragile term is always finitely reducible to a redex. This is the key to the observation that fragility is not only quite similar to root-activeness but is, in fact, essentially the same concept.

Proposition 5.5.47 (root-activeness = fragility)
Let $\mathcal{R}$ be an orthogonal ITRS and t a total term in $\mathcal{R}$. Then t is root-active iff t is fragile.

Proof. The "only if" direction is easy: If t is root active, then there is a reduction sequence S of length ω starting in t with infinitely many steps taking place at the root. Hence, $S\colon t \twoheadrightarrow^{\omega}_{\mathcal{R}} \bot$ is a destructive reduction sequence and t a fragile term.

Consider the converse direction: To this end, we assume that t is fragile and show that, for each reduction $t \to^{*} s$, there is a reduction $s \to^{*} t'$ to a redex t'. By Lemma 5.5.45, also s is fragile. Hence, there is a destructive reduction sequence $S\colon s \twoheadrightarrow \bot$ starting in s. According to Proposition 5.5.46, we can assume that S has length ω. Therefore, there is some $n < \omega$ such that $S|_{[0,n)}\colon s \to^{*} t'$ for a redex t'. □

Finally, we have gathered all tools necessary in order to prove the converse direction of the equivalence of PRS reductions and Böhm reductions. Since root-activeness is a well-established notion, we prefer to consider the Böhm reduction w.r.t. the set of root-active terms instead of the equivalent set of total fragile terms.

Theorem 5.5.48 (PRS reductions = Böhm reductions)
Let $\mathcal{R}$ be an ITRS and $\mathcal{B}$ the Böhm reduction of $\mathcal{R}$ w.r.t. $\mathcal{RA}$. Then there is a PRS reduction $s \twoheadrightarrow^{p}_{\mathcal{R}} t$ in $\mathcal{R}$ iff there is an MRS reduction $s \twoheadrightarrow^{m}_{\mathcal{B}} t$ in $\mathcal{B}$.

Proof. The "only if" direction follows immediately from Proposition 5.5.47 and Proposition 5.5.42.

Now consider the converse direction: Let $s \twoheadrightarrow^{m}_{\mathcal{B}} t$ be an MRS reduction in $\mathcal{B}$. Due to Lemma 5.4.38, there is a term $s' \in \mathcal{T}^{\infty}(\Sigma, \mathcal{V})$ such that there are MRS reductions $s \twoheadrightarrow^{m}_{\mathcal{R}} s'$ and $s' \twoheadrightarrow^{m}_{\bot} t$. As $\mathcal{RA}$ is a set of meaningless terms, we can assume, by Lemma 5.4.40, that all steps in $s' \twoheadrightarrow^{m}_{\bot} t$ occur at pairwise disjoint occurrences of root-active terms. By Proposition 5.5.47, root-active terms $r \in \mathcal{RA}$ are fragile, i.e. we have a destructive reduction $r \twoheadrightarrow^{p}_{\mathcal{R}} \bot$ starting in r. Thus, we can construct a PRS reduction $s' \twoheadrightarrow^{p}_{\mathcal{R}} t$ by replacing each step $C[r] \to_{\bot} C[\bot]$ in $s' \twoheadrightarrow^{m}_{\bot} t$ with the corresponding reduction $C[r] \twoheadrightarrow^{p}_{\mathcal{R}} C[\bot]$. By combining this reduction with the MRS reduction $s \twoheadrightarrow_{\mathcal{R}} s'$, which, according to Corollary 5.2.3, is also a PRS reduction, we obtain a PRS reduction $s \twoheadrightarrow^{p}_{\mathcal{R}} t$. □

With this equivalence, PRS reductions inherit a number of important properties that are enjoyed by Böhm reductions. Most prominently these properties include infinitary confluence and infinitary normalisation:

Theorem 5.5.49 (infinitary confluence)
Every orthogonal ITRS is infinitarily confluent.

Proof. Let $\mathcal{R}$ be an orthogonal ITRS. According to Theorem 5.5.48, it holds that $\twoheadrightarrow^{p}_{\mathcal{R}} = \twoheadrightarrow^{m}_{\mathcal{B}}$, where $\mathcal{B}$ is the Böhm reduction of $\mathcal{R}$ w.r.t. $\mathcal{RA}$. Since, by Proposition 5.4.36, $\mathcal{RA}$ is a set of meaningless terms, we can employ Theorem 5.4.42 to obtain that the Böhm reduction $\mathcal{B}$ is infinitary confluent w.r.t. MRS reductions. That is, $\twoheadrightarrow^{m}_{\mathcal{B}}$ satisfies the diamond property. Hence, so does $\twoheadrightarrow^{p}_{\mathcal{R}}$ which means that $\mathcal{R}$ is infinitarily confluent w.r.t. PRS reductions □

Theorem 5.5.50 (infinitary normalisation)
Every orthogonal ITRS is infinitarily normalising.

Proof. Similar to the proof of Theorem 5.5.49, however, referring to Theorem 5.4.41 instead of Theorem 5.4.42. □

Since PRS reductions in orthogonal systems essentially consist of a prefix which is an MRS reduction and a suffix consisting of nested destructive reductions, we can employ the Compression Lemma for MRS reductions (cf. Theorem 5.4.3) and the Compression Lemma for destructive reductions (cf. Proposition 5.5.46) to obtain the Compression Lemma for PRS reductions:

Theorem 5.5.51 (Compression Lemma)
Let $\mathcal{R}$ be an orthogonal ITRS. If $s \twoheadrightarrow t$, then $s \twoheadrightarrow^{\leq\omega} t$.

Proof. Let $s \twoheadrightarrow^{p}_{\mathcal{R}} t$. According to Theorem 5.5.48, we have $s \twoheadrightarrow^{m}_{\mathcal{B}} t$ for the Böhm reduction $\mathcal{B}$ of $\mathcal{R}$ w.r.t. $\mathcal{RA}$ and, therefore, by Lemma 5.4.38, we have reduction sequences $S\colon s \twoheadrightarrow^{m}_{\mathcal{R}} s'$ and $T\colon s' \twoheadrightarrow^{m}_{\perp} t$. Due to Theorem 5.4.3, we can assume that S is of length at most ω. If T is the empty reduction sequence, then we are done. If not, then T is a complete development of pairwise disjoint occurrences of root-active terms according to Lemma 5.4.40. Hence, each step is of the form $C[r] \to_{\perp} C[\perp]$ for some root-active term r. By Proposition 5.5.47, for each such term r, there is a destructive reduction $r \twoheadrightarrow^{p}_{\mathcal{R}} \perp$ which we can assume, in accordance with Proposition 5.5.46, to be of length ω. Hence, each step $C[r] \to_{\perp} C[\perp]$ can be replaced by the reduction $C[r] \twoheadrightarrow^{p}_{\mathcal{R}} C[\perp]$. Concatenating these reductions results in a reduction sequence $T'\colon s' \twoheadrightarrow^{p}_{\mathcal{R}} t$ of length at most $\omega \cdot \omega$. If $S\colon s \twoheadrightarrow^{m}_{\mathcal{R}} s'$ is of finite length, we can interleave the reduction steps in T' such that we obtain a reduction $T''\colon s' \twoheadrightarrow^{p,\omega}_{\mathcal{R}} t$ of length ω. Then we have $S \cdot T''\colon s \twoheadrightarrow^{p,\omega}_{\mathcal{R}} t$. If $S\colon s \twoheadrightarrow^{m}_{\mathcal{R}} s'$ has length ω, we construct a reduction $s \twoheadrightarrow^{p}_{\mathcal{R}} t$ as follows: As illustrated above, T' consists of destructive reduction sequences taking place at some pairwise disjoint positions. These steps can be interleaved into the reduction sequence S resulting into a reduction sequence $s \twoheadrightarrow^{p}_{\mathcal{R}} t$ of length ω. The argument for that is similar to that employed in the successor case of the induction proof of the Compression Lemma in [KKSdV95a]. $\square$

Chapter 6

Term Graph Rewriting

The purpose of this chapter is to introduce rewriting on term graphs. The more general topic of graph rewriting, from which term graph rewriting stems, was first studied in the late 1960's [PR69, Pra71]. Soon, graphs were used to represent terms in order to be able to use graph rewriting to efficiently implement term rewriting [Wad71]. The ability of term graph rewriting to simulate term rewriting shall be the focus of this chapter. This employment of term graph rewriting is particularly attractive for infinitary rewriting as it in some cases allows to simulate an infinite reduction sequence on terms by a finite reduction sequence on term graphs. Section 6.2 discusses both these aspects.

Since the first appearance of graph rewriting many different approaches have been suggested and studied. This variety of formalisations breaks down into three major approaches: The *operational approaches* [PR69, BvEG⁺87, Ech08] define the rewriting of graphs by explicit removal and insertion of subgraphs as specified by the rewrite rules. On the other hand, the *algebraic methods* [EPS73, Rao84, CMR⁺97] use pushout constructions in appropriately defined categories in order to formalise the semantics of a rewrite rule. The third way of defining rewriting in graphs, called *equational term graph rewriting* [AK96], interprets a graph as a set of equations of terms, where the nodes of the graph are essentially represented as variables in these equations. The rewrite rules are then applied to the equations that make up a term graph.

In this thesis we prefer the operational view of term graph rewriting. In particular, we favour the approach of Barendregt et al. [BvEG⁺87]. We define the basic notions of finitary and infinitary term graph rewriting in Section 6.1.

6.1 Term Graph Rewriting Systems

As mentioned in the introduction, we consider term graph rewriting along the lines of Barendregt et al. [BvEG⁺87]. The most important reason for this decision is that there is already a very good understanding of the relation of this style of term graph rewriting and infinitary term rewriting – at least for the metric model[1]. Secondly, it is favourable for being relatively close to an implementation.

Just as for terms, we assume to have a countably infinite set $\mathcal{V}$ of variables which is always chosen such that it is disjoint from the signature Σ. Variables are used to parametrise the rules of a term graph rewriting system. Therefore, the extended signature $\Sigma_\mathcal{V} = \Sigma \uplus \mathcal{V}$ is considered, where the variables are interpreted as nullary symbols. In analogy to terms we use the notation $\mathcal{G}_\mathcal{C}(\Sigma, \mathcal{V})$ and $\mathcal{G}^\infty_\mathcal{C}(\Sigma, \mathcal{V})$ instead of $\mathcal{G}_\mathcal{C}(\Sigma_\mathcal{V})$ and $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\mathcal{V})$, respectively.

[1] One can argue (cf. [CD97]) that the double-pushout approach [EPS73] and the equational approach are better suited for a partial order model of infinitary term rewriting. We do not, however, make an attempt to analyse the relation between our partial order approach to infinitary term rewriting and these variants of term graph rewriting.

A term graph rewrite rule is represented by a single graph. The left- and right-hand side of a rule are distinguished by two (not necessarily distinct) root nodes. This allows that structures can also be shared between the left- and the right-hand side of a rule.

Definition 6.1.1 (term graph rewriting system, [BvEG⁺87])
Let Σ be a signature.

(i) A *term graph rewrite rule* ρ over Σ is a tuple $(N, \mathsf{lab}, \mathsf{suc}, l, r)$, where $g = (N, \mathsf{lab}, \mathsf{suc})$ is a finite $\Sigma_{\mathcal{V}}$-graph and $l, r \in N$. l and r are called the *left* and the *right root* of the rule, respectively. All nodes in g must be reachable from l or r. ρ_l and ρ_r denote the term graphs $(g|l, l)$ resp. $(g|r, r)$, called the *left-* resp. the *right-hand side* of ρ.

Additionally, we require that, for each variable $v \in \mathcal{V}$, there is at most one node in g labelled with v, and each node labelled by a variable must be reachable from the left root l.

If, instead of requiring that the underlying graph g is finite, we only require the left-hand side ρ_l to be finite, then ρ is called an *infinitary term graph rewrite rule*.

(ii) A *term graph rewriting system* (*GRS*) over Σ is a pair $\mathcal{R} = (\Sigma, R)$ consisting of a signature Σ and a set of term graph rewrite rules R over Σ. If, instead, R is a set of infinitary term graph rewrite rules, then $\mathcal{R}$ is called an *infinitary term graph rewriting system* (*IGRS*).

The restriction concerning $\mathcal{V}$-nodes is necessary for the definition of the semantics of term graph rewrite rules. In principle, there is no need to have a set of variable symbols. Since no two distinct nodes are allowed to be labelled by the same variable, one could have used also only a single symbol, say v, that distinguishes variable nodes from all other nodes.[2] The use of explicit variable symbols is only for convenience and to define the translation of IGRSs into ITRSs in a straightforward manner.

When presenting examples, we prefer using a graphical representation of term graph rewrite rules. To this end, we employ a graphical representation of the underlying graph and indicate the left and right node by a rewrite arrow going from the left to the right root of the rule.

Example 6.1.2
Consider the term rewrite rule $a(x, s(y)) \to s(a(x, y))$. This rule can be translated into a term graph rewrite rule as follows:

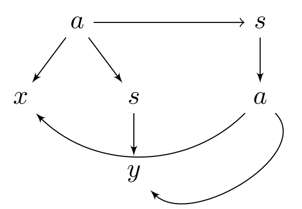

By translating a term rewrite rule ρ into a term graph rewrite rule, we mean finding a term graph rewrite rule ρ' such that $\rho\colon \mathcal{U}(\rho'_l) \to \mathcal{U}(\rho'_r)$. There might be, of course, multiple such translations. For the example above, however, there is only one.

Note that the rewrite arrow in the graphical representation of a term graph rewrite rule always points from the left root to the right root. Keep in mind that this is only a means to indicate the left and the right root, and should not be considered as part of the underlying graph. Nevertheless, it is also possible that the right root of the rule is reachable from the left root in the underlying graph. This is, for example, necessary when representing collapsing term rewrite rules such as $a(x, 0) \to x$:

[2] This was done in the original definition by Barendregt et al. [BvEG⁺87]. In their definition variable nodes are nodes without any label.

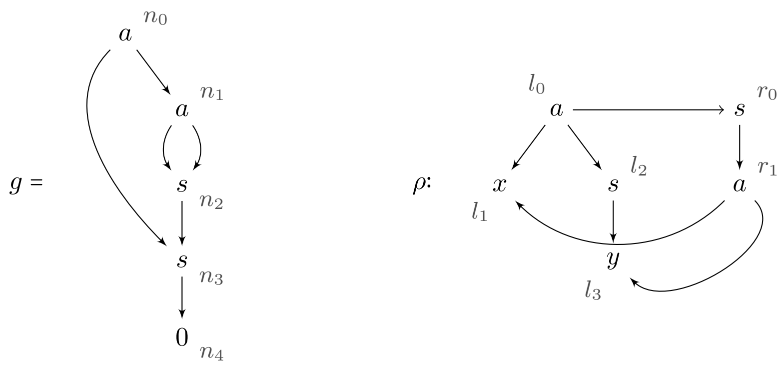

Figure 6.1: Example term graph and term graph rewrite rule.

