

TECHNISCHE UNIVERSITÄT MÜNCHEN
INSTITUT FÜR INFORMATIK

On Normalization and Type Checking for Tree Transducers

Dipl.-Inform. Sylvia Friese

Vollständiger Abdruck der von der Fakultät für Informatik der
Technischen Universität München zur Erlangung des akademischen
Grades eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitzender: Univ.-Prof. Alfons Kemper, Ph.D.
Prüfer der Dissertation: 1. Univ.-Prof. Dr. Helmut Seidl
2. Univ.-Prof. Dr. Klaus U. Schulz,
Ludwig-Maximilians-Universität München

Die Dissertation wurde am 24.03.2011 bei der Technischen
Universität München eingereicht und durch die Fakultät für
Informatik am 07.07.2011 angenommen.

Abstract

Tree transducers are an expressive formalism for reasoning about tree structured data. Practical applications range from XSLT-like document transformations to translations of natural languages. Important problems for transducers are to decide whether two transducers are equivalent, to construct normal forms, give semantic characterizations, and type checking, i.e., to check whether the produced outputs satisfy given structural constraints. This thesis addresses these problems for important classes of tree transducers. Constructive solutions are provided and classes of transducers for which these algorithms run in polynomial time, are identified.

Equivalence testing, normalization, and semantic characterization are often solved together by the use of a Myhill-Nerode theorem. This identifies necessary and sufficient semantic properties for transformations definable by a specific class of transducers. The theorem also implies that a unique normal form of those transducers exists. Moreover, it implies that, given a transducer, the normal transducer can be constructed. This immediately leads to the question: Are there classes of tree transducers for which a Myhill-Nerode theorem exists? We give an affirmative answer for the class of deterministic bottom-up tree transducers. A semantic characterization of transformations definable by these transducers is presented, and, moreover, it is evidenced that for every deterministic bottom-up tree transducer, a unique equivalent transducer can be constructed, which is minimal. The construction is based on a sequence of normalizing transformations, which, among others, guarantee that non-trivial output is produced as early as possible. For a deterministic bottom-up transducer where every state produces either none or infinitely many outputs, the minimal transducer can be constructed in polynomial time.

One of the useful properties of tree walking transducers is decidability of type checking: Given a transducer and input and output types, it can be checked statically whether each document adhering to the input type is necessarily transformed by the transducer into documents adhering to the output type. Here, a “type” means a regular set of trees specified by a finite-state tree automaton. Usually, type checking of tree transducers is extremely expensive; already for simple top-down tree transducers it is known to be EXPTIME-complete. Are there expressive classes of tree transducers for which type checking can be performed in polynomial time? Most of the previous approaches are based on inverse type inference. In contrast, the approach presented here uses forward

type inference. This means to infer, given a transducer and an input type, the corresponding set of output trees. In general, this set is not a type, i.e., it is not regular. However, it can be decided if its intersection with a given type is empty. Using this approach, we show that type checking can be performed in polynomial time if (1) the output type is specified by a deterministic tree automaton, and (2) the tree walking transducer visits every input node only a bounded number of times. If the transducer is additionally equipped with accumulating call-by-value parameters, then the complexity of type checking also depends (exponentially) on the number of such parameters. For this case, a fast approximative type checking algorithm is presented, based on context-free tree grammars. Finally, the approach is generalized from trees to forest walking transducers, which additionally support concatenation as a built-in output operation.

Zusammenfassung

Baumübersetzer sind ein ausdrucksstarker Formalismus, um baumstrukturierte Daten zu analysieren. Praktische Anwendungen reichen von XSLT-artigen Dokumentumstrukturierungen zu Übersetzungen natürlicher Sprache. Bedeutende Problemstellungen für Übersetzer sind, zu entscheiden, ob zwei Übersetzer äquivalent sind, eine Normalform konstruiert werden kann, es eine semantische Charakterisierung gibt und Typüberprüfung, d.h., zu überprüfen, ob die erzeugten Ausgaben gegebene strukturelle Bedingungen erfüllen. Diese Dissertation befasst sich mit diesen Fragen für wesentliche Baumübersetzerklassen. Wir stellen konstruktive Lösungen bereit und identifizieren Übersetzerklassen, für die diese Algorithmen nur polynomielle Zeit benötigen.

Äquivalenztest, Normalisierung und semantische Charakterisierung sind oft mittels eines Myhill-Nerode Theorems gemeinsam lösbar. Es zeigt notwendige und hinreichende semantische Eigenschaften für Übersetzungen einer bestimmten Klasse von Übersetzern auf und impliziert eine eindeutige Normalform für diese Übersetzer. Zusätzlich beinhaltet das Theorem, dass der entsprechende Übersetzer in Normalform aus einem beliebigen Übersetzer konstruiert werden kann. Es stellt sich also die Frage, ob es Klassen von Baumübersetzern gibt, für die ein Myhill-Nerode Theorem existiert. Für deterministische Aufwärts-Baumübersetzer (deterministic bottom-up tree transducers) geben wir eine positive Antwort. Wir präsentieren eine semantische Charakterisierung für Übersetzungen, die durch solche Übersetzer beschrieben werden können. Desweiteren zeigen wir, dass für jeden deterministischen Aufwärts-Baumübersetzer ein eindeutiger äquivalenter Übersetzer konstruiert werden kann, der minimal ist. Diese Konstruktion basiert auf einer Folge von Normalisierungen, welche unter anderem garantieren, dass nicht-triviale Ausgaben so früh wie möglich erzeugt werden. Wenn jeder Zustand eines deterministischen Aufwärts-Baumübersetzers entweder keine oder unendlich viele Ausgaben produziert, kann der minimale Übersetzer in polynomieller Zeit konstruiert werden.

Eine der nützlichen Eigenschaften von Zwei-Wege-Baumübersetzern (tree walking transducers) ist, dass das Problem der Typüberprüfung entscheidbar ist: Wenn ein Übersetzer und ein Eingabe- sowie ein Ausgabetypp vorgegeben sind, kann statisch geprüft werden, ob der Übersetzer Dokumente, die dem Eingabetyp entsprechen, grundsätzlich in Dokumente des Ausgabetyps übersetzt. Dabei verstehen wir unter "Typ" eine reguläre Baummenge, die durch einen endlichen Automaten spezifiziert ist. Normalerweise ist Typüberprüfung von

Baumübersetzern extrem aufwändig. Es ist bekannt, dass dieses Problem bereits für einfache Abwärts-Baumübersetzer (top-down tree transducers) EXPTIME-vollständig ist. Es stellt sich die Frage, ob es ausdrucksstarke Baumübersetzerklassen gibt, für die Typüberprüfung in polynomieller Zeit durchgeführt werden kann. Die meisten bisherigen Ansätze basierten auf inverser Typinferenz. Wir benutzen hier den entgegengesetzten Ansatz mittels vorwärtsgerichteter Typinferenz. Das heißt, für einen gegebenen Übersetzer und einen Eingabetyp leiten wir die Menge aller zugehörigen Ausgabebäume her. Im Allgemeinen ist diese Menge kein Typ, d.h., es ist keine reguläre Menge. Trotzdem ist es entscheidbar, ob die Schnittmenge mit einem gegebenen Ausgabebetyp leer ist. Mit diesem Ansatz zeigen wir, dass man Typüberprüfung in polynomieller Zeit durchführen kann, wenn (1) der Ausgabebetyp mittels eines deterministischen Automaten gegeben ist und (2) der Zwei-Wege-Baumübersetzer jeden Knoten eines Eingabebaums höchstens begrenzt oft besucht. Wenn der Baumübersetzer zusätzlich mit Wertparametern (call-by-value) ausgestattet ist, hängt die Komplexität ausserdem (exponentiell) von der Anzahl der Parameter ab. In diesem Fall wird ein schneller approximativer Typüberprüfungsalgorithmus präsentiert, der auf kontextfreien Baumgrammatiken basiert. Zum Schluß wird dieser Ansatz von Bäumen auf Zwei-Wege-*Wald*übersetzer verallgemeinert. Diese Übersetzer erlauben zusätzlich Konkatenation als Operation auf Ausgabebäumen.

Contents

1	Introduction	1
1.1	Tree Transducers	5
1.2	Questions	6
1.3	Results	7
1.4	Related Work	8
1.5	Own Publications	11
2	Preliminaries	13
2.1	Trees and Forests	13
2.2	Nodes and Paths	16
2.3	Tree Monoid	18
2.4	Trees with State Calls	22
I	Myhill-Nerode Theorem for DBTTs	23
3	Introduction	25
4	Bottom-Up Tree Transducers	29
4.1	Trim Transducers	35
5	Partition	37
5.1	Partitions of DBTTs	40
5.1.1	Finite Index	41
5.1.2	Path-Finite Partition	43
5.2	Bottom-Up Partition	48
5.2.1	Trim Partition	49
5.2.2	Proper Partition	50
5.2.3	Earliest Partition	54
5.2.4	Unified Partition	58
5.2.5	Congruence Relation	64
5.2.6	Minimal Partition	72

6 Myhill-Nerode Theorem	83
6.1 Minimization of Bottom-Up Transducers	84
6.1.1 Proper Transducers	85
6.1.2 Earliest Transducers	89
6.1.3 Unified Transducers	93
6.1.4 Minimal Transducers	100
6.2 From Minimal Transducers to Partitions	108
6.3 From Partitions to Minimal Transducers	111
7 Conclusion	131
7.1 Future Work	132
II Type Checking Tree Walking Transducers	133
8 Introduction	135
9 Tree Walking Transducers	139
9.1 Notes and References	147
10 Type Checking	149
10.1 Type Checking by Forward Type Inference	149
10.2 Tree Automata	150
10.3 Basic Properties of BTAs	152
10.4 Notes and References	153
11 Type Checking of Tree Walking Transducers	155
11.1 Intersecting 2TTs with Output Types	155
11.2 Deciding Emptiness of 2TTs	157
11.3 Efficient Subcases	164
11.4 Conclusion	166
11.5 Notes and References	167
12 Macro Tree Walking Transducers	169
12.1 Type Checking 2MTTs	174
12.2 Deciding Emptiness of 2MTTs	178
12.3 Input-Linear 2MTTs	178
12.4 Notes and References	182
13 Macro Forest Walking Transducers	185
13.1 Intersecting 2MFTs with Output Types	188
13.2 Deciding Emptiness of 2MFTs	192
13.3 Notes and References	194
14 Conclusion	197