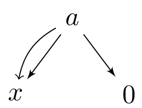

Since a variable is restricted to have only one occurrence in a rule, the only way of referencing it in the right-hand side is by sharing. The example above illustrates this. As we will see shortly, there is no explicit mechanism for duplication in the semantics of the term graph rewriting framework we are considering – hence, this restriction. This can be seen more explicitly when translating a duplicating term rewrite rule such as $d(x) \to a(x, x)$ to a term graph rewrite rule:

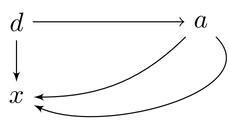

Next, we define how rewriting is performed on term graphs according to term graph rewrite rules introduced above. The construction that is involved in a rewriting step is quite technical. To this end, we want to convey the intuition of it beforehand. In principle, it follows the idea of term rewriting. The application of a term rewrite rule $l \to r$ to a term t consists of the following steps: At first the left-hand side l is matched with a subterm $t|_\pi$ by finding a substitution σ with $t|_\pi = l\sigma$. Then the subterm $t|_\pi$ is removed from t and replaced by the according instance $r\sigma$ of the right-hand side r.

The construction for term graphs is somewhat similar. For the sake of illustration, let us consider the term graph g and the term graph rewrite rule ρ shown in Figure 6.1. Equally to term rewriting, the first step is the *matching* of the left-hand side ρ_l of ρ with a sub-term graph of g. The equivalent to substitutions for term graphs are $\mathcal{V}$-homomorphisms. $\mathcal{V}$-homomorphisms instantiate variables as they allow nodes labelled with a variable to be mapped to an arbitrarily labelled node. In the example, we match the left-hand side with the sub-term graph rooted at n_1:

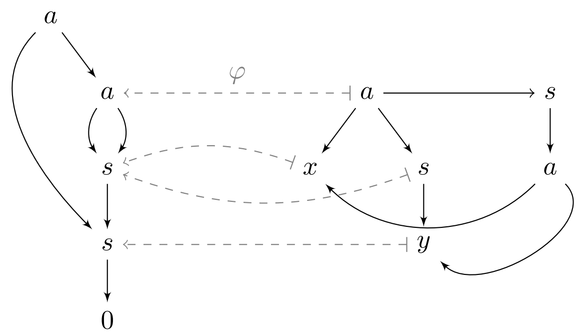

The sub-term graph $g|n_1$ is a *redex* of the rule ρ.

During the next step, the *build step*, all nodes in the term graph rewrite rule not reachable from the left root are added to the term graph g including all edges between them. Moreover, edges that start in these nodes having an endpoint in the rule's left-hand side, such as the two edges starting in r_1, are also copied, where the endpoint in g is then the corresponding image by φ.

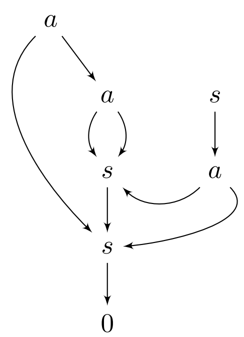

Next, in the *redirection step*, all edges ending in the root of the redex, viz. the node n_1, are redirected to the right root of ρ (resp. its image by φ if existent), viz. r_0. That is, the right-hand side of the rewrite rule is embedded into the graph exactly at that position where the redex has resided before:

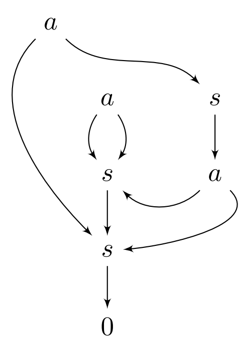

Eventually, in the *garbage collection step*, the root node of the new term graph is set, viz. the node n_0, and all nodes not reachable from the root node are removed.

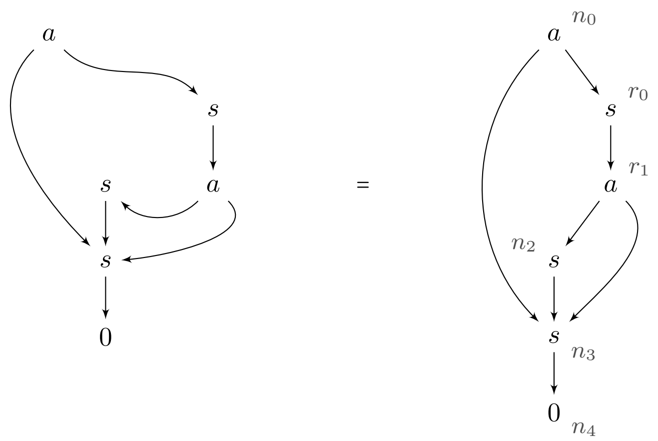

The graph on the right-hand side depicts the same term graph. It is only restructured for readability and contains the node names in order to compare it to the original term graph g and the rule ρ.

Below we will give the formal definition of the construction of the result of a term graph rewrite rule application.

Definition 6.1.3 (application of a term graph rewrite rule, [BvEG⁺87])
Let $\rho = (N^\rho, \mathsf{lab}^\rho, \mathsf{suc}^\rho, l^\rho, r^\rho)$ be a term graph rewrite rule over signature Σ, $g \in \mathcal{G}^\infty(\Sigma, \mathcal{V})$ and $n \in N^g$. ρ is called *applicable* to g at n if there is a $\mathcal{V}$-homomorphism $\varphi\colon \rho_l \to_{\mathcal{V}} g|n$. φ is called the *matching $\mathcal{V}$-homomorphism* of the rule application, and $g|n$ is called a *ρ-redex*. Next, we define the *result* of the application of the rule ρ to g at n using the $\mathcal{V}$-homomorphism φ. This is done by constructing the intermediate graphs g_1 and g_2, and the final result g_3.

(i) The graph g_1 is obtained from g by adding the part of ρ not contained in the left-hand side:

$$N^{g_1} = N^g \uplus (N^\rho \setminus N^{\rho_l})$$

$$\mathsf{lab}^{g_1}(m) = \begin{cases} \mathsf{lab}^g(m) & \text{if } m \in N^g \\ \mathsf{lab}^\rho(m) & \text{if } m \in N^\rho \setminus N^{\rho_l} \end{cases}$$

$$\mathsf{suc}_i^{g_1}(m) = \begin{cases} \mathsf{suc}_i^g(m) & \text{if } m \in N^g \\ \mathsf{suc}_i^\rho(m) & \text{if } m, \mathsf{suc}_i^\rho(m) \in N^\rho \setminus N^{\rho_l} \\ \varphi(\mathsf{suc}_i^\rho(m)) & \text{if } m \in N^\rho \setminus N^{\rho_l}, \mathsf{suc}_i^\rho(m) \in N^{\rho_l} \end{cases}$$

(ii) Let l and r be the nodes in g_1 corresponding to l^ρ and r^ρ, respectively. That is, $l = n$, and $r = \varphi(r^\rho)$ if $r^\rho \in N^{\rho_l}$ and $r = r^\rho$ otherwise. g_2 is obtained from g_1 by redirecting edges ending in l to r:

$$N^{g_2} = N^{g_1}$$

$$\mathsf{lab}^{g_2} = \mathsf{lab}^{g_1}$$

$$\mathsf{suc}^{g_2}(m) = \begin{cases} \mathsf{suc}^{g_1}(m) & \text{if } \mathsf{suc}^{g_1}(m) \neq l \\ r & \text{if } \mathsf{suc}^{g_1}(m) = l \end{cases}$$

(iii) The term graph g_3 is obtained by setting the root node $\widehat{r}$, which is r if $l = r^g$, and otherwise r^g. That is, $g_3 = (g_2|\widehat{r}, \widehat{r})$. This also means that all nodes not reachable from $\widehat{r}$ are removed.

Remark 6.1.4. Recall that, according to Lemma 4.2.5, the $\mathcal{V}$-homomorphism is unique if it exists. Hence, if a rule is applicable to a term graph at a certain node, then the result is uniquely defined.

Subsequently, we will consider rewriting on canonical term graphs only. Therefore, we also adopt the convention described in Remark 4.3.9, i.e. we consider $\mathcal{C}(g_3)$ as the result of the application of a rule instead of g_3.

With the definition of the application of term graph rewrite rules to term graphs we can define the semantics of a term graph rewritings system, viz. the ARS it induces.

Definition 6.1.5 (semantics of IGRSs)
Let $\mathcal{R} = (\Sigma, R)$ be an IGRS.

(i) A *prestep* of $\mathcal{R}$ is a triple (g, n, ρ) consisting of a term graph $g \in G$, a node $n \in N^g$, and a rule $\rho \in R$.

(ii) A prestep $\varphi = (g, n, \rho)$ is called a *step* if ρ is applicable to g at n.

(iii) The ARS *induced* of $\mathcal{R}$, denoted $\mathcal{A}_\mathcal{R}$, is the ARS given by the tuple $(A, \Phi, \mathsf{src}, \mathsf{tgt})$, where A is the set of canonical term graphs $\mathcal{G}^\infty_\mathcal{C}(\Sigma, \mathcal{V})$, and Φ is the set of steps of $\mathcal{R}$. Let $\varphi = (g, n, \rho)$ be a step of $\mathcal{R}$. We define $\mathsf{src}(\varphi) = g$ and $\mathsf{tgt}(\varphi) = g'$, where g' is the result of the application of ρ to g at n.

Notation 6.1.6. Just as for term rewriting system, also every IGRS $\mathcal{R} = (\Sigma, R)$ can be associated with its induced ARS $\mathcal{A}_\mathcal{R}$. That is why we identify $\mathcal{R}$ with $\mathcal{A}_\mathcal{R}$ and consider IGRSs as a special case of ARSs. In particular, we will write $s \to_\mathcal{R} t$ instead of $s \to_{\mathcal{A}_\mathcal{R}} t$. Since the steps of an IGRS additionally contain information about the rule that was applied and the node where it was applied, we use the notation $g \to_{n,\rho} h$ for a step (g, n, ρ) whenever this is appropriate.

Compared to rewriting on terms the notion of term graph rewriting that we use here is slightly different: Instead of removing the redex from the term graph and replacing it with the corresponding instance of the right-hand side of the rule, only those parts of the redex are removed which do not correspond to a part of the left-hand side which is shared with the right-hand side. These parts are then replaced by those parts of the rule which occur only in the right-hand side, i.e. are not reachable from the left root. In this way, as much of the original sharing is maintained as possible. Alternatively, one could have defined term graph rewriting also in accordance with term rewriting: If a rule ρ is applicable to a term graph g at n with matching $\mathcal{V}$-homomorphism φ, then the result of the reduction step is the term graph $g[\rho'_r]_n$, where ρ'_r is the instance of ρ's right-hand side ρ_r according to φ. The instance ρ'_r can be defined by replacing each variable node n in ρ_r by (a copy of) the sub-term graph $g|\varphi(n)$.

In order to illustrate this, let us reconsider the term graph g and the term graph rewrite rule ρ in Figure 6.1. The result of the application of ρ to g at n_1 is shown in Figure 6.2a. We have already seen how this came about. Figure 6.2c, on the other hand, shows the term graph $g[\rho'_r]_{n_1}$, where ρ'_r is the instance of the right-hand side of ρ according to the matching $\mathcal{V}$-homomorphism φ. The primed nodes n'_2, n'_3 and n'_4 constitute a copy of the sub-term graph $g|\varphi(l_1)$ that corresponds to the x-node in ρ. Similarly, the double-primed nodes n''_3 and n''_4 constitute a copy of the sub-term graph $g|\varphi(l_2)$ that corresponds to the y-node in ρ. One can see, that the only difference between the different constructions is the sharing behaviour. In other words: Both alternatives have the same unravelling. The construction by explicit replacements has, however, less sharing. The following lemma confirms this in a more general setting – independent of how the instantiation of the right-hand side is defined:

Lemma 6.1.7 (rewriting by replacement)
Let $g, h \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$, $n \in N^g$, ρ an infinitary term graph rewrite rule over Σ and $g \to_{n,\rho} h$. Then there is a $\mathcal{V}$-homomorphism $\varphi\colon g[\rho_r]_n \to_\mathcal{V} h$.

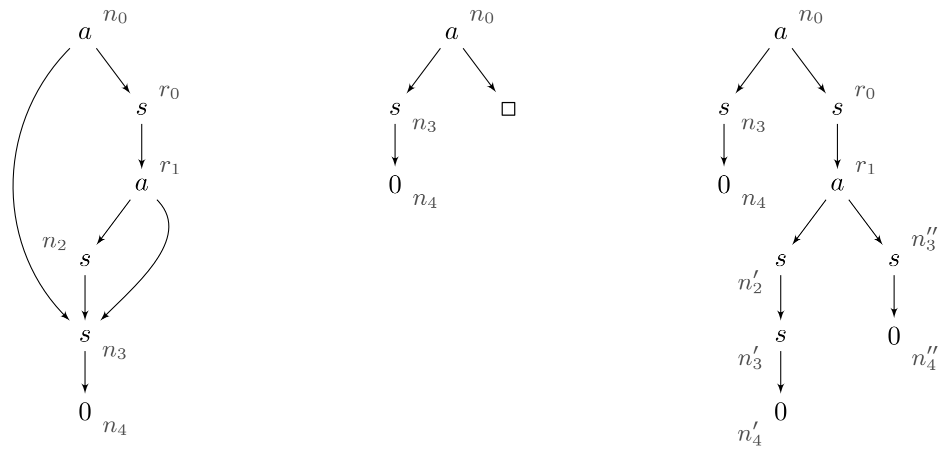

(a) Reduct of g. (b) The context $g[]_{n_1}$. (c) Alternative reduct of g.

Figure 6.2: Alternative definition of term graph rewriting.

Proof. Let $g' = g[\rho_r]_n$ and φ the matching $\mathcal{V}$-homomorphism of the rule application. Define the function $\psi\colon N^{g'} \to N^h$ as follows:

$$\psi(m) = \begin{cases} m & \text{if } m \in N^g \text{ or } m \in N^\rho \setminus N^{\rho_l} \\ \varphi(m) & \text{if } m \in N^{\rho_l} \end{cases}$$

By a careful case analysis following the construction of h, one can show that ψ is indeed a $\mathcal{V}$-homomorphism from g' to h. □

Next, we want to extend the ARS semantics of term graph rewriting to an MRS and a PRS semantics allowing transfinite reduction sequences. To this end, we need to generalise the notions of height and context of a rewrite step to the setting of term graphs. We have to be careful when we do this. A straightforward candidate for the height of a reduction step $\varphi\colon g \to_{n,\rho} h$ is $2^{-\mathsf{depth}_g(n)}$. Similarly, we can easily define the context of φ as $g[\bot]_n$, the replacement of the redex by $\bot$. These candidates for heights and contexts of reduction steps, however, do not yield an MRS or a PRS, respectively. The following example illustrates this:

Example 6.1.8
Consider the term graph rewrite rule

$$\rho\colon \quad \begin{tikzcd} f \arrow[r] \arrow[d] & c \\ x \end{tikzcd}$$

This rule gives rise to the following reduction step:

φ:

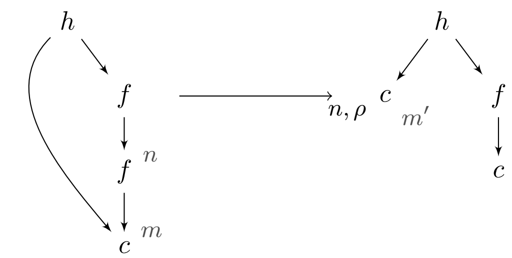

Let g be the term graph before the reduction and g' the term graph afterwards. The depth of the node n at which the reduction takes place is 2. Hence, $2^{-\mathsf{depth}_g(n)} = \frac{1}{4}$. However, the c-node m in g and the corresponding c-node m' in g' have different acyclic sharing. Since these two nodes are at depth 1, the similarity of g and g' is $\mathsf{sim}(g, g') = 1$. Hence, $\mathbf{d}(g, g') = \frac{1}{2}$. Therefore, $2^{-\mathsf{depth}_g(n)}$ cannot be defined to be the height of the rewriting step as this violates the condition on MRS that the height is greater or equal to the distance of the involved objects.

$g[\bot]_n$ is the term graph

$$\begin{tikzcd}
 & h \arrow[dl] \arrow[dr] & \\
c & & f \arrow[d] \\
 & & \bot
\end{tikzcd}$$

For $g[\bot]_n$ to be a legitimate context for the rewrite step φ, we must have $g[\bot]_n \leq_\bot g, g'$. We indeed have (also in general) $g[\bot]_n \leq_\bot g'$. Yet, it does not hold that $g[\bot]_n \leq_\bot g$. Also here the problem is the different acyclic sharing of the respective c-nodes.

The following definition introducing the deletion of a whole sub-term graph and the depth of a sub-term graph provides the necessary tools in order to solve the problem illustrated above:

Definition 6.1.9 (sub-term graph depth, sub-term graph deletion)
Let $g \in \mathcal{G}^\infty(\Sigma_\bot, \mathcal{V})$ and $n \in N^g$.

(i) The *depth* of the sub-term graph $g|n$ in g, denoted $g|n\text{-}\mathsf{depth}(g)$, is the minimal depth of nodes $m \in N^{g|n}$ in g, i.e.

$$g|n\text{-}\mathsf{depth}(g) = \min\left\{\mathsf{depth}_g(m) \,\middle|\, m \in N^{g|n}\right\}.$$

(ii) The *deletion* of the sub-term graph $g|n$ from g, denoted $g \setminus g|n$, is the term graph $(N, \mathsf{lab}, \mathsf{suc}, r^g)$ given by

$$\begin{aligned}
M^{g|n} &= \left\{m \in N^{g|n} \,\middle|\, \exists m' \in N^g \setminus N^{g|n}, i \in \mathbb{N}\colon \mathsf{suc}_i^g(m') = m\right\}\\
N &= (N^g \setminus N^{g|n}) \cup M^{g|n}\\
\mathsf{lab}(m) &= \begin{cases} \mathsf{lab}^g(m) & \text{if } m \in N^g \setminus N^{g|n}\\ \bot & \text{if } m \in M^{g|n}\end{cases}\\
\mathsf{suc}(m) &= \begin{cases} \mathsf{suc}^g(m) & \text{if } m \in N^g \setminus N^{g|n}\\ \varepsilon & \text{if } m \in M^{g|n}\end{cases}
\end{aligned}$$

In the construction of the deletion of a sub-term graph, the set $M^{g|n}$ contains exactly those nodes in the subterm graph which are on the fringe to the rest of the term graph. These are the only nodes in the sub-term graph which are preserved by the deletion. Their sole purpose is to fill the holes that are caused by the deletion of the sub-term graph. That is why they are relabelled with $\bot$.

Example 6.1.10
Reconsider the reduction step discussed in Example 6.1.8. Instead of taking only the depth of n for the definition of the height of the reduction step we are considering the depth of the redex $g|n$ in g: $g|n\text{-}\mathsf{depth}(g)$ is 1 since $\mathsf{depth}_g(m)$ is 1. Defining the height of φ as $2^{-g|n\text{-}\mathsf{depth}(g)}$ yields a height of $\frac{1}{2}$. This satisfies the condition for heights in MRSs as $2^{-g|n\text{-}\mathsf{depth}(g)}$ is greater or equal to the distance $\mathbf{d}(g, g') = \frac{1}{2}$.

Similarly, instead of simply replacing the node n by a $\bot$-node we consider the deletion of the whole redex $g|n$ from g as the context of the reduction step. This yields the context

$$\begin{tikzcd} & h \arrow[dl] \arrow[dr] & \\ \bot & & f \arrow[d] \\ & & \bot \end{tikzcd}$$

Clearly, we have that $g \setminus g|n \leq_\bot g, g'$ and, thus, $g \setminus g|n$ satisfies the condition for contexts in PRSs.

The following lemma confirms that the above observation for the example also holds in general.

Lemma 6.1.11 (sub-term graph deletion)
Let $\varphi\colon g \to_n h$ be a reduction step in an IGRS. Then it holds that $g \setminus g|n \leq_\bot g, h$.

Proof. Let $g' = g \setminus g|n$. Define the function $\varphi\colon N^{g'} \to N^g$ by $m \mapsto m$. It is straightforward to check that φ is a strong $\bot$-homomorphism from g' to g.

Define $\psi\colon N^{g'} \to N^h$ by

$$\psi(m) = \begin{cases} m & \text{if } m \in N^g \setminus N^{g|n} \\ \mathsf{suc}_i^h(m') & \text{if } m \in M^{g|n}, i \in \mathbb{N}, m' \in N^g \setminus N^{g|n} \text{ with } \mathsf{suc}_i^g(m') = m \end{cases}$$

It can be easily checked that ψ is well-defined. Showing that it is a strong $\bot$-homomorphism is straightforward. □

With the above lemma we can safely define the MRS and the PRS that is induced by an IGRS:

Definition 6.1.12 (transfinite semantics of IGRSs)
Let $\mathcal{R} = (\Sigma, R)$ be a IGRS.

(i) The MRS *induced* by $\mathcal{R}$, denoted $\mathcal{M}_\mathcal{R}$, is given by the tuple

$$(\mathcal{G}^\infty_\mathcal{C}(\Sigma, \mathcal{V}), \Phi, \mathsf{src}, \mathsf{tgt}, \mathbf{d}, \mathsf{hgt}),$$

where $(\mathcal{G}^\infty_\mathcal{C}(\Sigma, \mathcal{V}), \Phi, \mathsf{src}, \mathsf{tgt})$ is the ARS $\mathcal{A}_\mathcal{R}$ induced by $\mathcal{R}$, $\mathbf{d}$ is the ultrametric on $\mathcal{G}^\infty_\mathcal{C}(\Sigma, \mathcal{V})$, and hgt is defined as

$$\mathsf{hgt}(\varphi) = 2^{-g|n\text{-}\mathsf{depth}(g)} \quad \text{for each } \varphi\colon g \to_n h$$

(ii) The PRS *induced* by $\mathcal{R}$, denoted $\mathcal{P}_\mathcal{R}$, is given by the tuple

$$(\mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot, \mathcal{V}), \Phi, \mathsf{src}, \mathsf{tgt}, \leq_\bot, \mathsf{cxt}),$$

where $(\mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot, \mathcal{V}), \Phi, \mathsf{src}, \mathsf{tgt})$ is the ARS $\mathcal{A}_{\mathcal{R}'}$ induced by the IGRS $\mathcal{R}' = (\Sigma_\bot, R)$, $\leq_\bot$ is the partial order on $\mathcal{G}^\infty(\Sigma_\bot, \mathcal{V})$, and cxt is defined as

$$\mathsf{cxt}(\varphi) = g \setminus g|n \quad \text{for each } \varphi\colon g \to_n h$$

The following proposition confirms that the above definition indeed yield an MRS and a PRS. Moreover, it shows that both transfinite semantics are complete.

Proposition 6.1.13 (transfinite semantics yields complete URS/PRS)
Each IGRS $\mathcal{R}$ induces a complete URS $\mathcal{M}_\mathcal{R}$ and a complete PRS $\mathcal{P}_\mathcal{R}$.

Proof. At first consider $\mathcal{M}_\mathcal{R}$: $(\mathcal{G}^\infty_\mathcal{C}(\Sigma,\mathcal{V}),\mathbf{d})$ forms, according to Proposition 4.5.16 and Proposition 4.5.19, a complete ultrametric space. Hence, it remains to be shown that $\mathbf{d}(g,h) \le \mathsf{hgt}(\varphi)$ for each reduction step $\varphi\colon g \to_{n,\rho} h$. According to Lemma 6.1.11, we have $g \setminus g|n \le_\bot g, h$ and, therefore, $g \setminus g|n \le_\bot g \sqcap_\bot h$. Hence, we have, due to Lemma 4.2.11, that $\bot\text{-}\mathsf{depth}(g \setminus g|n) \le \bot\text{-}\mathsf{depth}(g \sqcap_\bot h)$, which is equivalent to $\bot\text{-}\mathsf{depth}(g \setminus g|n) \le \mathsf{sim}(g,h)$.

Since g is a total term graph, the only $\bot$-nodes in $g \setminus g|n$ are precisely those in $M^{g|n}$. It is clear that, for each node $m \in N^{g|n}$, there is a node $m' \in M^{g|n}$ with $\mathsf{depth}_g(m') \le \mathsf{depth}_g(m)$. Since $M^{g|n}$ is also a subset of $N^{g|n}$, we have

$$\begin{aligned}\bot\text{-}\mathsf{depth}(g \setminus g|m) &= \min\left\{\mathsf{depth}_{g\setminus g|n}(m) \,\middle|\, m \in M^{g|n}\right\}\\ &= \min\left\{\mathsf{depth}_g(m) \,\middle|\, m \in M^{g|n}\right\}\\ &= \min\left\{\mathsf{depth}_g(m) \,\middle|\, m \in N^{g|n}\right\}\\ &= g|n\text{-}\mathsf{depth}(g)\end{aligned}$$

Hence, we have that $g|n\text{-}\mathsf{depth}(g) \le \mathsf{sim}(g,h)$ and we can conclude that

$$\mathbf{d}(g,h) = 2^{-\mathsf{sim}(g,h)} \le 2^{-g|n\text{-}\mathsf{depth}(g)} = \mathsf{hgt}(\varphi).$$

Next, consider $\mathcal{P}_\mathcal{R}$: Since $(\mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot,\mathcal{V}),\le_\bot)$ forms, according to Proposition 4.4.18, a complete semilattice, it remains to be shown that $\mathsf{cxt}(\varphi) \le_\bot g, h$ holds for each rewrite step $\varphi\colon g \to_n h$. This follows immediately from Lemma 6.1.11. □

Recall that for ITRSs the PRS semantics always extends the corresponding MRS semantics. For term graphs, we can at least show that the PRS semantics of IGRSs *weakly* extends its MRS semantics:

Proposition 6.1.14 (PRS semantics of IGRSs weakly extends MRS semantics)
For each IGRS $\mathcal{R}$, its induced PRS $\mathcal{P}_\mathcal{R}$ weakly extends its induced MRS $\mathcal{M}_\mathcal{R}$.

Proof. We have to show that items (1) - (4) of Definition 3.3.10(ii) hold true. (1) holds by Proposition 4.5.22. Items (2) - (4) follow immediately from Definition 6.1.12. □

Also here we can apply the theory established in Section 3.3.2 in order to identify MRS reductions with total PRS reductions:

Corollary 6.1.15 (total PRS reductions = MRS reductions)
Let $\mathcal{R}$ be an IGRS. Then the following holds for reduction sequences in $\mathcal{R}$:

(i) $S\colon s \hookrightarrow^p \dots$ is total iff $S\colon s \hookrightarrow^m \dots$.

(ii) $S\colon s \hookrightarrow^p t$ is total iff $S\colon a \hookrightarrow^m b$.

Proof. Follows immediately from Proposition 6.1.14 and Proposition 3.3.11. □

It is not known whether the PRS semantics also *strongly* extends the MRS semantics. We conjecture, however, that it does.

6.2 Simulating Term Rewriting

As mentioned in the introduction to this chapter, the original aim of term graph rewriting was to efficiently implement term rewriting. It is efficient in the sense that it allows substructures to be shared when they are needed multiple times instead of being copied. We have seen this in the translation of the duplicating term rewrite rule $d(x) \to a(x,x)$ to a term graph rewrite rule. The goal is to be able to simulate a term reduction $s \to^\star t$ by a term graph reduction $g \to^\star h$ such that $\mathcal{U}(g) = s$ and $\mathcal{U}(h) = t$.

The particular utility of this idea for infinitary rewriting is that it is possible to represent infinite terms by finite term graphs. Instead of *horizontal sharing* that is used for *sharing of subexpressions* [Wad71], *vertical sharing*, i.e. cyclic term graphs, allows to represent infinite terms by finite term graphs. A term t is *represented* by a term graph g if $t = \mathcal{U}(g)$. The simplest example for vertical sharing is the term graph consisting of a single f-node and a loop which represents the term f^ω. The principle motivation in this setting is to be able to simulate a transfinite reduction $s \twoheadrightarrow t$ involving possibly infinite terms by a finite reduction $g \to^* h$ on finite term graphs. This is, however, not always possible since not every infinite term is representable by a finite term graph.

Definition 6.2.1 (rational term)
A term $t \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$ is called *rational* if there is a finite term graph $g \in \mathcal{G}(\Sigma, \mathcal{V})$ such that $t = \mathcal{U}(g)$.

Infinite rational terms can be represented by finite term graphs through cycles. These cycles are, however, only capable of repeating the same structure over and over again. Hence, rational terms are precisely those which have only finitely many different subterms:

Theorem 6.2.2 (rational terms, [Cou83])
A term t is rational iff it has finitely many distinct subterms, i.e. $\mathcal{S}(t)$ is finite.

Example 6.2.3 (rational terms)
Consider the infinite term $c(0, c(0, c(0, \ldots$ that represents the infinite list $[0, 0, 0, \ldots]$. This term has only two different subterms, viz. 0 and the term $c(0, c(0, c(0, \ldots$ itself. Hence, this term can be represented by a term graph consisting of only two nodes:

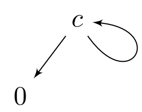

That is, $c(0, c(0, c(0, \ldots$ is a rational term. By contrast, the term $c(0, c(s(0), c(s^2(0), \ldots,$ that represents the infinite list $[0, s(0), s^2(0), \ldots,$ has infinitely many subterms, viz. for each $n \in \mathbb{N}$ the terms $s^n(0)$ and $c(s^n(0), c(s^{(n+1)}(0), c(s^{(n+2)}(0), \ldots$. It is not possible to represent this term by a finite term graph. For each subterm, one needs a node in the term graph:

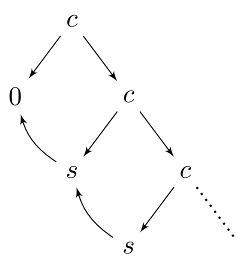

We have intuitively used the notion of translating a term rewrite rule to a term graph rewrite rule. The following definition makes this explicit:

Definition 6.2.4 (unravelling of IGRSs)
Let $\mathcal{R} = (\Sigma, R)$ be an IGRS.

(i) Let $\rho \in R$. The *unravelling* of ρ, denoted $\mathcal{U}(\rho)$, is the term rewrite rule $\mathcal{U}(\rho_l) \to \mathcal{U}(\rho_r)$.