Chapter 1

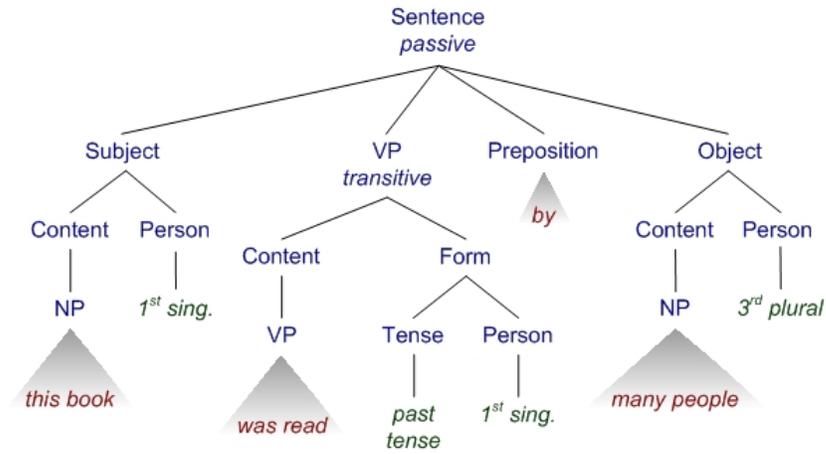
Introduction

Tree transducers are a formal model to describe transformations of tree structured data. Such transformations occur in various fields. Early formal approaches were inspired by compiler development. Internal representations (of the syntax or the semantics) of programs are trees. Separating the syntax of a program language from assigning semantics to the source language led to the analysis of the model of syntax-directed translations [Iro61]. More generally, program analysis, like evaluations and optimizations, are tree transformations. Further works also based on the theory of tree languages, finite tree automata, and the associated algebraic formulation (e.g., [Tha69, Rou70, Tha70, CDG⁺07]). The formal models of tree transducers are more abstract. Input trees are not only syntax trees of program languages, but trees over a ranked alphabet. Already in the year 1969, Thatcher stated that the abstract models of tree transducers are not only generalizations of transformations of program languages, but also of translations of natural languages:

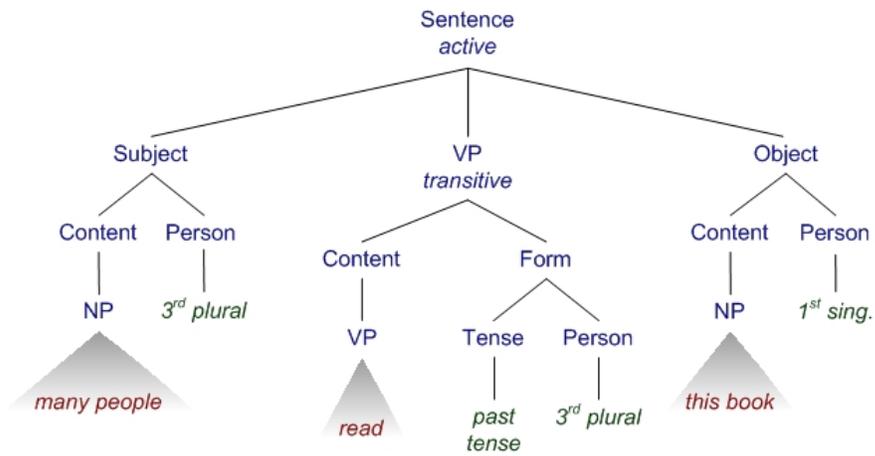
“It appears to be an area with considerable promise of application to questions of syntax and semantics of programming languages and to the analysis of natural languages.” [Tha69]

In the beginning of this century, tree transducers were rediscovered in the context of natural language processing [KG05]. In addition to the string- and phrase-based models, tree-based transformation models are considered in many areas, e.g., in machine translation [Wu97, ADB00, Eis03], i.e., automatic translation from one natural language (like Chinese) into another (like English). Further tasks are, for instance, automatically answering a human-language question [HG01, EM03a], generating natural language from information stored in computer databases [BR00], and automatic summarization [MM99, KM02], i.e., to automatically abstract the relevant information of (multiple) documents.

Due to the inherent syntax and semantics of natural language, sentences will be represented by syntax trees, for instance like the two in Figure 1.1. A transformation model may, for instance, translate passive voice into active



(a) Passive voice: "This book was read by many people."



(b) Active voice: "Many people read this book."

Figure 1.1: Syntax trees of the sentence "This book was read by many people." and its active counterpart "Many people read this book.".

voice, transforming the syntax tree in Figure 1.1(a) into the syntax tree in Figure 1.1(b). The worldwide process of globalization increases the demand for natural language processing tools constantly. An international software product has to be adapted to different markets, languages, and cultures. The World Wide Web (WWW) provides a lot of information. Asking for specific information in human language, causes that a computer has to transform the question to understand it. Furthermore, the computer has to extract the correct information from a huge amount of data, which may be given in different languages. Finally, the computer has to summarize these informations and transform the correct answer into human language [FBCC⁺10].

Lots of work on tree transformations was initiated also by the exchange of data via the WWW. In the end of the 1990s, the extensible markup language XML became the current standard data exchange format for the WWW [ABS00, BPS98]. It allows to represent data in a sequential form preserving the underlying structure. Associated parts are wrapped by matching pairs of open and close tags. For instance, a company structure can be presented by an XML-document like the one in Figure 1.2(a). Such an XML document can be seen as a sequential representation of a tree. The tree, which is described by the given example document is shown in Figure 1.2(b). XML is a meta language. It does not dictate, which tags are allowed and how to use them. Services, which require the data in a specific structure, have to declare the type of *valid* documents. Such types can be defined using an XML type definition language like DTD (Document Type Definition, [BPS08]), XML Schema [FW04], or RELAX NG (Regular Language Description for XML New Generation, [CM01]).

Today, data with an inherent structure is used, exchanged, and restructured everywhere. An internet shop has to present its range of products dependent on the requests of different customers. Emails need an exchange format, which is platform-independent. A company uses data in different services, which need the data in different representations. Booking a flight via an online travel agency will cause a request to web services of different airlines. The airlines, in turn, will respond a list of suitable flights with detailed information about time, prices, and more. The travel agency then has to transform the structure of the data to present it on its webpage.

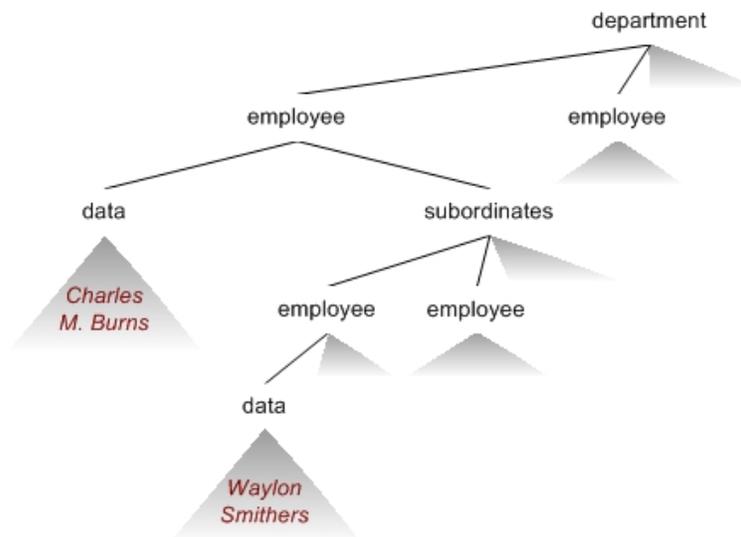
As a consequence, there is a widespread use of transformation and query languages for XML, like the current W3C recommendations XSLT (Extensible Stylesheet Language Transformations, [Kay03]) and XQuery [BCF⁺03]. XSLT is a complex language. A system, which automatically infers an XSLT program from a given set of examples would lighten the work of web programmers. Moreover, it is necessary to check if an XSLT stylesheet produces the desired output for documents of a given input type. On the one hand, it is possible that the transformation is erroneous, especially if it is generated automatically, on the other hand, one may check if the stylesheet also produces valid output for a new input type or a more restricting output type. Assume that the travel agency gets flight informations from a new airline in a new type. The type differs slightly from types of other airlines. The agency will check their existing transformations before developing a new one.

```

<department>
  <employee>
    <data>
      <name> Charles Montgomery Burns </name>
      <id> ... </id> ...
    </data>
    <subordinates>
      <employee>
        <data><name> Waylon Smithers </name> ... </data> ...
      </employee>
      <employee> ... </employee> ...
    </subordinates>
  </employee>
  <employee> ... </employee> ...
</department>

```

(a) XML-document for a company structure.



(b) The company structure seen as a tree.

Figure 1.2: Company structure as XML-document and as a tree.

These two important examples show that data in many areas can be described using tree structures. Consequently, there is a high demand to understand the various transformations between different structures. To analyze transformations of tree structured data, we consider the formal model of tree transducers.