(ii) The *unravelling* of $\mathcal{R}$, denoted $\mathcal{U}(\mathcal{R})$, is the ITRS $(\Sigma, \mathcal{U}(R))$, where

$$\mathcal{U}(R) = \{\mathcal{U}(\rho) \mid \rho \in R\}.$$

So we can say that a term rewrite rule ρ is translated to a term graph rewrite rule ρ' if $\rho = \mathcal{U}(\rho')$. It is obvious that a term rewrite rule might have several different translations and we will see the ramifications of the different choices that are available.

As we have mentioned previously, cycles in term graphs can be used represent infinite terms. Cycles may also be introduced by term graph rewriting. The following definition aims to capture when this happens.

Definition 6.2.5 (cyclicity of IGRSs)
Let $\mathcal{R} = (\Sigma, R)$ be an IGRS.

(i) A rule $\rho \in R$ is called *cyclic* if the underlying graph g of ρ is cyclic or contains a non-empty path from the right root to the left root; otherwise, it is called *acyclic*.

(ii) $\mathcal{R}$ is called *cyclic* if some rule in $\mathcal{R}$ is cyclic. Otherwise $\mathcal{R}$ is called *acyclic*.

The above definition of cyclicity of term graph rewrite rules is non-standard. Besides the standard notion of cyclicity of the underlying graph it also includes an implicit cyclicity that is caused by an edge from the right-hand side of the rule to the rule's left root. The following example illustrates this implicit cyclicity:

Example 6.2.6
Consider the term graph rewrite rule

ρ:

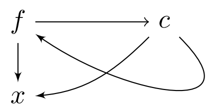

Note that the unravelling of ρ is $\mathcal{U}(\rho)$: $f(x) \to c(x, f(x))$ or $\mathcal{U}(\rho)$: $f(x) \to x : f(x)$ if c is written as an infix :. If : is interpreted as the list constructor, then the reducts of $f(t)$ by $\mathcal{U}(\rho)$ are increasingly long lists containing t. In terms of infinitary rewriting, $f(t)$ strongly converges to $[t, t, \dots]$ in ω steps.

If we apply ρ to the term $f(t)$, we obtain the reduction step

φ:

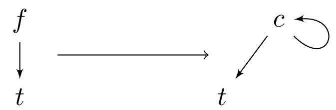

The cycle in the result term h arises in the redirection step of the rewrite construction:

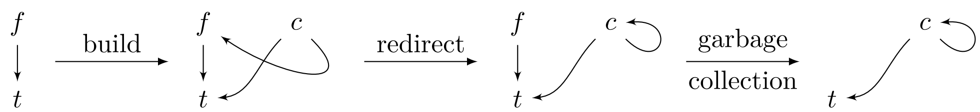

Although neither the term $f(t)$ nor the applied rule ρ contains a cycle, the result h of the reduction step is cyclic. Also note that the unravelling $\mathcal{U}(h)$ of the result is the infinite list $[t, t, \dots]$ that is obtained by infinitary term rewriting.

The phenomenon that we have seen above occurs whenever there is a non-empty path from the right root to the left root. It is an artifact of the construction of the result term graph. One could have defined the rewriting construction slightly different such that this phenomenon does not appear. Farmer and Watro observed this behaviour also in their variant of term graph rewriting [FW90]. It occurs when the right-hand side contains a reference to the left-hand side. Thereby, the right-hand side *captures* the redex during the application of the rule. Therefore, it is called *redex capturing*:

Definition 6.2.7 (redex capturing)
Let $\rho = (g, l, r)$ be a term graph rewrite rule. ρ is said to *capture its redexes* if g contains a non-empty path from r to l. The *cyclisation* of ρ is the term graph rewrite rule $\rho' = (g', l, r)$ given by

$$N^{g'} = N^g \qquad \mathsf{lab}^{g'} = \mathsf{lab}^g$$

$$\mathsf{suc}^{g'}(n) = \begin{cases} r & \text{if } \mathsf{suc}^g(n) = l \text{ and } n \in N^{g|r} \\ \mathsf{suc}^g(n) & \text{otherwise} \end{cases}$$

That is, edges going from the right-hand side to the left root are redirected to the right root.

It is clear that the cyclisation of the rules of a IGRS does not change its semantics. It just makes the implicit cycles caused by redex capturing explicit

Fact 6.2.8 (rewriting is invariant to cyclisation)
Let ρ be a term graph rewrite rule and ρ' its cyclisation. Then it holds that

$$g \to_{n,\rho} h \quad \textit{iff} \quad g \to_{n,\rho'} h$$

Example 6.2.9
Reconsider the term graph rewrite rule ρ from Example 6.2.6. As we have seen, ρ captures its redexes. The cyclisation of ρ is the term graph rewrite rule

ρ':
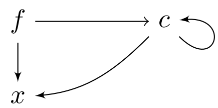

ρ' yields the same reductions as ρ. For example, if applied to $f(t)$, we obtain the same reduction step φ. The construction, however, is slightly different, of course:

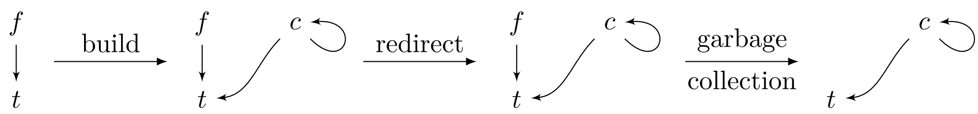

Redex capturing is not always desirable, e.g. when we want to simulate finitary rewriting. As we have seen in Example 6.2.6, a single reduction step in an GRS $\mathcal{R}$ might take infinitely many steps in the TRS $\mathcal{U}(\mathcal{R})$.

On the other hand, this is certainly desirable whenever we want to represent transfinite term reduction sequences by finite term graph reduction sequences. The most prominent example of an application of this method is the implementation of the fixed point combinator Y in function programming languages (cf.[PJ87] and [Tur79]):

Example 6.2.10
In an applicative language, Y is defined by the rewrite rule $\rho_0\colon Y f \to f(Y f)$. The f is interpreted as a variable and function application is represented by juxtaposition. Written as a term graph rewrite rule ρ_0 becomes

ρ_1:
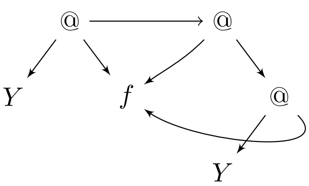

The symbol @ represents function application explicitly. One can see that in ρ_1 (as well as in ρ_0) the left-hand side occurs as a subterm in the right-hand side. Hence, it can be more concisely represented by

ρ_2:

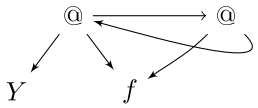

ρ_2 captures its redexes. Note that we still have $\mathcal{U}(\rho_1) = \mathcal{U}(\rho_2) = \rho_0$. The (equivalent) cyclisation of ρ_2 is the term graph rewrite rule

ρ_3:

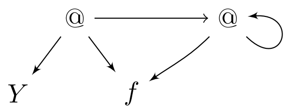

The term graph rewrite rule ρ_3 is used in most functional programming languages to implement the fixed point combinator. However, note that $\mathcal{U}(\rho_3) \neq \rho_0$.

The approach illustrated in Example 6.2.10 can be applied to any term rewrite rule, in particular term rewrite rules of the form $l \to C[l, \dots, l]$ whose right-hand side contains the left-hand side as a subterm – possibly multiple occurrences of it. This rule can then be translated straightforwardly to a term graph rewrite rule similar to ρ_1 in Example 6.2.10. Then one can compute a variant of the obtained rule, i.e. one with the same unravelling, having maximal sharing. This can be achieved by employing a congruence closure algorithm (cf. [DST80] and [Bah07]). Eventually, a term graph rewrite rule is obtained which captures its redexes. For further optimisation, one can then compute its cyclisation.

Next we want to discuss the adequacy of term graph rewriting for simulating term rewriting. That is, we would like to have for each term rewriting system $\mathcal{R}$ and each of its translation into a graph rewriting system $\mathcal{G}$ with $\mathcal{U}(\mathcal{G}) = \mathcal{R}$, that $\mathcal{G}$ simulates $\mathcal{R}$. Ideally, we this includes both soundness, i.e. that $g \twoheadrightarrow_{\mathcal{G}} h$ implies $\mathcal{U}(g) \twoheadrightarrow_{\mathcal{R}} \mathcal{U}(h)$, and completeness, i.e. that $\mathcal{U}(g) \twoheadrightarrow_{\mathcal{R}} \mathcal{U}(h)$ implies $g \twoheadrightarrow_{\mathcal{G}} h$. One can easily see that full completeness can not be satisfied since term graph rewriting is in general coarser that term rewriting. In term rewriting, each redex can be contracted individually. In term graph rewriting, however, a redex in a term graph can, through sharing, represent multiple redexes (in its unravelling). Contracting a redex in a term graph might, therefore, correspond to contractions of multiple redexes in the corresponding term. To this end, a weaker notion of completeness is considered.

The analysis of this adequacy of term graph rewriting is restricted to orthogonal systems. Orthogonality of term graph rewriting systems is similar to the corresponding concept in term rewriting:

Definition 6.2.11 (left-linearity, orthogonality)
Let $\mathcal{R}$ be an IGRS and ρ an infinitary term graph rewrite rule.

(i) ρ is called *left-linear* if its left-hand side is a term tree. $\mathcal{R}$ is called *left-linear* if each rule in $\mathcal{R}$ is left-linear.

(ii) $\mathcal{R}$ is called *orthogonal* if it is left-linear, and its unravelling $\mathcal{U}(\mathcal{R})$ does not have any critical pairs.

Non-left linearity in term graph rewriting is in some sense stronger than in term rewriting: Equality in term rewriting can be enforced by having multiple occurrences of the same variable in the left hand side. Equality in term graph rewriting can be enforced by sharing. This is however not restricted to variables. Since, for the matching of the left-hand side with the redex, a $\mathcal{V}$-homomorphism is employed, matching requires the redex to have at least the "same amount" of sharing as the left-hand side.

$$\begin{tikzcd}
\mathcal{U}(g) \arrow[r, "*", "\mathcal{U}(\mathcal{R})"'] & t \arrow[r, dashed, "*", "\mathcal{U}(\mathcal{R})"'] & \mathcal{U}(h) \\
g \arrow[rr, dashed, "*", "\mathcal{R}"'] & & h
\end{tikzcd}$$

(a) Acyclic term graphs.

$$\begin{tikzcd}
\mathcal{U}(g) \arrow[r, twoheadrightarrow, "r", "\mathcal{U}(\mathcal{R})"'] & t \arrow[r, twoheadrightarrow, dashed, "r", "\mathcal{U}(\mathcal{R})"'] & \mathcal{U}(h) \\
g \arrow[rr, dashed, "*", "\mathcal{R}"'] & & h
\end{tikzcd}$$

(b) Cyclic term graphs.

Figure 6.3: Weak completeness of term graph rewriting.

Remark 6.2.12. Since term graph rewrite rules are only allowed to have at most a single v-node for each variable $v \in \mathcal{V}$, restricting the left-hand side to have a tree structure, which means that it is isomorphic to its own unravelling, causes the unravelling of left-linear term graph rewrite rules to be left-linear term rewrite rules. Hence, also the unravelling of orthogonal IGRSs results in orthogonal ITRSs.

At first we deal with the simulation of finitary term rewriting. Since cycles in term graphs can cause a contraction of a single redex to represent infinitely many contractions in the corresponding term rewriting system, we have to restrict the analysis to acyclic (I)GRSs and acyclic term graphs.

The soundness of term graph rewriting is straightforward:

Theorem 6.2.13 (soundness of acyclic term graph rewriting, [KKSdV94]) *Let $\mathcal{R}$ be a left-linear acyclic IGRS over Σ and $g \in \mathcal{G}^\infty_\mathcal{C}(\Sigma, \mathcal{V})$ an acyclic term graph. If $g \to^*_\mathcal{R} h$, then $\mathcal{U}(g) \to^*_{\mathcal{U}(\mathcal{R})} \mathcal{U}(h)$.*

As we have already mentioned, full completeness is beyond the abilities of term graph rewriting. Hence, we have to be satisfied with a weaker notion. It includes that normal forms are preserved, and that term reductions can be simulated by a term graph reduction which might overshoot the mark, but only so far such that term rewriting can catch up with it; cf. Figure 6.3a.

Theorem 6.2.14 (completeness of acyclic term graph rewriting, [KKSdV94]) *Let $\mathcal{R}$ be an orthogonal acyclic IGRS over Σ and $g \in \mathcal{G}^\infty_\mathcal{C}(\Sigma, \mathcal{V})$ an acyclic term graph. Then the following holds:*

(i) g is a normal form in $\mathcal{R}$ iff $\mathcal{U}(g)$ is a normal form in $\mathcal{U}(\mathcal{R})$.

(ii) If there is a finite reduction $\mathcal{U}(g) \to^*_{\mathcal{U}(\mathcal{R})} t$, then there are finite reductions $g \to^*_\mathcal{R} h$ and $t \to^*_{\mathcal{U}(\mathcal{R})} \mathcal{U}(h)$; cf. Figure 6.3a.

This overshooting of term graph rewriting can be seen in the example illustrated in Figure 6.4: Let ρ be the term graph rewrite rule $c \to d$. The term graph in the lower left corner represents the term $f(c, c)$ by sharing a single c-node. The reduction step in the upper left corner contracts the left occurrence of c by the term rewrite rule $\mathcal{U}(\rho)$: $c \to d$. Since the shared c-node in the corresponding term graph represents two $\mathcal{U}(\rho)$-redexes, the result of its contraction by the term graph rewrite rule ρ yields a term graph representing the term $t(d, d)$. That is, both $\mathcal{U}(\rho)$-redexes were contracted. Hence, in order to catch up on the term rewriting side, also the second $\mathcal{U}(\rho)$-redex has to be contracted.

Next, we consider term graph rewriting of cyclic term graphs. The problem that arises in cyclic term graphs is that a node might not only represent finitely many term redexes through horizontal sharing, but it can also represent infinitely many term redexes through vertical sharing. Hence, even a single term graph rewriting step might simulate infinitely many term rewriting steps. That is why the adequacy of cyclic term graph rewriting is regarded w.r.t. infinitary term rewriting. The soundness of this kind of simulation is straightforward.

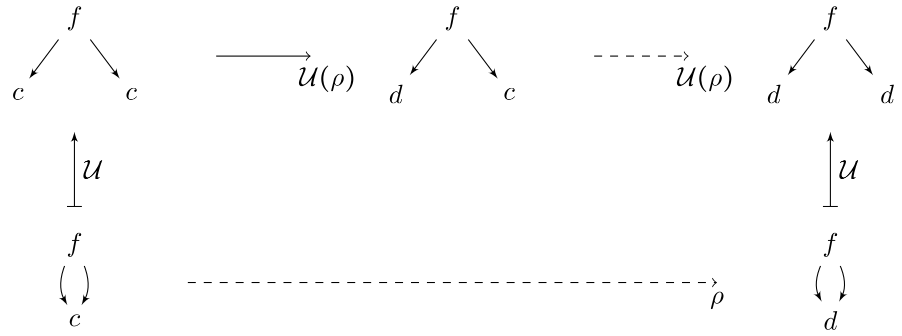

Figure 6.4: Example for weak completeness of term graph rewriting.

Unfortunately, orthogonal term graph rewriting is not adequate for arbitrary infinitary term rewriting. In order to obtain an adequacy result, one has to consider a subset of all possible transfinite reductions:

Definition 6.2.15 (rational reduction, [KKSdV94])
Let $\mathcal{R}$ be an ITRS over Σ and t a rational term in $\mathcal{T}^\infty(\Sigma, \mathcal{V})$.

(i) Let $U \subseteq \mathcal{P}(t)$ be a set of occurrences in t. U is called *rational* if the labelling $t^{(U)}$ of t is still rational.

(ii) The set of *rational reduction sequences* in $\mathcal{R}$ is the smallest set $\mathcal{S}$ satisfying the following conditions:

(a) If T is a complete development of a set of rational redex occurrences in a rational term and T is of length at most ω, then $T \in \mathcal{S}$.

(b) If $T_0, \ldots, T_n \in \mathcal{S}$, then also $\prod_{i \le n} T_i \in \mathcal{S}$.

We use the notation $S\colon s \twoheadrightarrow^r t$ to indicate that a reduction sequence $S\colon s \twoheadrightarrow t$ is rational.

Indeed, cyclic term graph rewriting is also sound w.r.t. rational term graph rewriting:

Theorem 6.2.16 (soundness of cyclic term graph rewriting, [KKSdV94])
Let $\mathcal{R}$ be a left-linear IGRS over Σ and $g \in \mathcal{G}^\infty_\mathcal{C}(\Sigma, \mathcal{V})$. If $g \to^*_\mathcal{R} h$, then $\mathcal{U}(g) \twoheadrightarrow^r_{\mathcal{U}(\mathcal{R})} \mathcal{U}(h)$.

Another aspect one has to take into account is the fact that (infinitary) confluence is essential for weak completeness of term graph reduction w.r.t. (infinitary) term rewriting, whenever the signature has a symbol of arity 2 or higher:

Let $\mathcal{R}$ be an IGRS and t_0 a term that is infinitarily reducible to t_1 and t_2 in the ITRS $\mathcal{U}(\mathcal{R})$. Consider the term graph depicted in the lower left corner of Figure 6.5. Its unravelling is the term $f(t_0, t_0)$ which is infinitarily reducible to $f(t_1, t_2)$. If weak completeness holds, then there is a term graph reduction as shown in the lower half of Figure 6.5 and an infinitary term reduction to the unravelling of the result term graph, shown in the upper right corner of the picture. Provided that there is no rule in $\mathcal{R}$ having a left root labelled with f, the reduction $f(t_1, t_2) \twoheadrightarrow f(t_3, t_3)$ is only possible if there are reductions $t_1 \twoheadrightarrow t_3$ and $t_2 \twoheadrightarrow t_3$. Note that we assume here arbitrary transfinite reductions and not only regular ones. This abstract example should only illustrate the motivation for the restrictions that are imposed to the analysis of completeness of cyclic term graph reduction.

The same argument can be made for acyclic term graph rewriting and its weak completeness w.r.t. finitary term rewriting. Orthogonal ITRSs are finitarily confluent (cf. Theorem 2.3.31). They are, however, not necessarily infinitarily confluent (cf. Example 5.4.22). We have seen that orthogonal ITRSs are infinitarily confluent iff they are almost non-collapsing (cf. Theorem 5.4.28). Therefore, we also need a corresponding notion for IGRSs:

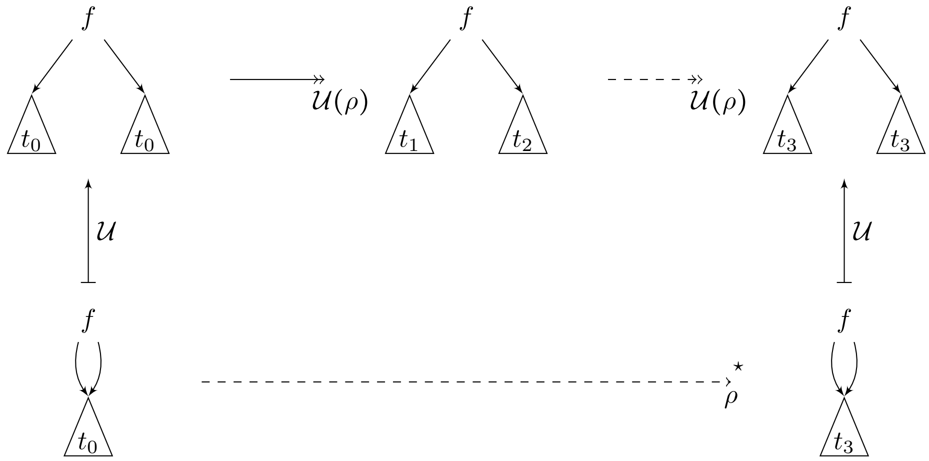

Figure 6.5: Necessity of (infinitary) CR for weak completeness of term graph rewriting.

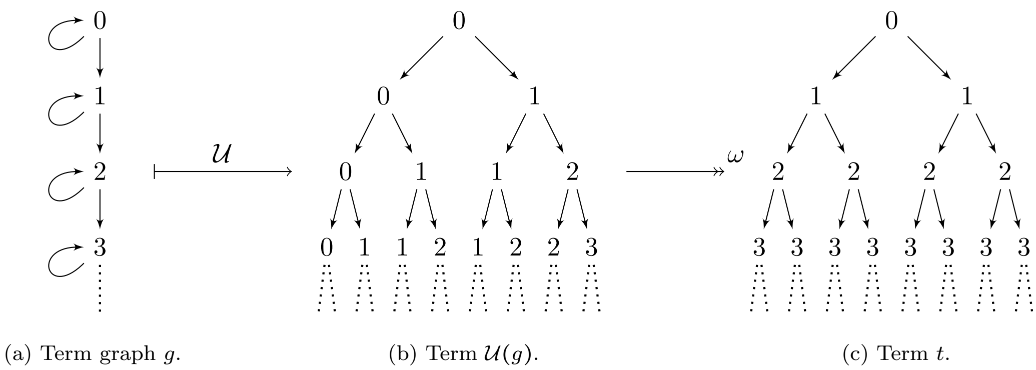

Figure 6.6: Infinitary term graph rewriting vs. infinitary term rewriting.

Definition 6.2.17 (almost non-collapsing)
An IGRS $\mathcal{R}$ is called *almost non-collapsing* if its unravelling $\mathcal{U}(\mathcal{R})$ is almost non-collapsing.

The unravelling of an almost non-collapsing orthogonal IGRS is by definition also almost non-collapsing and orthogonal (cf. Remark 6.2.12), and is, therefore, infinitarily confluent according to Theorem 5.4.28.

The following theorem shows that weak completeness holds for almost non-collapsing IGRSs:

Theorem 6.2.18 (weak completeness of acyclic term graph rewriting, [KKSdV94])
Let $\mathcal{R}$ be an orthogonal and almost non-collapsing IGRS over Σ and $g \in \mathcal{G}^\infty_\mathcal{C}(\Sigma, \mathcal{V})$. Then the following holds:

(i) g is a normal form in $\mathcal{R}$ iff $\mathcal{U}(g)$ is a normal form in $\mathcal{U}(\mathcal{R})$.

(ii) If there is a rational reduction $\mathcal{U}(g) \twoheadrightarrow^r_{\mathcal{U}(\mathcal{R})} t$, then there are reductions $g \to^*_\mathcal{R} h$ and $t \twoheadrightarrow^r_{\mathcal{U}(\mathcal{R})} \mathcal{U}(h)$; cf. Figure 6.3b.

Weak completeness does not hold for arbitrary transfinite term reductions – even if we allow transfinite term graph reductions. The following example illustrates this:

Example 6.2.19 ([KKSdV94])
Consider the signature Σ which contains for each natural number $n \in \mathbb{N}$ a binary symbol $\overline{n}$ and the GRS $\mathcal{R}$ over Σ containing for each $n \in \mathbb{N}$ a rule

ρ_n:
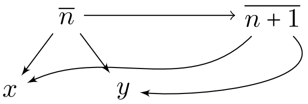

The unravelling of ρ_n is the term rewrite rule $\mathcal{U}(\rho_n)\colon \overline{n}(x,y) \to \overline{n+1}(x,y)$. That is, each symbol n in a term can be "incremented" by applying the term graph rewrite rule $\mathcal{U}(\rho_n)$.

Now consider the term graph g shown in Figure 6.6a and its unravelling $\mathcal{U}(g)$ shown in Figure 6.6b. It is easy to see that one can reduce $\mathcal{U}(g)$ to the term t in the TRS $\mathcal{U}(\mathcal{R})$ in ω steps. Each node at depth n in t has the label $\overline{n}$ for all $n \in \mathbb{N}$. Hence, for each $n \in \mathbb{N}$ there are only finitely many nodes in t labelled with $\overline{n}$. In order to have weak completeness, there has to be a term graph h with $g \twoheadrightarrow_{\mathcal{R}} h$ and $t \twoheadrightarrow_{\mathcal{U}(\mathcal{R})} \mathcal{U}(h)$. However, g does solely (infinitarily) reduce in $\mathcal{R}$ to term graphs which only differ from g in the labelling. In particular, each reduct also contains loops which means that there is some n such that the unravelling of the reduct contains infinitely many nodes labelled with $\overline{n}$. On the other hand, each term that t can be (infinitarily) reduced to has only finitely many nodes labelled with $\overline{n}$ for each $n \in \mathbb{N}$.

The example above uses an infinite signature. One can, however translate this to a system with only finitely many symbols.[3] To this end, we consider the signature $\Sigma' = \{f/2, g/1, h/1\}$. In this way, each symbol n in Σ can be encoded in Σ' by $h(g^n(f(\cdot,\cdot)))$. The rewrite rules in $\mathcal{R}$ can then be mimicked by the single rule

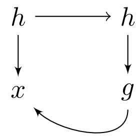

[3]In [KKSdV94] a similar translation to a system over a finite signature is given. For their translation, however, the counterexample does not work anymore since it allows terms of the form $g^n(f(\cdot,\cdot))$ to be reduced to g^ω by an transfinite reduction.

Chapter 7

Conclusions

In this thesis we have discussed various fields of infinitary rewriting. The focus of our work was set on term rewriting and related areas such as term graph rewriting. It should be pointed out though that also other reduction systems were analysed for the properties of their transfinite reductions. This includes infinitary versions of λ-*calculi* [Nak75, KKSdV95b, KKSdV97, BDC99, KvOdV99, Blo04, BK09] and *combinatory reduction systems* [KS05b, KS05a, KS06, Ket08, Ket09] as well as *stream definitions* [Sij89, EGH$^+$07, Zan09]. Moreover, most of the results that we have developed and presented in this thesis are restricted to (almost) orthogonal systems. Yet, in the literature also non-orthogonal systems were considered (cf. [IN91, GL06]). The analysis of infinitary systems that we have conducted here was primarily concerned with the fundamental characteristics of transfinite reduction sequences as well as with infinitary termination and confluence properties. Recently, however, also the modular behaviour of infinitary properties was investigated [Kah09, Sim06]. In the following, we briefly summarise the results of our work, and suggest promising directions for further research. Finally, we also outline possible applications.

7.1 Summary

In this thesis we have discussed several different approaches for modelling transfinite reduction sequences. This analysis was chiefly restricted to the well-known metric approach (MRSs) and our novel method using partial orders (PRSs). We have seen that both approaches yield transfinite reductions which are intuitive and share some of the properties that we know from finite reduction sequences. In addition, we have shown that termination and confluence properties lifted to the infinitary setting exhibit a behaviour similar to that known from the corresponding properties in the finitary setting.

We have also given a comparison between weakly and strongly convergent reductions which both models – the metric and the partial order model – are able to distinguish. It was argued that the strong variant describes the intuition of infinitary term and term graph rewriting more accurately as it models the result of infinite reductions in terms of stable parts of the intermediate objects. In particular, we think that this intuition of infinitary rewriting is manifested most strikingly in strongly convergent PRS reductions: The result of an infinite reduction sequence is quite literally the largest part of the term resp. term graph that eventually remains untouched. Our finding that MRS reductions are precisely the total PRS reductions – at least in the setting of term and term graph rewriting – shows that the metric approach provides the same intuition and can be considered simply as a more restrictive variant of the partial order method.

Most importantly, we were able to establish an equivalence of strongly convergent PRS reductions of a term rewriting system and strongly convergent MRS reductions in the corresponding Böhm reduction w.r.t. root-active terms. This is another indication which shows that the partial order method constitutes a natural concept of transfinite reductions. The use of sets of meaningless terms and its induced Böhm reductions is a powerful but seemingly rather artificial construction. The equivalence of both notions of reduction shows that Böhm reductions occur naturally if the more intuitive model using partial orders is employed.

Surprisingly, despite their finer structure, strongly convergent PRS reductions possess – compared to strongly convergent MRS reductions – more advantageous properties. In particular, orthogonal systems, unlike in the metric model, allow arbitrary complete developments and are infinitary confluent, which generalises results known from finitary orthogonal term rewriting. Moreover, orthogonal systems are infinitarily normalising. Similarly to strongly convergent MRS reductions, the partial order approach allows to simulate every transfinite reduction in at most ω steps and to approximate the result of a transfinite reduction arbitrarily well by a finite reduction.

In order to apply the two models of transfinite reductions also to term graph rewriting, we have introduced a partial order and a metric on term graphs. Both are extensions of the corresponding concepts on terms and we were able to establish that they also have the same fundamental properties as their term counterparts. More precisely, the metric was shown to be a complete ultrametric and the partial order was shown to form a complete semilattice. With the help of these tools we have introduced corresponding models of infinitary term graph rewriting. Beyond that, we have argued the ability of term graph rewriting to implement infinitary term rewriting in restricted cases.

In a nutshell, our analysis has shown that employing partial orders to describe transfinite reductions constitutes a powerful model which is superior to the classic metric method in many aspects including its formal properties and its ability to capture the intuition of transfinite reductions. Nonetheless, it nicely subsumes the metric notion of transfinite reductions. Therefore, we think that this framework constitutes an attractive basis for further research and an interesting object of investigation itself.

7.2 Future Work and Possible Applications

Since several new concepts were developed in this thesis, which could not be analysed in their full depth, there are a number of open questions and promising new directions for further research. First and foremost, beyond abstract properties, the partial order framework for transfinite reductions was analysed only in the setting of orthogonal term rewriting systems – and only for its strongly convergent reductions. We have seen that weakly convergent MRS reductions have comparatively unsatisfying properties. We did not investigate weakly convergent PRS reductions and it might be the case that in this setting weak convergence has nicer properties.

This thesis also introduced infinitary term graph rewriting employing both a metric and a partial order on term graphs. We did not study any particular properties of infinitary term graph rewriting. It might be promising to investigate whether some results known from finitary term graph rewriting extend to infinitary term rewriting. This includes, in particular, confluence and the existence of complete developments for orthogonal systems (cf. [BvEG+87]). It is also interesting to explore whether some results known for infinitary term rewriting such as the compression property and infinitary normalisation can be generalised to the term graph setting. Moreover, we do not know whether the PRS semantics of IGRSs not only weakly but also strongly extends the corresponding MRS semantics as it does in the case of ITRSs. We do, however, conjecture that it does.