1.1 Tree Transducers

The formal model of tree transducers enhances finite-state word automata in two dimensions. First, they work on trees rather than on strings. Secondly, tree transducers produce output. Here, we consider tree-to-tree transducers over ranked trees, i.e., input and output are trees with finite signatures. The most simple models of tree transducers are top-down and bottom-up tree transducers, cf. Chapter 7 of the textbook [CDG⁺07]. Input trees are transformed starting at the root (top-down) or at the leaves (bottom-up). They were introduced in the 1970s by Rounds and Thatcher [Rou70, Tha70], and Thatcher [Tha73], respectively. Their expressive powers are incomparable, both for nondeterministic and deterministic transducers [Eng75]. Deterministic and nondeterministic top-down tree transducers (DTTts and TTTs, respectively) are able to produce different outputs for the same occurrence of an input subtree, i.e., a (D)TTT copies an input tree and afterwards processes the copies differently. This is not possible by deterministic or nondeterministic bottom-up tree transducers (DBTTs and BTTs, respectively). However, both DBTTs and BTTs can decide whether to generate any output for an input subtree or not after processing the subtree. Moreover, in the nondeterministic case, BTTs can produce copies of the same output tree after nondeterministic processing of the input tree, in contrast to DTTts, TTTs, and DBTTs. Consider, for instance, the example transformation depicted in Figure 1.1: Transforming sentences in passive voice into sentences in active voice. For an input tree like the one in Figure 1.1(a), a (deterministic) bottom-up tree transducer produces an output tree (Figure 1.1(b)) only if the subtree “Preposition(...by...)” is proper. Thereto, it processes this subtree without producing any output for it. This is neither possible by deterministic nor by nondeterministic top-down tree transducers. On the other hand, there are transformations, which are not definable by bottom-up tree transducers, but by top-down tree transducers (see [Eng75] for examples). Especially, if the transformation produces different outputs for the same subtree.

An enhanced model of tree transducers is the model of tree walking transducers (or two-way tree transducers, 2TTs) [Eng09]. 2TTs are essentially the same as the k -pebble tree transducers of [MSV00] for the case $k = 0$. In contrast to the top-down and bottom-up models, the processing of the input tree by a 2TT is not restricted to one direction. After transforming a node, the transducer can stay at this node, proceed with its father, or one of its children. Thus, in contrast to top-down and bottom-up tree transducers, also in the deterministic case, a tree walking transducer does not have to terminate for

every input tree. A 2TT is similar to an attribute grammar, which operates on derivation trees, and has trees as semantic domain (with tree top-concatenation as only semantic operation). A further generalization is the model of macro tree walking transducers (2MTTs). It can be seen as the k -pebble macro tree transducer of [EM03b], for the case that $k = 0$. The 2MTT generalizes the 2TT by adding formal context-parameters to the attributes, i.e., each attribute is seen as a function, which can take parameters of type output tree. The addition of parameters extends the expressive power of tree walking transducers [EM03b]. Tree transducers work on ranked trees, but for practical applications such as XSLT-transformations we have often to deal with unranked forests. A formal model is the model of macro forest walking transducers (2MFTs), which is an extension of 2MTTs. The 2MFT work on forests instead of ranked trees. For instance, such a transducer can handle department structures like the one in Figure 1.2(b), with arbitrary numbers of employees. Macro forest walking transducers are able to concatenate output forests as an additional operation. Such transducers are very expressive and can simulate most features of transformation languages such as XSLT.

This thesis answers important computational problems for deterministic bottom-up tree transducers and the different models of walking transducers.

1.2 Questions

There are a lot of open questions in the context of tree transducers. The most important computational problems are:

1. Type checking: Does a given transducer produce the desired output trees for proper input trees?
2. Equivalence testing: Do two transducers realize the same transformation?
3. Normalization: Can a unique transducer in normal form for a given transformation be constructed?
4. Learning by example: How could a transducer for a transformation be generated of which we only know examples?

The three latter problems are closely related because a common solution may often be constructed with the use of a Myhill-Nerode theorem. The original Myhill-Nerode theorem was stated in 1957/1958 [Myh57, Ner58] for regular word languages and deterministic finite-state automata. It gives a semantic characterization of word languages definable by finite-state automata and provides a unique (up to isomorphism) minimal deterministic finite string automaton for a regular language. Such a semantic characterization and a unique representation of a language or (in the context of transducers) a transformation often give rise to a learning algorithm [Gol78, OGV93, LMN10]. For sequential string-to-string transducers, a Myhill-Nerode style theorem and a learning algorithm exists [Cho03, OGV93]. It is based on the minimal earliest transducers, which

were introduced for strings by Mohri [Moh00]. Recently, it was shown that also for deterministic top-down tree-to-tree transducers, there is a Myhill-Nerode theorem [LMN10]. The proof relies on a new canonical normal form for such transducers, called the “earliest” normal form (inspired by the earliest string transducers of Mohri [Moh00]). Beside this, similar results are only known for string transducers and simple relabeling tree transducers.

The question arises whether there is a Myhill-Nerode theorem for deterministic bottom-up tree transducers.

1.3 Results

In the first part of this thesis, we show that a Myhill-Nerode theorem for deterministic bottom-up tree transducers exists. Even though the idea is similar, obtaining the normal form and a semantic characterization is quite different for DBTTs than for top-down transducers. First, we present a semantic characterization of transformations definable by DBTTs. We identify one important semantic property of these transformations: They are *path-finite*, which means that every path in input trees evokes only finitely many different paths in output trees. If, in addition, the congruence relation of the transformation has finite index, we show that the transformation is definable by a deterministic bottom-up tree transducer. In particular, the Myhill-Nerode theorem leads to a unique minimal DBTT. Given an arbitrary deterministic bottom-up tree transducer, we show that the unique DBTT can be constructed. Generating the “earliest” normal form for a given deterministic bottom-up tree transducer proceeds in four steps: (1) First, we make the transducer “proper”, i.e., we remove all output from states which only produce finitely many different outputs. The output for such states is postponed until the root node of the input tree. This is similar to the “proper normal form” of [AU71, EM03c] (which removes states that produce finitely many outputs, using regular look-ahead). (2) We make the transducer “earliest”, i.e., every state produces output as early as possible during translation. (3) We remove pairwise “ground context unifiers” (this is a technical property, achievable in quadratic time on the transducer). (4) Last, we minimize in the usual way (by merging states that are isomorphic). Steps (2)–(4) can be done in polynomial time, i.e., given a proper transducer, its unique earliest transducer is constructed in polynomial time. Hence, equivalence checking for proper transducers can be done in polynomial time. Constructing a proper transducer (Step 1) takes double exponential time in the worst case. The Myhill-Nerode style theorem can be used to build a Gold style learning algorithm [Gol78], as done for deterministic top-down tree transducers in [LMN10]. Furthermore, the normal form can be used to decide certain (semantic) subclasses of DBTTs; e.g., we can decide whether a given DBTT is equivalent to a relabeling, using the normal form. This provides an alternative proof of [Gaz06], for the deterministic case.

For the first problem, type checking a given transducer w.r.t. given input and output types, we are interested in exact solutions (in contrast to the approxi-

mative approach). The class of transformations for which exact type checking is possible, is surprisingly large [EM03b, MSV03, MBPS05]; the price to be paid for exactness is also extremely high. Even for simple top-down transformations, exact type checking is exponential-time complete [MN05], and for more complex transformations such as the k -pebble transducers of [MSV03], the problem is non-elementary, i.e., cannot be expressed by an iterated exponential function. For practical considerations, however, one is interested in useful subclasses of transformations for which exact type checking is provably tractable. In the second part of this dissertation, we investigate type checking of transformations formulated through tree walking transducers, macro tree walking transducers and macro forest walking transducers. Given suitable descriptions (types) of admissible inputs and outputs for a tree walking transducer, type checking the transducer means to test whether all outputs produced by the 2TT on admissible inputs are again admissible.

Our main result in Part II is: If valid output trees are described by deterministic tree automata, then exact type checking can be done in polynomial time for a large class of practically interesting transformations obtained by putting only mild restrictions on these kinds of transducers. Our approach is forward type inference, i.e., we infer the set of output trees produced by the transducer for valid input trees and compare this set with the given output type. More precisely, we determine if the intersection of the inferred outputs with the complement of the given output type is empty. The presented algorithms solve the exact type checking problem for arbitrary 2TTs, 2MTTs, and 2MFTs, with respect to regular tree (or forest) languages as types. In general, they run in exponential time, but we present subclasses for which they run in polynomial time. We show that if a 2TT visits the same node only constantly often, i.e., is *strictly b -bounded*, then it can be type checked in polynomial time. For the more complex classes of macro tree walking and macro forest walking transducers, we also present approximative algorithms.

1.4 Related Work

The Myhill-Nerode theorem for bottom-up tree automata is the straightforward generalization of the corresponding results for finite string automata. The minimal automaton and the congruence relation appear in [Bra68, Bra69, AG68, EW67], see also Chapter 1.5 in [CDG⁺07]. Special transformations can be defined by tree automata, like relabelings and recognizable tree relations [DT90]. For finite state transducers, a Myhill-Nerode style theorem exists for subsequential string transducers [Cho03, OGV93], which is based on the minimal earliest transducers for strings by Mohri [Moh00]. For tree-to-tree transducers, in [LMN10] a Myhill-Nerode theorem for top-down tree transducers was the basis for a Gold style learning algorithm [Gol78]. In 1980 it was shown that equivalence for deterministic transducers is decidable both in the top-down [Ési80] and bottom-up case [Zac80]. Later, a polynomial-time algorithm for single-valued bottom-up transducers was provided [Sei92].

A valid generalization of both deterministic bottom-up and deterministic top-down tree transducers is the deterministic top-down tree transducer with regular look-ahead [Eng77]. Even though the equivalence problem for DTTTs with regular look-ahead is easily reduced to the one for DTTTs [EMS09], it is an intriguing open problem whether DTTTs with regular look-ahead have a canonical normal form. Another related model of transformation is the attribute grammar [Knu68], seen as a tree transducer [Fül81, FV98]. For attributed tree transducers, decidability of equivalence is still an open problem, but for the special subclass of “nonnested, separated” attribute grammars (those, which can be evaluated in one strict top-down phase followed by one strict bottom-up phase), equivalence is known to be decidable [CFZ82]. This class strictly includes DTTTs (but not DBTTs [FV95]). There are several other interesting incomparable classes of tree translations for which equivalence is known to be decidable, but no normal form (and no complexity) is known, e.g., MSO-definable tree translations [EM06]. This class coincides with single-use restricted attribute grammars or macro tree transducers with look-ahead [EM99]. Is there a canonical normal form for such transducers?

Another interesting generalization are tree-to-*string* transducers. It is a long standing open problem [Eng80] whether or not deterministic top-down tree-to-string transducers (DTTSTs) have decidable equivalence. Recently, for the subclass of *linear sequential* DTTSTs, a unique normal form similar to the earliest normal form was proved [LLN⁺11]. Another recent result states that functional visibly pushdown transducers have decidable equivalence [FRR⁺10]. This class is closely related to non-copying DTTSTs. It raises the question whether our Myhill-Nerode theorem for DBTTs can be extended to functional (but nondeterministic) bottom-up tree transducers.

In its most general setting, the type checking problem for XML transformations is undecidable. Hence, general solutions are bound to be approximative, but seem to work well for practical XSLT transformations [MOS07]. Approximative type checking for XML transformations is typically based on (subclasses of) recognizable tree languages. Using XPath [CD99] as pattern language, XQuery [BCF⁺03] is a functional language for querying XML documents. It is strongly-typed and type checking is performed via type inference rules computing approximative types for each expression. Approximative type inference is also used in XDuce [HP03] and its follow-up version CDuce [Fri04]; navigation and deconstruction are based on an extension of the pattern matching mechanism of functional languages with regular expression constructs. Recently, Hosoya et al. proposed a type checking system based on the approximative type inference of [HP02] for parametric polymorphism for XML [HFC05]. Type variables are interpreted as markings indicating the parameterized subparts. In [MOS07], a sound type checking algorithm is proposed (originally developed for the Java-based language XACT [KMIS04]) based on an XSLT flow analysis that determines the possible outcomes of pattern matching operations; for the benefit of better performance, the algorithm deals with regular approximations of possible outputs.

Milo et al. [MSV03] propose the k -pebble tree transducer (k -PTT) as a for-

mal model for XML transformations, and show that exact type checking can be done for k -PTTs using *inverse* type inference. The latter means to start with an output type O of a transformation τ and then to construct the type of the inputs by backwards translating O through τ . Each k -pebble tree transducer can be simulated by compositions of $k + 1$ stay macro tree transducers (stay-MTTs) [EM03b], thus, type checking can be solved in time (iterated) exponential in the number of used pebbles. Intuitively, k -pebble tree transducers for $k = 0$ correspond to our 2TTs. In [Eng09], it was shown that inverse type inference for tree walking transducers can be done in exponential time, and can be done for k -fold compositions of 2TTs in k -fold exponential time. Maneth et al. show in [MBPS05] that inverse type inference can be done for a transformation language providing all standard features of most XML transformation languages using a simulation by at most three stay-MTTs. Inverse type inference is used in [MN04, MN05] to identify subclasses of top-down XML transformations, which have tractable exact type checking. We note that the classes considered there are incomparable to the ones considered in this thesis.

Tree transducers are a promising model to describe translations on natural languages, but the requirements to transformation models used in this context are complex [KG05]. Because of the ambiguity of natural language, most approaches in machine translation are on statistical models [MS99]. The overview [KG05] of (probabilistic) tree-based models in natural language processing figures out important requirements, like closure properties, on a reasonable formal syntax-based translation model (see also [Mal11]). Until now there is no model known which fulfills all of these properties. Most of the transducers, which are considered as translation models, are extensions of the basic models of bottom-up and top-down tree transducers (as required by Knight et al. in [KG05]). Thus, it is important to understand the basic models of tree transducers, even though the results are not directly applicable to machine translation models. Promising models to handle the probabilities are weighted transducers, i.e., rules are equipped with a weight (see [FV09] for a survey). In [Mal11], weighted (linear and nondeleting) extended top-down tree transducers are analyzed in the context of machine translation where extended means that the patterns on left-hand sides of rules are not restricted to one input symbol. They fulfill most of the requirements [Mal11], but are not closed under composition [AD82]. The different subclasses of transformations induced by extended top-down tree transducers are studied in [MGHK09]. The model of linear extended multi bottom-up tree transducers fulfills most of the requirements stated in [KG05], but it is an open problem if these transducers are efficiently trainable [ELM09, ELM08]. Instead of learning a transducer, the problem in this context is to train probabilistic transducers (given by sets of rules) on huge amounts of collected training data in the form of tree pairs. In [GKM08], training of probabilistic tree-to-tree and tree-to-string transducers is discussed. Especially, for the model of weighted extended top-down tree transducers a training algorithm is provided. In [Eis03] a training algorithm for synchronous tree substitution grammars is given where the model is trained on all possible derivations in proportion to their probabilities.

1.5 Own Publications

This thesis is based on the three original publications [FSM10, FSM11, MFS11].

- The paper [FSM10] was presented in 2010 at the 14th International Conference on Developments in Language Theory (DLT) in London, Ontario, Canada. It covers the minimization of deterministic bottom-up tree transducers detailed in Section 6.1 of Part I of this work.
- A longer version of the afore mentioned paper will be published in the International Journal of Foundations of Computer Science (IJFCS) [FSM11]. It shows some extended proofs also contained in this thesis.
- Part II is based on the chapter [MFS11] of the book “Modern Applications of Automata Theory” [DS11]. It is estimated to appear as Volume 2 of the IISc Research Monographs Series in September 2011.

These three publications are co-authored by Prof. Dr. Sebastian Maneth and my supervisor Prof. Dr. Helmut Seidl.

Chapter 2

Preliminaries

In this chapter, we designate some elementary notations of sets, words, and functions and define trees, forests, and other basic terms.

In the following \mathbb{N} denotes the set of all positive natural numbers and by $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$ we refer to the set of natural numbers including 0. For a constant $k \in \mathbb{N}_0$ we abbreviate by $[k]$ the set $\{1, \dots, k\} \subset \mathbb{N}$ where $[0] = \emptyset$.

For a set X , we denote by X^* the set of *words* or *strings* (including the empty word ϵ) over X . The concatenation of two words u and v over X is denoted by $u.v$ or uv . If there are words u, v , and w such that $u.v = w$, we say that u is a *prefix* and v a *suffix* of w . A subset $W \subseteq X^*$ of words is called *prefix-closed* if for all $u, v \in X^*$ holds: $u.v \in W$ implies $u \in W$. Note that ϵ is element of every non-empty, prefix-closed set.

For two sets X and Y and a subset $D \subseteq X$, a function $\varphi : D \rightarrow Y$ is a partial function from X to Y , written $\varphi : X \dashrightarrow Y$ where φ is undefined for every $x \in X \setminus D$. Then D is the domain of φ . We also use $\text{dom}(\varphi)$ to refer to the domain. The image of a subset $S \subseteq X$ under φ is the set of all images of elements in S for which φ is defined, i.e., $\varphi(S) = \{\varphi(x) \mid x \in D \cap S\}$.

2.1 Trees and Forests

Trees and forests are build over alphabets (i.e., finite sets) of symbols or node labels. An *unranked tree* over an alphabet Σ consists of a root node labeled by a symbol \mathbf{a} from Σ and a forest f , written $\mathbf{a}\langle f \rangle$. A *forest* (or *hedge*) is a sequence of an arbitrary number of unranked trees, written $u_1 u_2 \dots u_m$. The number m is called the *length* of the forest. The empty forest, i.e., a forest with length $m = 0$, is denoted by ϵ .

Definition 2.1 (Forests). Let Σ be an alphabet. The set \mathcal{F}_Σ of forests f over Σ is defined by the grammar rules

$$f ::= \epsilon \mid uf \quad \text{and} \quad u ::= \mathbf{a}\langle f \rangle$$

where $\mathbf{a} \in \Sigma$.

The tree in Figure 1.2(b) is such a forest (with only one non-empty tree) of the form $\mathbf{department}\langle f \rangle \epsilon$ where f is a forest of “**employee**-trees”. The length of f depends on the company structure and varies from company to company.

Rather than on forests, *tree* transducers work on *ranked* trees. There, we assume that a fixed rank is given for every element of Σ , i.e., $\Sigma = \bigsqcup_{m \in \mathbb{N}} \Sigma^{(m)}$ where $\Sigma^{(m)}$ is the set of all symbols with rank m and \bigsqcup denotes the union of distinct sets. Then, Σ is said to be a *ranked alphabet*. We define $\mathit{rank}(\mathbf{a}) = m$ for all symbols $\mathbf{a} \in \Sigma^{(m)}$ for $m \geq 0$. We also write $\mathbf{a}^{(m)}$ to denote a symbol \mathbf{a} of rank m . The maximal rank $\mathit{mr}(\Sigma)$ is the smallest number m such that $\Sigma^{(m)} \neq \emptyset$ and $\Sigma^{(m+i)} = \emptyset$ for all $i \geq 1$.

Definition 2.2 (Ranked Trees). Let Σ be a ranked alphabet. The set \mathcal{T}_Σ of ranked trees over Σ is defined by the grammar rules

$$u ::= \mathbf{a}(\overbrace{u, \dots, u}^{m \text{ times}}) \mid \mathbf{b}$$

where $\mathbf{a} \in \Sigma^{(m)}$ and $\mathbf{b} \in \Sigma^{(0)}$.

In the following, we use the term ‘tree’ as a synonym for *ranked tree*. In the company structure example in Figure 1.2, the symbol **employee** has rank 2. Every node with label **employee** in the tree has exactly two children which have the labels **data** and **subordinates**, respectively. On the other side, for the symbol **department** there is no fixed rank. In different companies, there are different many employees. Thus, in order to define transducers on arbitrary trees and forests, e.g., on XML documents, we rely on ranked tree representations of forests, e.g., through binary trees. The empty forest then is represented by a leaf with label **e** (where **e** is a new symbol that does not appear in the document). The content of an element node **a** is coded as the left child of **a**, while the forest of right siblings of **a** is represented as the right child. This is the well-known *first-child next-sibling encoding*, cf. the textbooks [CDG⁺07, Knu97]. Accordingly, the ranks of symbols are either 0 or 2. Figure 2.1 shows the binary encoding of the tree of the company structure presented in Figure 1.2(b) in the introduction. In the following, we refer to the unranked tree of this example structure by u_B and its encoding by $u_{B'}$. Several other encodings of unranked trees or forests by ranked trees have been used in the literature, for instance the *extension encoding* was presented in [CNT04, Tak75]. The encoding *encode* in [MSV03] is similar to the first-child next-sibling encoding.