We have seen that, on the one hand, finitary (cyclic) term graph rewriting is adequate for simulating rational term rewriting but that, on the other hand, not even infinitary term graph rewriting is adequate for simulating infinitary term rewriting. This raises the question which subset of infinitary term rewriting is infinitary term graph rewriting able to simulate adequately. A more practically oriented question is how to transform term rewriting systems into term graph rewriting systems in order to be able to simulate at least some transfinite reductions. We have informally discussed a very simple approach which is folklore in functional programming language implementations (cf. [PJ87, Tur79]). Finding other techniques that allow employing term graph rewriting to implement infinitary term rewriting is of great importance for its practical relevance.

In this thesis we favoured the term graph rewriting framework of Barendregt et al. [BvEG⁺87]. It has been shown to be at least to some extent adequate in order to simulate strongly convergent MRS reductions of ITRSs. As we have mentioned, there are many other approaches to term graph rewriting. There are some indications suggesting that the *double-pushout approach* [EPS73] and the *equational approach* [AK96] are more appropriate for simulating strongly convergent PRS reductions. Corradini et al. [Cor93, CD97, CG95] also investigated infinite reductions based on a partial order on terms and showed that term graph rewriting in the double-pushout framework is able to adequately simulate their notion of infinitary parallel term rewriting. We conjecture similar results for strongly convergent PRS reductions of ITRSs and term graph rewriting in the double-pushout and the equational framework.

For implementing infinitary term rewriting using term graph rewriting, it is necessary to know the closure properties of the set of rational terms. In particular, it would be advantageous to establish criteria on ITRSs which assure that at least normalising transfinite reductions preserve rationality, i.e. normal forms reachable from rational terms – possibly through transfinite reductions – are also rational. For quite restrictive systems having only constants as left-hand sides, so-called *regular equations*, this preservation holds [Cou83]. A comparable result was presented by Kennaway et al. [KKSdV94] who showed that rational reductions preserve rational terms. The intention of introducing infinitary term graph rewriting is to provide a tool that helps identifying closure properties of rational terms.

We have introduced a partial order on term graphs which extends the usual partial order on terms. The partial order on term graphs also has similar properties, most importantly it also forms a complete semilattice. Böhm trees can be defined using direct approximants [Lév78, Ket06]. This technique employs the partial order structure on terms. With the partial order on term graphs one might be able to define a similar notion of *"Böhm graphs"* for term graph rewriting systems.

As we have mentioned, also λ-calculi and combinatory reduction systems were investigated for their transfinite reductions. The research on these higher-order systems is currently almost entirely limited to the metric model. Only Blom [Blo04] considers a partial order model similar to ours. However, this approach is only able to model strongly convergent reductions. It is desirable to apply our partial order framework of infinitary rewriting to these higher-order systems. Unfortunately, the partial order on terms is not appropriate for higher-order terms. The reason for this is that it would allow reduction sequences to converge to higher-order terms which are deemed meaningless, e.g. terms having infinitely nested meta-variables or λ-abstractions. The same problem also arises for the metric model. Usually, the considered infinitary higher-order systems use various alternative metrics by employing different depth measurements which ignore certain edges of term trees. In order to apply the partial order model to higher-order systems, this idea has to be conveyed to the definition of the partial order on terms for the purpose of obtaining appropriate partial orders on higher-order terms.

Bibliography

[AK96] Zena M. Ariola and Jan Willem Klop. Equational term graph rewriting. *Fundamenta Informaticae*, 26(3-4):207–240, 1996.

[AKK+94] Zena M. Ariola, Richard Kennaway, Jan Willem Klop, M. Ronan Sleep, and Fer-Jan de Vries. Syntactic definitions of undefined: On defining the undefined. In Masami Hagiya and John C. Mitchell, editors, *Theoretical Aspects of Computer Software*, volume 789 of *Lecture Notes in Computer Science*, pages 543–554. Springer Berlin / Heidelberg, 1994. International Symposium TACS '94 Sendai, Japan, April 19–22, 1994, Proceedings.

[AM96] Sergio Antoy and Aart Middeldorp. A sequential reduction strategy. *Theoretical Computer Science*, 165(1):75–95, 1996.

[AN80] André Arnold and Maurice Nivat. The metric space of infinite trees. Algebraic and topological properties. *Fundamenta Informaticae*, 3(4):445–476, 1980.

[Bah07] Patrick Bahr. Implementation of a fast congruence closure algorithm. Technical report, Technische Universität Dresden, Dresden, 2007.

[Bar84] Henk P. Barendregt. *The Lambda Calculus: Its Syntax and Semantics*, volume 103 of *Studies in Logic and the Foundations of Mathemantics*. Elsevier Science, revised edition edition, 1984.

[Bar92] Henk P. Barendregt. Representing 'undefined' in the lambda calculus. *Journal of Functional Programming*, 2:367–374, 1992.

[BDC99] Alessandro Berarducci and Mariangiola Dezani-Ciancaglini. Infinite λ-calculus and types. *Theoretical Computer Science*, 212(1-2):29–75, 1999.

[Ber96] Alessandro Berarducci. Infinite λ-calculus and non-sensible models. In A. Ursini and P. Aglianó, editors, *Logic and algebra*, number 180 in Lecture Notes in Pure and Applied Mathematics, pages 339–378. CRC Press, 1996.

[BK09] Henk P. Barendregt and Jan Willem Klop. Applications of infinitary lambda calculus. *Information and Computation*, 207(5):559–582, 2009. From Type Theory to Morphological Complexity: Special Issue dedicated to the 60th Birthday Anniversary of Giuseppe Longo.

[Blo04] Stefan Blom. An approximation based approach to infinitary lambda calculi. In Vincent van Oostrom, editor, *Rewriting Techniques and Applications*, volume 3091 of *Lecture Notes in Computer Science*, pages 221–232. Springer Berlin / Heidelberg, 2004. RTA '04, Aachen, Germany, June 3-5, 2004, Proceedings.

[BN98] Franz Baader and Tobias Nipkow. *Term Rewriting and All That*. Cambridge University Press, 1998.

[BvEG+87] Henk P. Barendregt, Marko C.J.D. van Eekelen, John R.W. Glauert, Richard Kennaway, Marinus J. Plasmeijer, and M. Ronan Sleep. Term graph rewriting. In Philip C. Treleaven Jaco de Bakker, A. J. Nijman, editor, *Parallel Architectures and Languages Europe, Volume II: Parallel Languages*, volume 259 of *Lecture Notes in Computer Science*, pages 141–158. Springer Berlin / Heidelberg, 1987. PARLE '87, Eindhoven, The Netherlands, June 15-19, 1987, Proceedings.

[CD97] Andrea Corradini and Frank Drewes. (Cyclic) term graph rewriting is adequate for rational parallel term rewriting. Technical Report TR-14-97, Universita di Pisa, Dipartimento di Informatica, 1997.

[CG95] Andrea Corradini and Fabio Gadducci. CPO models for infinite term rewriting. In Vangalur S. Alagar and Maurice Nivat, editors, *Algebraic Methodology and Software Technology*, volume 936 of *Lecture Notes in Computer Science*, pages 368–384. Springer Berlin / Heidelberg, 1995. 4th International Conference, AMAST '95 Montreal, Canada, July 3–7, 1995 Proceedings.

[Chu32] Alonzo Church. A set of postulates for the foundation of logic. *Annals of Mathematics*, 33:346–366, 1932.

[CMR+97] Andrea Corradini, Ugo Montanari, Francesca. Rossi, Hartmut Ehrig, Reiko Heckel, and Michael Löwe. Algebraic approaches to graph transformation, part I: Basic concepts and double pushout approach. In Grzegorz Rozenberg, editor, *Handbook of Graph Grammars and Computing by Graph Transformation, Volume 1: Foundations*, pages 163–245. University of Pisa, 1997.

[Cor93] Andrea Corradini. Term rewriting in CT_Σ. In Marie-Claude Gaudel and Jean-Pierre Jouannaud, editors, *TAPSOFT'93: Theory and Practice of Software Development*, volume 668 of *Lecture Notes in Computer Science*, pages 468–484. Springer Berlin / Heidelberg, 1993. Orsay, France, April 13-17, 1993 Proceedings.

[Cou83] Bruno Courcelle. Fundamental properties of infinite trees. *Theoretical Computer Science*, 25(2):95–169, 1983.

[CR36] Alonzo Church and J. Barkley Rosser. Some properties of conversion. *Transactions of the American Mathematical Society*, 39(3):472–482, May 1936.

[DK89] Nachum Dershowitz and Stéphane Kaplan. Rewrite, rewrite, rewrite, rewrite, rewrite... In *16th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 250–259, New York, NY, USA, 1989. ACM. POPL '89, Austin, Texas, January ,1989, Proceedings.

[DKP89] Nachum Dershowitz, Stéphane Kaplan, and David A. Plaisted. Infinite normal forms. In Giorgio Ausiello, Mariangiola Dezani-Ciancaglini, and Simona Ronchi Della Rocca, editors, *Automata, Languages and Programming, 16th International Colloquium,*, volume 372 of *Lecture Notes in Computer Science*, pages 249–262. Springer Berlin / Heidelberg, 1989. ICALP '89, Stresa, Italy, July 11-15, 1989, Proceedings.

[DKP91] Nachum Dershowitz, Stéphane Kaplan, and David A. Plaisted. Rewrite, rewrite, rewrite, rewrite, rewrite, ... *Theoretical Computer Science*, 83(1):71–96, 1991.

[DST80] Peter J. Downey, Ravi Sethi, and Robert Endre Tarjan. Variations on the common subexpression problem. *Journal of the ACM*, 27(4):758–771, 1980.

[Ech08] Rachid Echahed. On term-graph rewrite strategies. In Jürgen Giesl, editor, *7th International Workshop on Reduction Strategies in Rewriting and Programming*, volume 204, pages 99–110, 2008. WRS '07, Paris, France, June 25, 2007, Proceedings.

[EGH+07] Jörg Endrullis, Clemens Grabmayer, Dimitri Hendriks, Ariya Isihara, and Jan Willem Klop. Productivity of stream definitions. In Erzsébet Csuhaj-Varjú and Zoltán Ésik, editors, *Fundamentals of Computation Theory*, volume 4639 of *Lecture Notes in Computer Science*, pages 274–287. Springer Berlin / Heidelberg, 2007. FCT '07, Budapest, Hungary, August 27-30, 2007, Proceedings.

[EPS73] Hartmut Ehrig, Michael Pfender, and Hans Jürgen Schneider. Graph-grammars: An algebraic approach. In *14th Annual Symposium on Switching and Automata Theory*, pages 167–180, Washington, DC, USA, 1973. IEEE Computer Society. SWAT '73, The University of Iowa, USA, October 15-17, 1973, Proceedings.

[FRW90] William M. Farmer, John D. Ramsdell, and Ronald J. Watro. A correctness proof for combinator reduction with cycles. *ACM Transactions on Programming Languages and Systems*, 12(1):123–134, 1990.

[FW90] William M. Farmer and Ronald J. Watro. Redex capturing in term graph rewriting. *International Journal of Foundations of Computer Science*, 1:369–386, 1990.

[GL06] Bernhard Gramlich and Salvador Lucas. Generalizing newman's lemma for left-linear rewrite systems. In Frank Pfenning, editor, *Term Rewriting and Applications*, volume 4098 of *Lecture Notes in Computer Science*, pages 66–80. Springer Berlin / Heidelberg, 2006. RTA '06, Seattle, WA, USA, August 12-14, 2006, Proceedings.

[GTWW77] Joseph A. Goguen, James W. Thatcher, Eric G. Wagner, and Jesse B. Wright. Initial algebra semantics and continuous algebras. *Journal of the ACM*, 24(1):68–95, 1977.

[HL91a] Gérard P. Huet and Jean-Jacques Lévy. Computations in orthogonal rewriting systems, I. In Jean-Louis Lassez and Gordon Plotkin, editors, *Computational Logic - Essays in Honor of Alan Robinson*, chapter 11, pages 395–414. The MIT Press, 1991.

[HL91b] Gérard P. Huet and Jean-Jacques Lévy. Computations in orthogonal rewriting systems, II. In Jean-Louis Lassez and Gordon Plotkin, editors, *Computational Logic - Essays in Honor of Alan Robinson*, chapter 12, pages 415–443. The MIT Press, 1991.

[HP85] David Harel and Amir Pnueli. On the development of reactive systems. In Krzysztof R. Apt, editor, *Logics and models of concurrent systems*, pages 477–498. Springer-Verlag New York, Inc., New York, NY, USA, 1985.

[IN91] Paola Inverardi and Monica Nesi. Infinite normal forms for non-linear term rewriting systems. In Andrzej Tarlecki, editor, *Mathematical Foundations of Computer Science 1991*, volume 520 of *Lecture Notes in Computer Science*, pages 231–239. Springer Berlin / Heidelberg, 1991. MFCS'91, Kazimierz Dolny, Poland, September 9-13, 1991, Proceedings.

[IN95] Paola Inverardi and Monica Nesi. Deciding observational congruence of finite-state ccs expressions by rewriting. *Theoretical Computer Science*, 139(1-2):315–354, 1995.

[Jos89] Mark B. Josephs. The semantics of lazy functional languages. *Theoretical Computer Science*, 68(1):105–111, 1989.

[Kah09] Stefan Kahrs. Modularity of convergence in infinitary rewriting. In Ralf Treinen, editor, *Rewriting Techniques and Applications*, volume 5595 of *Lecture Notes in Computer Science*, pages 179–193. Springer Berlin / Heidelberg, 2009. RTA '09, Brasília, Brazil, June 29-July 1, 2009, Proceedings.

[KdV03] Richard Kennaway and Fer-Jan de Vries. Infinitary rewriting. In Terese [Ter03], chapter 12, pages 668–711.

[KdV05] Jan Willem Klop and Roel C. de Vrijer. Infinitary normalization. In Sergei N. Artëmov, Howard Barringer, Artur S. d'Avila Garcez, Luís C. Lamb, and John Woods, editors, *We Will Show Them! Essays in Honour of Dov Gabbay*, volume 2, pages 169–192. College Publications, 2005.

[Ken92] Richard Kennaway. On transfinite abstract reduction systems. Technical report, CWI (Centre for Mathematics and Computer Science), Amsterdam, 1992.

[Ket06] Jeroen Ketema. *Böhm-Like Trees for Rewriting*. PhD thesis, Vrije Universiteit Amsterdam, 2006.

[Ket08] Jeroen Ketema. On normalisation of infinitary combinatory reduction systems. In Andrei Voronkov, editor, *Rewriting Techniques and Applications*, volume 5117 of *Lecture Notes in Computer Science*, pages 172–186. Springer Berlin / Heidelberg, 2008. RTA '08, Hagenberg, Austria, July 15-17, 2008, Proceedings.

[Ket09] Jeroen Ketema. Comparing Böhm-like trees. In Ralf Treinen, editor, *Rewriting Techniques and Applications*, volume 5595 of *Lecture Notes in Computer Science*, pages 239–254. Springer Berlin / Heidelberg, 2009. RTA '09, Brasília, Brazil, June 29-July 1, 2009, Proceedings.