A tree u' , which occurs in a tree u is called *subtree* of u . The set $\mathbf{Subtrees}(u)$ of all subtrees of a tree u is recursively defined as follows:

$$\begin{aligned} \mathbf{Subtrees}(\mathbf{b}) &= \{\mathbf{b}\} \\ \mathbf{Subtrees}(\mathbf{a}(u_1, \dots, u_m)) &= \{\mathbf{a}(u_1, \dots, u_m)\} \cup \bigcup_{i \in [m]} \mathbf{Subtrees}(u_i) \end{aligned}$$

where \mathbf{b} and \mathbf{a} are labels of rank 0 and m , respectively. For instance, the tree

$$\mathbf{Form}(\mathbf{Tense}(\mathbf{past \ tense}), \mathbf{Person}(1^{\text{st}} \ \mathbf{sing.}))$$

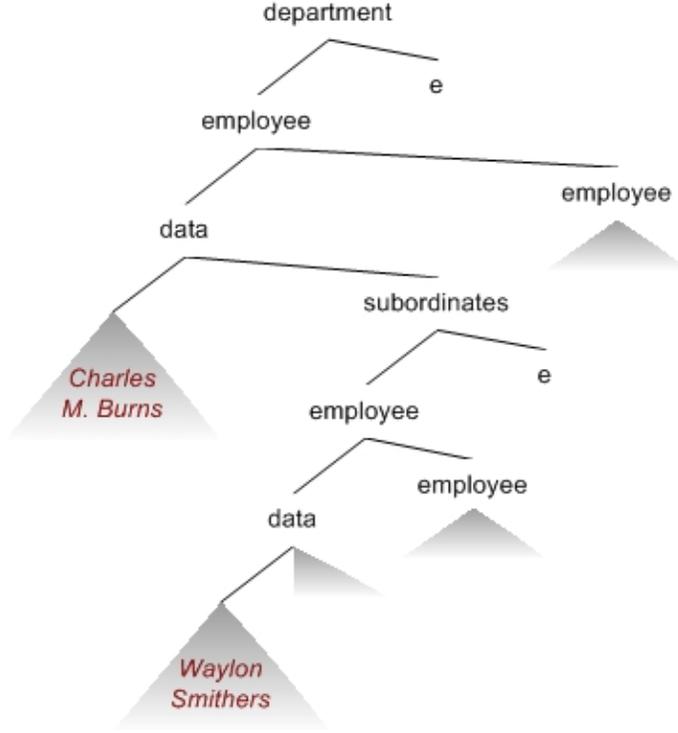


Figure 2.1: The binary encoding u_B' of the tree in Figure 1.2(b).

is a subtree of the syntax tree in Figure 1.1(a). The height of this (sub)tree is 3. We define the *height* of a tree recursively as

$$\begin{aligned} \text{height}(\mathbf{b}) &= 1 && \text{for } \mathbf{b} \in \Sigma^{(0)}, \\ \text{height}(\mathbf{a}(u_1, \dots, u_m)) &= 1 + \max\{\text{height}(u_1), \dots, \text{height}(u_m)\} && \text{for } \mathbf{a} \in \Sigma^{(m)}. \end{aligned}$$

The height of a tree is the maximal length of a path from the root to a leaf plus one (cf. Section 2.2).

We often consider trees with variables. For example, the right-hand sides of rules of deterministic bottom-up tree transducers are such trees (Chapter 4). Let $X = \{x_1, x_2, \dots\}$ denote a countable set of distinct variables of rank 0. Trees with variables of X are trees in $\mathcal{T}_\Sigma(X) = \mathcal{T}_{\Sigma \cup X}$. Let $t \in \mathcal{T}_\Sigma(X)$ be such a tree. We abbreviate by $t[t_1, \dots, t_k]$ the substitution $t[t_1/x_1, \dots, t_k/x_k]$ of trees t_i for the variables x_i ($i \in [k]$) in the tree t . The output of bottom-up tree transducers is built by replacing the variables of trees in $\mathcal{T}_\Sigma(X)$ by trees. For a finite set of m distinct variables, we also write $\mathcal{X}_m = \{x_1, \dots, x_m\}$. Talking about the semantics of transducers, X is often a singleton and, in this case, we also use variable y instead of x_1 .

2.2 Nodes and Paths

The same tree may occur several times as a subtree. For instance, the subtree $\mathbf{Person}(1^{st} \text{ sing.})$ occurs twice in the syntax tree of the sentence in passive voice in Figure 1.1(a). It is necessary to distinguish between a subtree and an occurrence of a subtree. Assigning unique coordinates to the nodes of a ranked tree enables us to indicate a certain occurrence of a subtree by the coordinates of its root.

Definition 2.3 (Nodes). The set $\mathbf{Nodes}(u) \subseteq \mathbb{N}^*$ of all *nodes* ϑ in a ranked tree $u \in \mathcal{T}_\Sigma$ is defined as

$$\begin{aligned} \mathbf{Nodes}(\mathbf{b}) &= \{\epsilon\} \\ \mathbf{Nodes}(\mathbf{a}(u_1, \dots, u_m)) &= \{\epsilon\} \cup \{i\vartheta \mid i \in [m], \vartheta \in \mathbf{Nodes}(u_i)\} \end{aligned}$$

where \mathbf{b} and \mathbf{a} are symbols of rank 0 and m , respectively, and $u_i \in \mathcal{T}_\Sigma$ for $i \in [m]$.

Note that this definition also holds for trees in $\mathcal{T}_\Sigma(X)$ for a set X of variables. The two occurrences of the subtree $\mathbf{Person}(1^{st} \text{ sing.})$ in the tree in Figure 1.1(a) are now uniquely indicated by their root nodes 1.2 and 2.2.2. Here, we write 1.2 instead of 12 for the second child of the first child of the root, to distinguish between this node and the twelfth son of the root. The *direction* $\eta(\vartheta)$ of a node ϑ indicates whether ϑ is the root of the tree or a particular child, i.e., we define $\eta(\epsilon) = 0$ and $\eta(\vartheta'j) = j$. Both nodes 1.2 and 2.2.2 have the direction 2, they are the second child of their fathers. The *depth* $|\vartheta|$ of a node ϑ is the length of the string describing ϑ , i.e., $|\epsilon| = 0$ and $|i\vartheta'| = 1 + |\vartheta'|$. The *size* $|u|$ of a tree u is defined as the number of nodes, i.e., $|u| = |\mathbf{Nodes}(u)|$. For a ranked tree u and a given node $\vartheta \in \mathbf{Nodes}(u)$, the tree $u[\vartheta]$ is called the *subtree of u located at ϑ* and is defined as

$$\begin{aligned} u[\epsilon] &= u \\ \mathbf{a}(u_1, \dots, u_m)[i\vartheta] &= u_i[\vartheta] \quad \text{for } i \in [m]. \end{aligned}$$

With $lab_u(\vartheta)$ we refer to the label of the node ϑ in a tree u , or $lab(\vartheta)$ if u is given by the context. Let u_p be the tree in Figure 1.1(a). Then, we get

$$u_p[1.2] = u_p[2.2.2] = \mathbf{Person}(1^{st} \text{ sing.}).$$

The label of these nodes is $lab_{u_p}(1.2) = lab_{u_p}(2.2.2) = \mathbf{Person}$.

Similar notions also apply to forests. In particular, the set $\mathbf{Nodes}(f)$ of nodes of a forest f is defined as

$$\begin{aligned} \mathbf{Nodes}(\epsilon) &= \emptyset \\ \mathbf{Nodes}(\mathbf{a}(f_1)f_2) &= \{0\} \cup \{0\vartheta \mid \vartheta \in \mathbf{Nodes}(f_1)\} \cup \{(i+1)\vartheta \mid i\vartheta \in \mathbf{Nodes}(f_2)\}. \end{aligned}$$

Note that the definition of nodes of a ranked tree differs from the definition of nodes in a forest consisting of one tree only. Accordingly, also the definition of *direction* differs. For a node ϑ in a forest we define the *direction* $\eta(\vartheta)$, which

now indicates whether ϑ is at the top-level, has a left sibling or both. Thus, $\eta(0) = 0$, $\eta(i) = 1$ for $i > 0$, $\eta(\vartheta'0) = 2$ for $\vartheta' \neq \epsilon$ and $\eta(\vartheta) = 3$ otherwise. The subforest $f[\vartheta]$ at a node ϑ in a forest is defined by:

$$f[i] = \begin{cases} f & \text{if } i = 0 \\ f'[i-1] & \text{if } i > 0 \wedge f = uf' \end{cases}$$

Moreover, for $\vartheta = i\vartheta'$ with $\vartheta' \neq \epsilon$,

$$f[\vartheta] = \begin{cases} f_1[\vartheta'] & \text{if } i = 0 \wedge f = \mathbf{a}\langle f_1 \rangle f_2 \\ f'[(i-1)\vartheta'] & \text{if } i > 0 \wedge f = uf' . \end{cases}$$

The *label* $lab_f(\vartheta)$ of ϑ in the forest f is defined by $lab_f(0) = \epsilon$ if $f = \epsilon$, and

$$lab_f(i\vartheta') = \begin{cases} \mathbf{a} & \text{if } i = 0 \wedge \vartheta' = \epsilon \\ lab_{f_1}(\vartheta') & \text{if } i = 0 \wedge \vartheta' \neq \epsilon \\ lab_{f_2}((i-1)\vartheta') & \text{if } i > 0 \end{cases}$$

if $f = \mathbf{a}\langle f_1 \rangle f_2$. Note that the label at a node in a forest thus either is from Σ or equals the empty forest ϵ . Consider for instance the unranked tree u_B (in Figure 1.2(b)) seen as a forest. We observe that the **subordinate**-node has the coordinates 0.0.1. The direction of this node is $\eta(0.0.1) = 3$. The subforest describing the personal data of Waylon Smithers (with **name** as root label of the first tree) is $u_B[0.0.1.0.0.0]$. Its father has the label $lab_{u_B}(0.0.1.0.0) = \mathbf{data}$. The direction of this node in u_B is $\eta(0.0.1.0.0) = 2$. Note that the subforest indicated by this node, i.e., $u_B[0.0.1.0.0]$ consists of the tree rooted at this **data**-node (the personal data of Waylon Smithers) and all its siblings.

In Part I we are not only interested in the node itself but also in the path from the root to the node including the sequence of labels at this path. A Σ -*path* is a string over the set $\{(\mathbf{a}, i) \mid \mathbf{a} \in \Sigma, 1 \leq i \leq \text{rank}(\mathbf{a})\}$.

Definition 2.4 (Paths). The set $\text{Paths}(u)$ of *paths* of the tree u is inductively defined as:

$$\begin{aligned} \text{Paths}(\mathbf{b}) &= \{\epsilon\} \\ \text{Paths}(\mathbf{a}(u_1, \dots, u_m)) &= \{\epsilon\} \cup \{(\mathbf{a}, i).p \mid i \in [m], p \in \text{Paths}(u_i)\} \end{aligned}$$

The set of all paths over Σ is $\text{Paths}_\Sigma = \bigcup_{u \in \mathcal{T}_\Sigma} \text{Paths}(u)$.

The *node* of a path p is its stepwise projection to the second component, i.e.,

$$\text{node}(\epsilon) = \epsilon \quad \text{and} \quad \text{node}((\mathbf{a}, i).p') = i.\text{node}(p') .$$

The *subtree* $u[p]$ of a tree u at a path p and the *label* $lab_u(p)$ of p in u are given by the node of p in u : $u[p] = u[\text{node}(p)]$ and $lab_u(p) = lab_u(\text{node}(p))$ if $p \in \text{Paths}(u)$, otherwise both are undefined. The *length* $|p|$ of a path p is given by the depth of its node, it is $|p| = |\text{node}(p)|$. We say that a path p *belongs to* a tree u if $p \in \text{Paths}(u)$.

Let $u_F = u_p[2.2]$ the subtree of the syntax tree u_p with root label **Form** (cf. Figure 1.1(a)). The path $p = (\mathbf{Form}, 2)(\mathbf{Person}, 1)$ refers to the leaf with label ‘1st sing.’. The node of this path is $node(p) = 2.1$ and the length of this path is $|p| = 2$. Remember that the height of this tree $|u_F|$ defined by the length of the longest path from the root to a leaf plus one is 3 (cf. Page 15).

Let u be a tree in \mathcal{T}_Σ . For a symbol $\mathbf{a} \in \Sigma$, the set $\mathbf{Paths}_\mathbf{a}(u)$ denotes the set of all paths $p \in \mathbf{Paths}(u)$ with $lab_u(p) = \mathbf{a}$. For a set $\Sigma' \subseteq \Sigma$, the set $\mathbf{Paths}_{\Sigma'}(u) = \bigcup_{\mathbf{a} \in \Sigma'} \mathbf{Paths}_\mathbf{a}(u)$ is the set of Σ' -labeled paths of tree $u \in \mathcal{T}_\Sigma$. Note that these definitions hold also for trees with variables if $\mathcal{T}_\Sigma(X)$ is interpreted as $\mathcal{T}_{\Sigma \cup X}$. Mostly, we consider for a variable y the y -paths of a tree in $\mathcal{T}_\Sigma(y)$, i.e., the set $\mathbf{Paths}_y(t)$. We also say *variable path* if the variable is clear by the context.

2.3 Tree Monoid

In the first part of this thesis, we need to distinguish different kinds of trees: Trees without any variable, trees with at least one occurrence of a dedicated variable y , and trees with exactly one occurrence of y . The latter are called *context*. Together, trees with or without occurrences of the variable y build a monoid with substitution into the variable as binary operation.

First, we consider trees possibly containing a dedicated variable $y \notin \Sigma$ of rank 0. Let $\mathcal{T}_\Sigma(y)$ denote this set: $\mathcal{T}_\Sigma(y) = \mathcal{T}_{\Sigma \cup \{y\}}$. On $\mathcal{T}_\Sigma(y)$, we define a binary operation “.” by substitution into the variable y :

$$t_1 \cdot t_2 = t_1[t_2/y]$$

Note that every occurrence of y in t_1 will be replaced by the same tree t_2 . For example, $\mathbf{Form}(y, y) \cdot \mathbf{Person}(y) = \mathbf{Form}(\mathbf{Person}(y), \mathbf{Person}(y))$. The result of $t_1 \cdot t_2$ is a ground tree, i.e., does not contain y , iff $t_1 \in \mathcal{T}_\Sigma$ or $t_2 \in \mathcal{T}_\Sigma$. For instance, $3^{rd} \text{ plural} \cdot \mathbf{Person}(y) = 3^{rd} \text{ plural}$. Moreover, the operation “.” is associative with neutral element y . Therefore, the set $\mathcal{T}_\Sigma(y)$ together with the operation “.” and y forms a monoid. Let $\hat{\mathcal{T}}_\Sigma(y)$ denote the sub-monoid consisting of all trees, which contain at least one occurrence of y , i.e., $\hat{\mathcal{T}}_\Sigma(y) = \mathcal{T}_\Sigma(y) \setminus \mathcal{T}_\Sigma$. Then $\mathcal{T}_\Sigma(y) = \hat{\mathcal{T}}_\Sigma(y) \cup \mathcal{T}_\Sigma$, and we have:

Proposition 2.1. [Eng80]

1. Let $z, z', t_1, t_2, t'_1, t'_2 \in \mathcal{T}_\Sigma(y)$ with $t_1 \neq t_2$ and $t'_1 \neq t'_2$. Assume that the two equalities $z \cdot t_1 = z' \cdot t'_1$ and $z \cdot t_2 = z' \cdot t'_2$ hold. Then one of the following two assertions is true:

- (a) $z, z' \in \mathcal{T}_\Sigma$ and $z = z'$; or
- (b) both z and z' contain an occurrence of y , i.e., are from $\hat{\mathcal{T}}_\Sigma(y)$ and $z \cdot s = z' \cdot s$ for some $s \in \hat{\mathcal{T}}_\Sigma(y)$.

2. The sub-monoid $\hat{\mathcal{T}}_\Sigma(y)$ is free.

If there are $z, z' \in \mathcal{T}_\Sigma(y)$ and $t, t' \in \hat{\mathcal{T}}_\Sigma(y)$ with $z \cdot t = z' \cdot t'$, then already one of the two assertions (a) and (b) is true.

Unfortunately, this proposition does not hold for forests. Let $\mathcal{F}_\Sigma(y)$ be the set of forests possibly containing the variable y (similar to $\mathcal{T}_\Sigma(y)$). For the forests $f = \mathbf{b}yy$ and $f' = yy\mathbf{b}$ (i.e., forests with three trees each of height 1), we get that

$$f[y\mathbf{b}/y] = f'[y\mathbf{b}/y],$$

but neither $f = f'$ nor is there a forest $g \in \hat{\mathcal{F}}_\Sigma(y)$ such that $f[g/y] = f'$ or $f = f'[g/y]$ where $\hat{\mathcal{F}}_\Sigma(y)$ is the set of all forests with at least one occurrence of y . Since most of the results of Part I rely on Proposition 2.1, they cannot be adapted to forest transformations in a straightforward manner.

Furthermore, we also need the following relations. Consider the set

$$\hat{\mathcal{T}}_\Sigma(y)_\perp = \hat{\mathcal{T}}_\Sigma(y) \cup \{\perp\}$$

of all trees containing at least one occurrence of the variable y enhanced with an extra bottom element \perp (not in $\Sigma \cup \{y\}$). On this set, we define a partial ordering by $\perp \sqsubseteq s$ for all s , and $s_1 \sqsubseteq s_2$ for $s_1, s_2 \in \hat{\mathcal{T}}_\Sigma(y)$ iff $s_1 = s' \cdot s_2$ for a suitable $s' \in \hat{\mathcal{T}}_\Sigma(y)$. Note that the trees s_1, s_2 and s' are not ground, i.e., they contain at least one occurrence of y . For example, for the trees on Page 18, we have $\perp \sqsubseteq \mathbf{Form}(\mathbf{Person}(y), \mathbf{Person}(y)) \sqsubseteq \mathbf{Person}(y)$. The greatest element w.r.t. the ordering \sqsubseteq is y while the least element is given by \perp . With respect to this ordering, we observe:

1. Every $s \in \hat{\mathcal{T}}_\Sigma(y)$ has finitely many upper bounds only.
2. For every $s_1, s_2 \in \hat{\mathcal{T}}_\Sigma(y)_\perp$, there exists a least upper bound $s_1 \sqcup s_2$ in $\hat{\mathcal{T}}_\Sigma(y)_\perp$.

Since $\hat{\mathcal{T}}_\Sigma(y)_\perp$ also has a least element, namely \perp , we conclude that $\hat{\mathcal{T}}_\Sigma(y)_\perp$ is a *complete* lattice satisfying the ascending chain condition, i.e., every set $S \subseteq \hat{\mathcal{T}}_\Sigma(y)_\perp$ has a least upper bound $s = \bigsqcup S$, and there are no infinite strictly ascending sequences $\perp \sqsubset s_1 \sqsubset s_2 \sqsubset \dots$ where $s \sqsubset s'$ iff $s \sqsubseteq s'$ and $s \neq s'$. We call the least upper bound of a set $S \subseteq \hat{\mathcal{T}}_\Sigma(y)_\perp$ the *greatest common suffix* of all trees in S and denote it by $\text{gcs}(S)$. Thus, we have

$$\text{gcs}(S) = \bigsqcup S$$

w.r.t. the order \sqsubseteq on $\hat{\mathcal{T}}_\Sigma(y)_\perp$. We call a tree $s \in \hat{\mathcal{T}}_\Sigma(y)$ *irreducible* if $s \neq y$ and $s \sqsubseteq s'$ only holds for $s' \in \{y, s\}$. The tree $\mathbf{Person}(y)$ is irreducible.