[KKSdV91] Richard Kennaway, Jan Willem Klop, M. Ronan Sleep, and Fer-Jan de Vries. Transfinite reductions in orthogonal term rewriting systems. In Ronald V. Book, editor, *Rewriting Techniques and Applications*, volume 488 of *Lecture Notes in Computer Science*, pages 1–12. Springer Berlin / Heidelberg, 1991. RTA '91, Como, Italy, April 10-12, 1991, Proceedings.

[KKSdV94] Richard Kennaway, Jan Willem Klop, M. Ronan Sleep, and Fer-Jan de Vries. On the adequacy of graph rewriting for simulating term rewriting. *ACM Transactions on Programming Languages and Systems*, 16(3):493–523, 1994.

[KKSdV95a] Richard Kennaway, Jan Willem Klop, M. Ronan Sleep, and Fer-Jan de Vries. Transfinite reductions in orthogonal term rewriting systems. *Information and Computation*, 119(1):18–38, 1995.

[KKSdV95b] Richard Kennaway, Jan Willem Klop, Ronan Sleep, and Fer-Jan de Vries. Infinitary lambda calculi and böhm models. In Jieh Hsiang, editor, *Rewriting Techniques and Applications*, volume 914 of *Lecture Notes in Computer Science*, pages 257–270. Springer Berlin / Heidelberg, 1995. RTA '95 Kaiserslautern, Germany, April 5-7, 1995 Proceedings.

[KKSdV97] Richard Kennaway, Jan Willem Klop, M R. Sleep, and Fer-Jan de Vries. Infinitary lambda calculus. *Theoretical Computer Science*, 175(1):93–125, 1997.

[KP93] Gilles Kahn and Gordon D. Plotkin. Concrete domains. *Theoretical Computer Science*, 121(1-2):187–277, 1993.

[KS05a] Jeroen Ketema and Jakob Grue Simonsen. Infinitary combinatory reduction systems. In Jürgen Giesl, editor, *Term Rewriting and Applications*, volume 3467 of *Lecture Notes in Computer Science*, pages 438–452. Springer Berlin / Heidelberg, 2005. RTA '05, Nara, Japan, April 19-21, 2005, Proceedings.

[KS05b] Jeroen Ketema and Jakob Grue Simonsen. On confluence of infinitary combinatory reduction systems. In Geoff Sutcliffe and Andrei Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning*, volume 3835 of *Lecture Notes in Computer Science*, pages 199–214. Springer Berlin / Heidelberg, 2005. LPAR '05, Montego Bay, Jamaica, December 2-6, 2005, Proceedings.

[KS06] Jeroen Ketema and Jakob Grue Simonsen. Infinitary combinatory reduction systems. Technical Report D-558, University of Copenhagen, 2006.

[KvOdV99] Richard Kennaway, Vincent van Oostrom, and Fer-Jan de Vries. Meaningless terms in rewriting. *Journal of Functional and Logic Programming*, 1999(1):1–35, February 1999.

[Lau93] John Launchbury. A natural semantics for lazy evaluation. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 144–154, New York, NY, USA, 1993. ACM. POPL '93, Charleston, South Carolina, USA, January 1993, Proceedings.

[Lév76] Jean-Jacques Lévy. An algebraic interpretation of the $\lambda\beta k$-calculus; and an application of a labelled λ-calculus. *Theoretical Computer Science*, 2(1):97–114, 1976.

[Lév78] Jean-Jacques Lévy. *Reductions Correctes et Optimales dans le Lambda-Calcul.* PhD thesis, Université Paris, January 1978.

[Lev79] Azriel Levy. *Basic Set Theory.* Perspective in Mathematical Logic. Springer-Verlag, 2nd edition, 1979.

[Lon83] Giuseppe Longo. Set-theoretical models of λ-calculus: theories, expansions, isomorphisms. *Annals of pure and applied logic*, 24(2):153–188, 1983.

[Luc01] Salvador Lucas. Transfinite rewriting semantics for term rewriting systems. In Aart Middeldorp, editor, *Rewriting Techniques and Applications*, volume 2051 of *Lecture Notes in Computer Science*, pages 216–230. Springer Berlin / Heidelberg, 2001. RTA '01, Utrecht, The Netherlands, May 22-24, 2001, Proceedings.

[Mid97] Aart Middeldorp. Call by need computations to root-stable form. In *ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 94–105, New York, NY, USA, 1997. ACM. POPL '97, Paris, France, 15-17 January, 1997, Proceedings.

[Mun00] James Munkres. *Topology.* Prentice Hall, 2nd edition, 2000.

[Nak75] Reiji Nakajima. Infinite normal forms for the lambda-calculus. In Corrado Böhm, editor, *Lambda-Calculus and Computer Science Theory*, volume 37 of *Lecture Notes in Computer Science*, pages 62–82. Springer Berlin / Heidelberg, 1975. Symposium on Lambda-Calculus and Computer Science Theory, Rome, Italy, March 25-27, 1975, Proceedings.

[PJ87] Simon Peyton-Jones. *The Implementation of Functional Programming Languages.* Prentice Hall, 1987.

[Pla95] Marinus J. Plasmeijer. Clean: a programming environment based on term graph rewriting. In Andrea Corradini and Ugo Montanari, editors, *Joint COMPUGRAPH/SEMAGRAPH Workshop on Graph Rewriting and Computation*, volume 2, pages 215–221, 1995. SEGRAGRA 95, Volterra, Italy, August 28 - September 1, 1995, Proceedings.

[Plu99] Detlef Plump. Term graph rewriting. In Hartmut Ehrig, Gregor Engels, Hans-Jörg Kreowski, and Grzegorz Rozenberg, editors, *Handbook of Graph Grammars and Computing by Graph Transformation, Volume 2: Applications, Languages, and Tools*, pages 3–61. World Scientific Publishing Co., Inc., River Edge, NJ, USA, 1999.

[PR69] John L. Pfaltz and Azriel Rosenfeld. Web grammars. In Donald E. Walker and Lewis M. Norton, editors, *1st International Joint Conference on Artificial Intelligence*, pages 609–620, 1969. IJCAI '69, Washington, DC, USA, May 1969, Proceedings.

[Pra71] Terrence W. Pratt. Pair grammars, graph languages and string-to-graph translations. *Journal of Computer and System Sciences*, 5(6):560—-595, 1971.

[Rao84] Jean Claude Raoult. On graph rewritings. *Theoretical Computer Science*, 32(1-2):1–24, 1984.

[Rod98] Pieter Hendrik Rodenburg. Termination and confluence in infinitary term rewriting. *The Journal of Symbolic Logic*, 63(4):1286–1296, 1998.

[Sie65] Wacław Sierpiński. *Cardinal and Ordinal Numbers.* Państwowe Wydawnictwo Naukowe, 2nd edition, 1965.

[Sij89] Ben A. Sijtsma. On the productivity of recursive list definitions. *ACM Transactions on Programming Languages and Systems*, 11(4):633–649, 1989.

[Sim04] Jakob Grue Simonsen. On confluence and residuals in Cauchy convergent transfinite rewriting. *Information Processing Letters*, 91(3):141–146, 2004.

[Sim06] Jakob Grue Simonsen. On modularity in infinitary term rewriting. *Information and Computation*, 204(6):957–988, 2006.

[Sta80a] John Staples. Computation on graph-like expressions. *Theoretical Computer Science*, 10(2):171–185, 1980.

[Sta80b] John Staples. Optimal evaluations of graph-like expressions. *Theoretical Computer Science*, 10(3):297–316, 1980.

[Sta80c] John Staples. Speeding up subtree replacement systems. *Theoretical Computer Science*, 11(1):39–47, 1980.

[Ter03] Terese. *Term Rewriting Systems.* Cambridge University Press, 1st edition, 2003.

[Tur79] D.A. Turner. A new implementation technique for applicative languages. *Software: Practice and Experience*, 9(1):31–49, 1979.

[vESP97] Marko C. J. D. van Eekelen, Sjaak Smetsers, and Rinus Plasmeijer. Graph rewriting semantics for functional programming languages. In Dirk van Dalen and Marc Bezem, editors, *Computer Science Logic*, volume 1258 of *Lecture*

Notes in Computer Science, pages 106–128. Springer Berlin / Heidelberg, 1997. CSL '96 Annual Conference of the EACSL Utrecht, The Netherlands, September 21–27, 1996, Proceedings.

[Wad71] Christopher P. Wadsworth. *Semantics and pragmatics of the lambda calculus*. Ph.D. thesis, University of Oxford, 1971.

[Zan09] Hans Zantema. Well-definedness of streams by termination. In Ralf Treinen, editor, *Rewriting Techniques and Applications*, volume 5595 of *Lecture Notes in Computer Science*, pages 164–178. Springer Berlin / Heidelberg, 2009. RTA '09, Brasília, Brazil, June 29-July 1, 2009, Proceedings.

Index

WN^{∞}, *see* infinitary normalisation

abstract reduction system, 18
acyclic
  graph, 51
  path, 51
  predecessor, 75
  sharing, 64
  term graph, 52
  term graph rewrite rule, 150
  term graph rewriting system, 150
  term rewriting system, 88
acyclicity, 88
algebraic term graph rewriting, 139, 159
almost non-collapsing
  term graph rewriting system, 155
  term rewriting system, 106
almost orthogonal term rewriting system, 25
α-reduction sequence, 19
α-sequence, 11
ancestor, 100, 120
applicable
  term graph rewrite rule, 143
  term rewrite rule, 24
argument of a redex, 24
ARS, *see* abstract reduction system

Böhm reduction, 109
Böhm reduction sequence, 109
Böhm tree, 110
basis element, 14
basis for a topology, 14
bcpo, *see* bounded complete partial order
$\perp,\mathcal{U}$-instance, 109
bounded complete partial order, 9
build step, 142

canonical term graph, 57
capture a redex, 151
Cauchy sequence, 16
closed
  reduction sequence, *see* closed sequence
  sequence, 11
  set, 14
closure
  reflexive symmetric transitive, 19
  reflexive transitive, 19
  symmetric, 19
  symmetric transitive, 19
  transitive, 19
coincide, 22
collapsing
  position, 25
  redex, 25
  rule, 25
  tower, 25
    infinite, 101
    maximal, 25
combinatory reduction system, 157
COMP, *see* completeness
compatible elements, 8
complete
  development, 100, 123
  metric reduction system, 28
  partial order, 9
  partial reduction system, 36
  semilattice, 9
completeness, 20
composition
  of functions, 7
  of partial functions, 7
concatenation, 11
conflicting redex occurrences, 25
confluence, 19
  ground, 89
  infinitary, 42
constructor, 24
context, 22
context function, 36
continuous
  at point, 14
  function, 14
  sequence, 14
convergent sequence, 15
conversion sequence
  finite, 19

CR, *see* confluence
CR^∞_{hc}, *see* infinitary confluence modulo hypercollapsing terms
CR^∞, *see* infinitary confluence
critical pair, 24
    joinable, 24
    trivial, 24
cyclic
    graph, 51
    path, 51
    sharing, 64
    term graph, 52
    term graph rewrite rule, 150
    term graph rewriting system, 150
cyclisation, 151

defined symbol, 24
deletion of a sub-term graph, 146
Δ-homomorphism, 54
Δ-isomorphic, 54
Δ-isomorphism, 54
Δ-node, 50
dense set, 18
depth
    of a node, 52
    of a sub-term graph, 146
depth of Δ, 52
depth-increasing reduction, 113
descendant, 98, 117
destructive reduction sequence, 133
development, 100, 123
    complete, 100, 123
    total, 124
diamond property, 20
directed set, 8
disjoint
    positions, 21
    sequences, 11
domain
    of a partial function, 7
    of a substitution, 22
double-pushout approach, 159
DP, *see* diamond property

edge of a path, 125
equational term graph rewriting, 139, 159
equivalence relation, 8
extension
    of a metric reduction system, 46
    of a metric space, 45
    of a sequence, 11

fair reduction sequence, 112
finite
    conversion sequence, 19
    graph, 50
    ground term, 20
    jumps, 127
    sequence, 11
    substitution, 22
    term, 20
    term graph, 51
fragile term, 133
fringe node, 75
function, 7
    composition, 7
    context, 36
    continuous, 14
    continuous at point, 14
    domain, 7
    height, 28
    identity, 7
    inverse, 7
    labelling, 50
    monotone, 8
    range, 7
    source, 18
    successor, 50
    target, 18

garbage collection step, 142
GCR, *see* ground confluence
glb, *see* greatest lower bound
graph, 50
    acyclic, 51
    cyclic, 51
    finite, 50
greatest element, 8
greatest lower bound, 8
ground confluence, 89
ground term, 20
    finite, 20
GRS, *see* term graph rewriting system

Hausdorff space, 14
hc-equivalent, 106
height function, 28
homomorphic, 54
horizontal sharing, 64, 149
hypercollapsing term, 106

identity function, 7
IGRS, *see* infinitary term graph rewriting system
induced abstract reduction system
    of a term graph rewriting system, 144
    of a term rewriting system, 24
induced metric reduction system

  of a term graph rewriting system, 147
  of a term rewriting system, 28
induced partial reduction system
  of a term graph rewriting system, 147
  of a term rewriting system, 37
infinitarily normalising
  reduction sequence, 110
  reduction strategy, 111
infinitarily terminating, 44
infinitary confluence, 42
infinitary confluence modulo hypercollapsing terms, 106
infinitary ground term, 20
infinitary normalisation, 42
infinitary term, 20
infinitary term graph rewrite rule, 140
infinitary term graph rewriting system, 140, *see* term graph rewriting system
infinitary term rewrite rule, 23
infinitary term rewriting system, 23
infinitary unique normal form property w.r.t. reduction, 42
infinite collapsing tower, 101
infinite sequence, 11
innermost
  redex, 88
  reduction sequence, 88
innermost normalisation, 89
innermost termination, 89
inverse
  of a partial function, 7
inverse of a function, 7
irreducible, 1
isomorphic, 8
ITRS, *see* infinitary term rewriting system

joinable critical pair, 24

labelled
  signature, 122
  symbol, 122
  term, 122
  term rewriting system, 122
labelling function, 50
λ-calculus, 157
lazy evaluation, 2
lazy functional programming language, 2
least element, 8
least upper bound, 8
left root of a term graph rewrite rule, 140
left-hand side
  of a term graph rewrite rule, 140
  of a term rewrite rule, 23
left-linear
  term graph rewrite rule, 152
  term graph rewriting system, 152
  term rewrite rule, 25
  term rewriting system, 25
left-narrow partial order, 8
length of a sequence, 11
limit, 15
limit inferior, 13
limit ordinal, 10
linear, 25
lower bound, 8
lub, *see* least upper bound

matching an occurrence representation, 62
matching step, 141
matching term, 129
matching $\mathcal{V}$-homomorphism, 143
maximal collapsing tower, 25
meaningless terms, 108
metric, 15
metric completion, 18
metric reduction system, 28
  complete, 28
  extension, 46
  weak extension, 46
metric space, 15
  extension, 45
metric topology, 16
monotone
  function, 8
  order, 23
MRS, *see* metric reduction system
MRS reduction, 90

needed
  redex occurrence, 111
  reduction sequence, 111
needed reduction, 111
needed-fair reduction sequence, 112
neighbourhood, 14
NF, *see* normal form property
node
  fringe, 75
  of a graph, 50
  of a path, 125
  truncation, 75
non-conflicting redex occurrences, 25
non-Δ-node, 50
non-looping term rewriting system, 88
non-loopingness, 88
non-σ-node, 50
normal form, 1, 19
normal form property, 20
normalisation, 20

innermost, 89

occurrence
in a term, 21
in a term graph, 52
occurrence representation, 61
occurrences
rational, 154
ω-normalising
reduction sequence, 111
reduction strategy, 111
ω-termination, 88
open
reduction sequence, *see* open sequence
sequence, 11
set, 14
open ball, 16
operational term graph rewriting, 139
order, *see* partial order
order isomorphic, 8
order isomorphism, 8
ordinal, 9
ordinal space, 14
orthogonal
term graph rewriting system, 152
term rewriting system, 25
outermost-volatile position, 116
overlap, 24, 108
overlay, 24
overlay system, 25