Example 2.1. According to the example of syntax trees for sentences (Figure 1.1), consider the trees

$$\begin{aligned} t_{s3} &= \mathbf{Subject}(\mathbf{Content}(y), \mathbf{Person}(3^{rd} \text{ plural})) \\ t_{s1} &= \mathbf{Subject}(\mathbf{Content}(y), \mathbf{Person}(1^{st} \text{ sing.})) \\ t_{o1} &= \mathbf{Object}(\mathbf{Content}(y), \mathbf{Person}(1^{st} \text{ sing.})) . \end{aligned}$$

For the first tree, we observe that its subtree $\mathbf{Content}(y)$ is greater than t_{s3} , i.e., $t_{s3} \sqsubseteq \mathbf{Content}(y)$. This follows because we can decompose the tree as follows:

$$t_{s3} = \mathbf{Subject}(y, \mathbf{Person}(3^{rd} \text{ plural})) \cdot \mathbf{Content}(y)$$

Also y is greater than t_{s3} and for both other trees, we get the same upper bounds. Additionally, each of these trees is a upper bound of itself. There are no more upper bounds of these trees. Furthermore, $\mathbf{Content}(y) \sqsubseteq y$. Thus, we get that the greatest common suffix of the set $S = \{t_{s3}, t_{s1}, t_{o1}\}$ is

$$\mathbf{gcs}(S) = \mathbf{Content}(y) .$$

It is the least tree of the common upper bounds of t_{s3}, t_{s1} , and t_{o1} .

Now consider the trees

$$\begin{aligned} t_{sy} &= \mathbf{Subject}(\mathbf{Content}(y), \mathbf{Person}(y)) \\ t_{oy} &= \mathbf{Object}(\mathbf{Content}(y), \mathbf{Person}(y)) . \end{aligned}$$

Neither $\mathbf{Content}(y)$ nor $\mathbf{Person}(y)$ is an upper bound of one of these trees. Only y and t_{sy} are upper bounds of t_{sy} . Thus, this tree is irreducible. Analog, it follows that t_{oy} is irreducible. We get that the greatest common suffix of t_{sy} and t_{oy} is y . \triangleleft

We enhance the definition of the partial order \sqsubseteq to trees in $\mathcal{T}_\Sigma(y) \cup \{\perp\}$ by $\perp \sqsubseteq t$ for all $t \in \mathcal{T}_\Sigma(y) \cup \{\perp\}$, and $t_1 \sqsubseteq t_2$ for $t_1, t_2 \in \mathcal{T}_\Sigma(y)$ iff $t_1 = s \cdot t_2$ for a suitable $s \in \hat{\mathcal{T}}_\Sigma(y)$. Note that the prefix s of t_1 must contain y . With this condition, the tree t_2 really occurs in t_1 .

Two trees (with or without the variable y) are unifiable if they are prefixes of the same tree. Let $\top \notin \Sigma \cup \{y, \perp\}$ be a new symbol. Assume that $t_1, t_2 \in \mathcal{T}_\Sigma(y)$ are trees, and that there are trees $s_1, s_2 \in \mathcal{T}_\Sigma(y) \cup \{\top\}$ such that $t_1 \cdot s_1 = t_2 \cdot s_2$. Note that $t_i \cdot s_i = t_i$ if $t_i \in \mathcal{T}_\Sigma$. Then we call t_1, t_2 *unifiable* and $\langle s_1, s_2 \rangle$ a *unifier* of t_1, t_2 . Note that we substitute the variable y in t_1 and t_2 possibly differently. In terms of the common unification problem [BS01, KB70], we can interpret the trees t_1 and t_2 as trees with different variables, i.e., unifying t_1 and $t_2[y']$ for some new variable $y' \neq y$.

We consider the set

$$\mathbb{D}_\Sigma = (\{y\} \times \hat{\mathcal{T}}_\Sigma(y)) \cup (\hat{\mathcal{T}}_\Sigma(y) \times \{y\}) \cup (\mathcal{T}_\Sigma \cup \{\top\})^2 \cup \{\perp\}$$

of *candidate unifiers*. The set \mathbb{D}_Σ forms a complete lattice w.r.t. the ordering \leq defined by

- $\perp \leq d \leq \langle \top, \top \rangle$ for all $d \in \mathbb{D}_\Sigma$,
- $\langle t_1, t_2 \rangle \leq \langle t'_1, t'_2 \rangle$ if $t_i = t'_i \cdot u$ for all $i \in \{1, 2\}$ for some tree $u \in \mathcal{T}_\Sigma \cup \{y\}$, and
- $\langle t_1, t_2 \rangle \leq \langle t_1, \top \rangle$ and $\langle t_1, t_2 \rangle \leq \langle \top, t_2 \rangle$ if $t_1, t_2 \in \mathcal{T}_\Sigma$.

Then the *most-general unifier* $\text{mgu}(t_1, t_2) \in \mathbb{D}_\Sigma$ for trees $t_1, t_2 \in \mathcal{T}_\Sigma(y)$ is the greatest unifier of t_1, t_2 w.r.t. the ordering \leq . It is \perp if t_1, t_2 are not unifiable. Furthermore, for a set of pairs $P \subseteq \mathcal{T}_\Sigma(y)^2$ the most-general unifier $\text{mgu}(P)$ is the least upper bound of the unifiers of pairs in P , i.e.,

$$\text{mgu}(P) = \bigvee \{ \text{mgu}(t_1, t_2) \mid (t_1, t_2) \in P \} .$$

If t_1 is ground and t_1, t_2 are unifiable by some pair $\langle s_1, s_2 \rangle$ then $t_1 \cdot s'_1 = t_1 = t_2 \cdot s_2$ holds for all $s'_1 \in \mathcal{T}_\Sigma(y)$. It follows that t_1, t_2 are unifiable by each pair $\langle s'_1, s_2 \rangle$ with $s'_1 \in \mathcal{T}_\Sigma(y)$. Thus, the most-general unifier is either $\langle \top, s_2 \rangle$ if t_2 is not ground, or $\langle \top, \top \rangle$ if also $t_2 \in \mathcal{T}_\Sigma$.

Example 2.2. Again, we consider parts of syntax trees of sentences. Thereto, let Σ be a ranked alphabet containing the symbols $\text{Form}^{(2)}$, $\text{Person}^{(1)}$, $\text{Tense}^{(1)}$, 3^{rd} plural $^{(0)}$, and past tense $^{(0)}$ (the exponents give the ranks of the symbols). Consider the trees

$$\begin{aligned} t_1 &= \text{Form}(y, \text{Person}(3^{rd} \text{ plural})) \\ t_2 &= \text{Form}(\text{Tense}(\text{past tense}), \text{Person}(y)) \\ t_3 &= \text{Form}(\text{Tense}(y), y) \\ t_4 &= \text{Form}(\text{Tense}(y), \text{Person}(3^{rd} \text{ plural})) \end{aligned}$$

The trees t_1 and t_2 are unifiable by the unifier $\langle \text{Tense}(\text{past tense}), 3^{rd} \text{ plural} \rangle$. This pair is already the most-general unifier $\text{mgu}(t_1, t_2)$ of the two trees. The trees t_1 and t_3 have the most-general unifier

$$\text{mgu}(t_1, t_3) = \langle \text{Tense}(\text{Person}(3^{rd} \text{ plural})), \text{Person}(3^{rd} \text{ plural}) \rangle ,$$

whereas the trees t_2 and t_3 are not unifiable, i.e., $\text{mgu}(t_2, t_3) = \perp$. Finally, consider the trees t_4 and t_1 , for which $\langle \text{Tense}(\text{past tense}), \text{past tense} \rangle$ is a unifier. Also $\langle \text{Tense}(3^{rd} \text{ plural}), 3^{rd} \text{ plural} \rangle$ or $\langle \text{Tense}(y), y \rangle$ are possible unifiers of this pair. Thus, we observe that the most-general unifier is in $\hat{\mathcal{T}}_\Sigma(y) \times \{y\}$, it is

$$\text{mgu}(t_1, t_4) = \langle \text{Tense}(y), y \rangle .$$

Note that the unifier $\langle \text{Tense}(\text{Tense}(y)), \text{Tense}(y) \rangle$ of t_1 and t_4 is no candidate unifier, i.e., is not in \mathbb{D}_Σ , because both parts are in $\hat{\mathcal{T}}_\Sigma(y)$ and different to y . \triangleleft

For the semantics of transducers, we also need the definition of contexts over an alphabet Σ . A *context* c is a tree $c \in \hat{\mathcal{T}}_\Sigma(y)$, which contains exactly one occurrence of y . The set of all contexts is denoted by $\mathcal{C}_\Sigma(y)$ (or \mathcal{C}_Σ). The *length* of a context c is the length of the path from the root to y . If context c has length n , then there are irreducible trees c_1, \dots, c_n such that $c = c_1 \cdot c_2 \cdots c_n$. In the previous example, the trees t_1, t_2 , and t_4 are contexts, in contrast to t_3 . The length of t_4 is 2 because the y -path of t_4 is $(\text{Form}, 1)(\text{Tense}, 1)$.

Let \sim be an equivalence relation over a set $U \subseteq \mathcal{T}_\Sigma$. It is a *congruence* (w.r.t. to contexts in $\mathcal{C}_\Sigma(y)$) if for all $c \in \mathcal{C}_\Sigma(y)$ and all $u, u' \in U$ holds:

$$u \sim u' \text{ and } c \cdot u \in U \quad \implies \quad c \cdot u' \in U \text{ and } c \cdot u \sim c \cdot u'$$

We write U/\sim to denote the set of equivalence classes of \sim over U .

2.4 Trees with State Calls

In Part II, we consider (macro) tree walking transducers. The right-hand sides of rules of such transducers are trees containing *state calls*. The definitions related to trees like the height of a tree have to be enhanced to trees with state calls. Tree walking transducers (2TTs, cf. Section 9), which do not support accumulating parameters, have rules like, e.g.,

$$q_I(\text{department}) \rightarrow \text{staff}(q(\text{down}_1), \mathbf{e})$$

where q_I and q are states, **department** is an input symbol and **staff** and **e** are output symbols of rank 2 and 0, respectively. For 2TTs, state calls are of the form $q(op)$ where op stands for *up*, *stay*, or *down_i*, like $q(\text{down}_1)$ in the example rule. To enhance the definitions related to trees to right-hand sides of rules of 2TTs, we deal with $q(op)$ as a symbol of rank 0 for all states q and all operations op . The right-hand sides of tree walking transducers are trees over the alphabet $\Sigma \cup (Q \times Op)$ where Q is the set of states and Op the set of possible operations: $Op = \{up, stay\} \cup \{down_i \mid i \in \mathbb{N}\}$. We get $height(q(op)) = 1$ for all states q and all operations op . The right-hand side $\text{staff}(q(\text{down}_1), \mathbf{e})$ of the example rule has height 2.