P-reduction, 111
parallel-outermost reduction sequence, 113
partial function, 7
composition, 7
inverse, 7
range, 7
partial order, 8
bounded complete, 9
complete, 9
left-narrow, 8
well-founded, 8
partial reduction system, 36
complete, 36
partial term, 23
partial term graph, 63
partially ordered class, 8
partially ordered set, 8
path, 51, 52, 125
acyclic, 51
cyclic, 51
edge, 125
node, 125
trace, 127

pattern of a term rewrite rule, 24
polynomial termination, 88
position
disjoint, 21
in a term, 21
of a trace, 129
outermost-volatile, 116
prefix, 21
volatile, 116
predecessor
acyclic, 75
prefix
of a position, 21
of a sequence, 11
preservation of sharing, 64
prestep
of a graph rewriting system, 144
of a term rewriting system, 24
proper subsequence, 11
proper subterm, 21
PRS, *see* partial reduction system
PRS reduction, 90

quotient, 8

R-fair reduction sequence, 112
range
of a function, 7
of a partial function, 7
rational occurrences, 154
rational reduction sequence, 154
rational term, 149
reachable node, 51
redex
of a term graph rewriting system, 143
of a term rewriting system, 24
redex capturing, 150
redex occurrence, 24
redex pattern, 24
redirection step, 142
reduce, 19
reduct, 19
reduction, *see* reduction sequence
reduction sequence, 19
ω-normalising, 111
closed, *see* closed sequence
destructive, 133
infinitarily normalising, 110
open, *see* open sequence
strongly α-convergent, 30, 38
strongly continuous, 29, 38
strongly convergent, 29, 38
strongly divergent, 30, 38
total, 40

weakly α-convergent, 30, 38
weakly continuous, 29, 38
weakly convergent, 29, 38
weakly divergent, 30, 38
reduction strategy, 110
ω-normalising, 111
infinitarily normalising, 111
strongly convergent, 111
reflexive symmetric transitive closure, 19
reflexive transitive closure, 19
regular equations, 159
replacement
of a term, 22
of a term graph, 53
representation by a term graph, 149
residual, 98, 118
result term, 24
result term graph, 143
reverse of a relation, 19
rewrite rule, 23
ρ-redex
of a term graph rewriting system, 143
of a term rewriting system, 24
ρ-redex occurrence, 24
right-hand side
of a term graph rewrite rule, 140
of a term rewrite rule, 23
right-root of a term graph rewrite rule, 140
root node of a term graph, 51
root-active term, 107

segment of a sequence, 11
semilattice, 9
sequence, 11
Cauchy, 16
closed, 11
concatenation, 11
continuous, 14
convergent, 15
converges to, 15
disjoint, 11
finite, 11
infinite, 11
length, 11
open, 11
prefix, 11
segment, 11
suffix, 11
transfinite, 11
set
closed, 14
dense, 18
directed, 8
open, 14
set of meaningless terms, 108
sharing
acyclic, 64
cyclic, 64
horizontal, 64, 149
of subexpressions, 49
vertical, 64, 149
Σ-graph, 50
σ-node, 50
signature, 20
labelled, 122
similarity
of term graphs, 75
of terms, 22
simple termination, 88
SIN, *see* innermost termination
SN, *see* termination
SN^∞, *see* infinitarily terminating
source function, 18
step
of a term graph rewriting system, 144
of a term rewriting system, 24
stream definition, 157
strong Δ-homomorphism, 65
strong extension, 46
strongly α-convergent reduction sequence, 30, 38
strongly continuous reduction sequence, 29, 38
strongly convergent reduction sequence, 29, 38
strongly convergent reduction strategy, 111
strongly divergent reduction sequence, 30, 38
sub-term graph, 52
subgraph, 52
subsequence, 11
proper, 11
substitution, 22
domain, 22
finite, 22
subterm, 21
at position, 22
successor function, 50
successor ordinal, 10
suffix of a sequence, 11
symbol at position, 22
symbol of a trace, 129
symmetric closure, 19
symmetric transitive closure, 19
syntactic accident, 29

target function, 18

TARS, *see* transfinite abstract reduction system
term, 20
  finite, 20
  labelled, 122
  partial, 23
  rational, 149
  total, 23
term graph, 51
  acyclic, 52
  canonical, 57
  cyclic, 52
  finite, 51
  total, 63
term graph rewrite rule, 140
  acyclic, 150
  applicable, 143
  cyclic, 150
  left-hand side, 140
  right-hand side, 140
  unravelling, 149
term graph rewriting
  algebraic, 139, 159
  equational, 139, 159
  operational, 139
term graph rewriting system, 140
  acyclic, 150
  almost non-collapsing, 155
  cyclic, 150
  unravelling, 149
term rewrite rule, 23
  applicable, 24
  left-hand side, 23
  right-hand side, 23
term rewriting system, 23
  labelled, 122
term tree, 52
termination, 20
  ω-, 88
  innermost, 89
  polynomial, 88
  simple, 88
  top-, 25
tiling diagram, 102
top-termination, 25
topological space, 14
topology, 14
total
  development, 124
  reduction sequence, 40
  term, 23
  term graph, 63
trace of a path, 127
transfinite abstract reduction system, 42
transfinite induction, 10
transfinite recursion, 10
transfinite sequence, 11
transitive class, 9
transitive closure, 19
trivial critical pair, 24
TRS, *see* term rewriting system
truncation node, 75
truncation of a term graph, 75

$\mathcal{U}$-equivalent, 108
ultrametric, 15
ultrametric reduction system, 28
ultrametric space, 15
UN, *see* unique normal form property
underlying abstract reduction system
  of a metric reduction system, 28
  of a partial reduction system, 36
underlying graph, 51
unique normal form property, 20
unique normal form property w.r.t. reduction, 20
unlabelled
  symbol, 122
$\mathsf{UN}_{\to}$, *see* unique normal form property w.r.t. reduction
unravelling
  of a term graph, 60
  of a term graph rewrite rule, 149
  of a term graph rewriting system, 149
$\mathsf{UN}^{\infty}_{\to}$, *see* infinitary unique normal form property w.r.t. reduction
upper bound, 8
URS, *see* ultrametric reduction system

vertical sharing, 64, 149
volatile position, 116

weak extension of a metric reduction system, 46
weakly α-convergent reduction sequence, 30, 38
weakly continuous reduction sequence, 29, 38
weakly convergent reduction sequence, 29, 38
weakly divergent reduction sequence, 30, 38
weakly orthogonal term rewriting system, 25
well-founded partial order, 8
well-order, 8
well-ordered class, 8
well-ordered set, 8

WIN, *see* innermost normalisation
WN, *see* normalisation

Symbols

$\mathbb{N}^+$ 7
$\mathbb{N}$ 7
$\mathbb{R}$ 7
$\mathbb{R}^+$ 7
$\mathbb{R}_0^+$ 7
$\circ$ 7
$\mathsf{range}(f)$ 7
id_A 7
f^{-1} 7
$\mathsf{dom}(f)$ 7
$f(A)$ 7
$f^{-1}(B')$ 7
R^n 7
R^* 7
R^+ 7
$\bigsqcup$ 8
$\bigsqcap$ 8
$\sqcup$ 9
$\bigsqcup_{P(i)} x_i$ 9
$\sqcap$ 9
$\bigsqcap_{P(i)} x_i$ 9
On 9
$\mathsf{S}(\alpha)$ 10
ω 10
$(a_\iota)_{\iota<\alpha}$ 11
M^α 11
$M^{\leq\alpha}$ 11
$M^{<\alpha}$ 11
M^* 11
ε 11
$|S|$ 11
$\cdot$ 11
S/f 11
$S|_{[\beta,\gamma)}$ 11
$\prod_{\iota<\alpha} S_\iota$ 12
$\liminf_{\iota\to\alpha} a_\iota$ 13
$\mathcal{T}$ 14
$\mathcal{B}$ 14
$\lim_{\iota\to\alpha} x_\iota$ 15
$\mathbf{d}$ 15
$B_{\mathbf{d}}(m,\varepsilon)$ 16
$\mathcal{A}$ 18
src 18
tgt 18
src 18
tgt 18
$S\colon a \to^*_{\mathcal{A}} b$ 19
$S\colon a \to_{\mathcal{A}} b$ 19
$S\colon a \to^+_{\mathcal{A}} b$ 19
$S\colon b \leftarrow^*_{\mathcal{A}} a$ 19
$S\colon b \leftarrow_{\mathcal{A}} a$ 19
$S\colon b \leftarrow^+_{\mathcal{A}} a$ 19
$S\colon a \leftrightarrow^*_{\mathcal{A}} b$ 19
$S\colon a \leftrightarrow_{\mathcal{A}} b$ 19
$S\colon a \leftrightarrow^+_{\mathcal{A}} b$ 19
$S\colon a \to^*_{\mathcal{A}^s} b$ 19
$S\colon a \to_{\mathcal{A}^s} b$ 19
$S\colon a \to^+_{\mathcal{A}^s} b$ 19
CR 19
DP 20
SN 20
WN 20
COMP 20
NF 20
UN 20
$\mathsf{UN}_{\to}$ 20
Σ 20
$\mathsf{ar}(f)$ 20
$\mathcal{T}(\Sigma,\mathcal{V})$ 20
$\mathcal{V}$ 20
$\mathcal{T}(\Sigma)$ 20
$\mathcal{T}^\infty(\Sigma,\mathcal{V})$ 20
$\mathcal{T}^\infty(\Sigma)$ 20
$\mathcal{P}(t)$ 21
$\mathcal{S}(t)$ 21
$t|_\pi$ 22
$t(\pi)$ 22
$t[s]_\pi$ 22
$\mathcal{P}_\Sigma(t)$ 22
$C[,\dots,]$ 22
$C\langle,\dots,\rangle$ 22
$C\{,\dots,\}$ 22
$C[]$ 22
$\mathsf{dom}(\sigma)$ 22
$\mathsf{sim}(t,t')$ 22
$\leq_\perp$ 23

$\Sigma_{\perp}$ 23
$\mathcal{P}_{\setminus\perp}(t)$ 23
$\mathcal{P}_{\perp}(t)$ 23
$\rho\colon l \to r$ 23
$\mathcal{A}_{\mathcal{R}}$ 24
$s \to_{\pi,\rho} t$ 24
$\mathcal{M}$ 28
hgt 28
$\mathcal{M}_{\mathcal{R}}$ 28
$S\colon a \hookrightarrow^{\alpha}_{\mathcal{A}}$ 32
$S\colon a \twoheadrightarrow^{\alpha}_{\mathcal{A}}$ 32
$S\colon a \hookrightarrow^{\alpha}_{\mathcal{A}} b$ 32
$S\colon a \twoheadrightarrow^{\alpha}_{\mathcal{A}} b$ 32
$S\colon a \nearrow^{\alpha}_{\mathcal{A}}$ 32
$S\colon a \nearrow\!\!\!\nearrow^{\alpha}_{\mathcal{A}}$ 32
$\mathcal{P}$ 36
cxt 36
$\mathcal{P}_{\mathcal{R}}$ 37
CR^{∞} 42
WN^{∞} 42
$\mathsf{UN}^{\infty}_{\to}$ 42
$a \leftrightarrow^{w}_{\mathcal{T}} b$ 43
$a \leftrightarrow^{s}_{\mathcal{T}} b$ 43
NF^{∞} 43
UN^{∞} 43
$\mathsf{UN}^{\infty}_{\to}$ 43
WN^{∞} 43
SN^{∞} 44
SN^{∞} 44
lab 50
suc 50
suc_i 50
$\mathsf{ar}_g(n)$ 50
$\mathcal{G}^{\infty}(\Sigma)$ 51
$\mathcal{G}(\Sigma)$ 51
$\mathcal{P}(g)$ 52
$\mathcal{P}_g(n)$ 52
$\mathsf{depth}_g(n)$ 52
$\mathsf{depth}(g)$ 52
Δ-$\mathsf{depth}(g)$ 52
σ-$\mathsf{depth}(g)$ 52
$g|n$ 52
$\mathsf{node}_g(\pi)$ 52
$\mathsf{ar}_g(\pi)$ 52
$g|\pi$ 52
$\varphi\colon g \to_{\Delta} h$ 54
$\varphi\colon g \xrightarrow{\sim}_{\Delta} h$ 54
$\cong_{\Delta}$ 54
$\varphi\colon g \to_{\sigma} h$ 54
$\varphi\colon g \to h$ 54
$\varphi\colon g \xrightarrow{\sim}_{\sigma} h$ 54
$\varphi\colon g \xrightarrow{\sim} h$ 54
$\cong_{\sigma}$ 54
$\cong$ 54
$\mathcal{G}^{\infty}_{\mathcal{C}}(\Sigma)$ 57
$\mathcal{G}_{\mathcal{C}}(\Sigma)$ 57
$\sim_g$ 58
$\mathcal{U}(g)$ 60
$\mathcal{P}^a(g)$ 64
$\mathcal{P}^a_g(n)$ 64
$\leq_{\perp}$ 66
$\mathsf{sim}(g,h)$ 75
$\mathsf{Pre}^a_g(n)$ 75
$N^g_{<d}$ 75
$N^g_{=d}$ 75
$g|d$ 75
SIN 89
WIN 89
$\equiv_{hc}$ 106
$\mathsf{CR}^{\infty}_{hc}$ 106
$\equiv_{\mathcal{U}}$ 108
$\mathcal{U}_{\perp}$ 109
$\mathcal{B}_{\mathcal{R},\mathcal{U}}$ 109
$\to_{\mathcal{U},\perp}$ 109
$\mathcal{B}$ 109
$\to_{\perp}$ 109
$\mathcal{S}(t)$ 110
$U /\!/ S$ 117
Σ^l 122
$\|t\|$ 122
$\mathcal{R}^l$ 122
$\|\rho\|$ 122
$t^{(U)}$ 122
$\mathcal{P}(t,U,\mathcal{R})$ 127
$\mathsf{pos}(T)$ 129
$\mathsf{sym}(T)$ 129
$\mathcal{F}(t,U,\mathcal{R})$ 129
$\mathcal{A}_{\mathcal{R}}$ 144
$g \to_{n,\rho} h$ 144
$g|n$-$\mathsf{depth}(g)$ 146
$g \setminus g|n$ 146
$\mathcal{M}_{\mathcal{R}}$ 147
$\mathcal{P}_{\mathcal{R}}$ 147
$\mathcal{U}(\rho)$ 149
$\mathcal{U}(\mathcal{R})$ 149
$S\colon s \twoheadrightarrow^{r} t$ 154