In Chapter 12, we then will add parameters to state calls to obtain *macro tree walking transducers* (2MTTs). In this case, we fix a set $Y = \{y_1, y_2, \dots\}$ of formal *parameters*. These parameters are of rank 0. We use the same definition as for sets of variables X (Page 15). Then, states have a fixed rank, which determines the number of parameters. A rule of a 2MTT is for instance:

$$q_I(\text{department}, 0) \rightarrow \text{staff}(q(\text{down}_1, \mathbf{e}, \mathbf{e}), \mathbf{e})$$

On the right-hand side, we have the state q of rank 3. The exact definition of rules of 2MTTs with ranked states will be given in Chapter 12. In a 2MTT, a state call has the form $q(op, t_1, \dots, t_n)$ where t_i ($i \in [n]$) are ranked trees over $\Sigma \cup Y$ and further state calls. These right-hand sides are seen as trees over the alphabet $\Sigma \cup Y \cup (Q \times Op)$ where for all states q and all operations op , a label of the form (q, op) has the rank $rank(q) - 1$. Figure 2.2 illustrates the tree representation of the right-hand side of the example rule. There, the state call (q, down_1) has rank 2. In general, there may occur further state calls in the subtrees. Additionally, if the state on the left-hand side has a rank greater than 1, there may occur parameters y_i on the right-hand side. For these trees on right-hand sides of 2MTTs, we define the height recursively by:

$$height(q(op, t_1, \dots, t_n)) = 1 + \max(height(t_1), \dots, height(t_n))$$

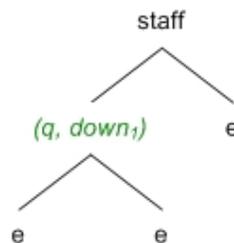


Figure 2.2: Tree representation of a right-hand side in a 2MTT.

Part I

Myhill-Nerode Theorem for Bottom-Up Tree Transducers

Chapter 3

Introduction

In Chapter 1, we pointed out the relevance of a Myhill-Nerode theorem for tree transducers. In this part, we now present a Myhill-Nerode theorem for deterministic bottom-up tree transducers (DBTTs). We provide a semantic characterization of tree transformations, which are definable by DBTTs. Thereto, we split a tree-to-tree transformation into a subtree function, which provides an output tree for every subtree of an input tree, and context functions, which map contexts to images. Such a division is called *partition*. In Figure 3.1, a partition for the example transformation, which translates sentences in passive voice into sentences in active voice (Figure 1.1), is depicted. The **Object**-subtree of the input tree is transformed into the **Subject**-subtree of the output tree. Furthermore, we get an image for the context of the **Object**-subtree, the remainder of the output tree. Note that the grammatical person of the verb phrase in the image of the context is determined by the subtree. Thus, different subtrees may induce different transformations of the same contexts. We will specify a unique partition for every transformation. Furthermore, for this unique partition, we will present necessary conditions that the transformation is definable by a DBTT. These conditions are that the partition has *finite index* and is *path-finite*. Roughly spoken, having finite index means that a partition has only finitely many different context functions and path-finiteness means that each path in input trees from the root to a node only affects finitely many different paths in output trees. Both properties are caused by, amongst others, the finiteness of the number of states of deterministic bottom-up tree transducers. We will prove that these two conditions are not only necessary but also sufficient, i.e., a transformation is definable by a DBTT if and only if its partition is path-finite and has finite index.

Moreover, the path-finite partition with finite index of a transformation leads to a unique minimal DBTT. Given an arbitrary deterministic bottom-up tree transducer, we show that this unique DBTT can be constructed. The construction is based on a sequence of normalizing transformations, which, among others, guarantee that non-trivial output is produced as early as possible. For a deterministic bottom-up tree transducer where every state produces either none

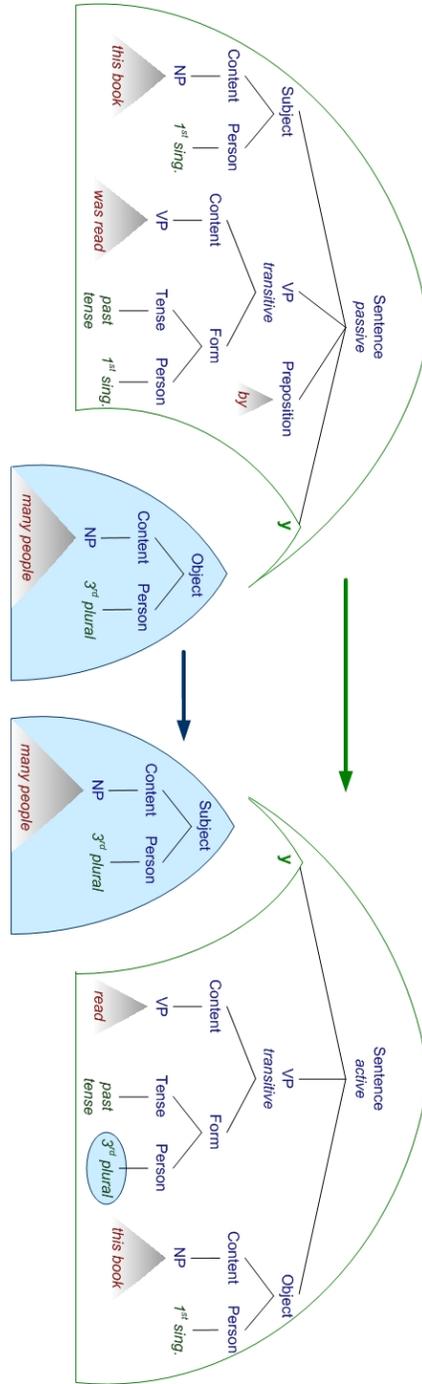


Figure 3.1: Transformation of a subtree and its context.

or infinitely many outputs, the minimal transducer can be constructed in polynomial time. This leads to a polynomial time algorithm to decide equivalence of deterministic bottom-up tree transducers.

Summarizing, we get the Myhill-Nerode theorem for deterministic bottom-up tree transducers. It says that for a tree transformation the following properties are equivalent:

1. The transformation is definable by a DBTT.
2. It exists a unique minimal DBTT describing this transformation.
3. The partition of the transformation is path-finite and has finite index, i.e., is a *bottom-up partition*.

Part I is organized as follows: Some basic notations were already given in the previous Chapter 2. In **Chapter 4**, we add the formal definition of deterministic bottom-up tree transducers and related notions and properties.

The semantic representation via partitions is given in **Chapter 5**. First, we define the basic terms that are used in this context such as subtree functions, context functions, and partitions. Every deterministic bottom-up tree transducer induces a dedicated partition. We deduce this partition and the inherent properties like path-finiteness. The necessity of these properties is proven. In the second section of this chapter (Section 5.2), we consider arbitrary transformations. We present a normal form for the partition of a transformation, which is the unique bottom-up partition if there is a path-finite partition of the transformation with finite index. This partition is disambiguated by some properties like producing non-trivial output as early as possible, i.e., by the subtree function, whereas trivial output is postponed to the context functions. In the end, the bottom-up partition is the partition of the same transformation with the minimal index fulfilling these properties.

In **Chapter 6** we prove the Myhill-Nerode theorem. First, we present the minimization of arbitrary deterministic bottom-up tree transducers (Section 6.1). It is shown that the minimal DBTT is unique and we prove that the construction can be done in polynomial time if states only produce infinitely many different output trees or none. The second part of the proof of the Myhill-Nerode theorem (Section 6.2) is that the partition of a minimal DBTT is already the bottom-up partition of its transformation. In the last section of this chapter (Section 6.3), it is shown that the bottom-up partition of a transformation induces a minimal DBTT, defining the same transformation. It is a technical procedure where we deduce the transition rules and states and prove their correctness.

We present conclusions and ideas of future work in **Chapter 7**. In particular, we give an outlook to a learning algorithm of deterministic bottom-up tree transducers.

Chapter 4

Bottom-Up Tree Transducers

In this part, we analyze deterministic bottom-up tree transducers. Such a finite state tree transducer starts at the leaves of a tree and processes the tree to the root. During the execution, the transducer builds a new tree starting with several subtrees. For instance, a bottom-up transducer for the example transformation in Figure 1.1 produces output trees for the four subtrees of the root node as illustrated in Figure 4.1 and memorizes crucial information (such as the person or the tense) in a state. Depending on these states and the label **Sentence passive** of the root node, it builds a tree using the four output trees (Figure 4.2).

For the following definition, we recall the notation of $\mathcal{X} = \{x_1, x_2, \dots\}$ for a countable set of variables and $\mathcal{X}_m = \{x_1, \dots, x_m\}$ for a finite subset of \mathcal{X} with m variables (Page 15). For a ranked alphabet Σ , $mr(\Sigma)$ denotes the maximal rank of symbols in Σ .

Definition 4.1. A *deterministic bottom-up tree transducer* (DBTT for short) is a tuple $T = (Q, \Sigma, \Delta, R, F)$ where

- Q is a finite set of states,
- Σ and Δ are ranked input and output alphabets, respectively, disjoint with Q ,
- $R : \bigcup_{m=0}^{mr(\Sigma)} (\Sigma^{(m)} \times Q^m) \dashrightarrow Q \times \mathcal{T}_\Delta(\mathcal{X})$ is a partial *transition function*, and
- $F : Q \dashrightarrow \mathcal{T}_\Delta(y)$ is a partial function mapping states to final outputs, called the *final function* of T .

For every input symbol $\mathbf{a} \in \Sigma$ of rank m and sequence of states q_1, \dots, q_m , the transition function R provides at most one transition (or rule), which is denoted by $\mathbf{a}(q_1, \dots, q_m) \rightarrow q(z)$ where $q \in Q$ and $z \in \mathcal{T}_\Delta(\mathcal{X}_m)$.

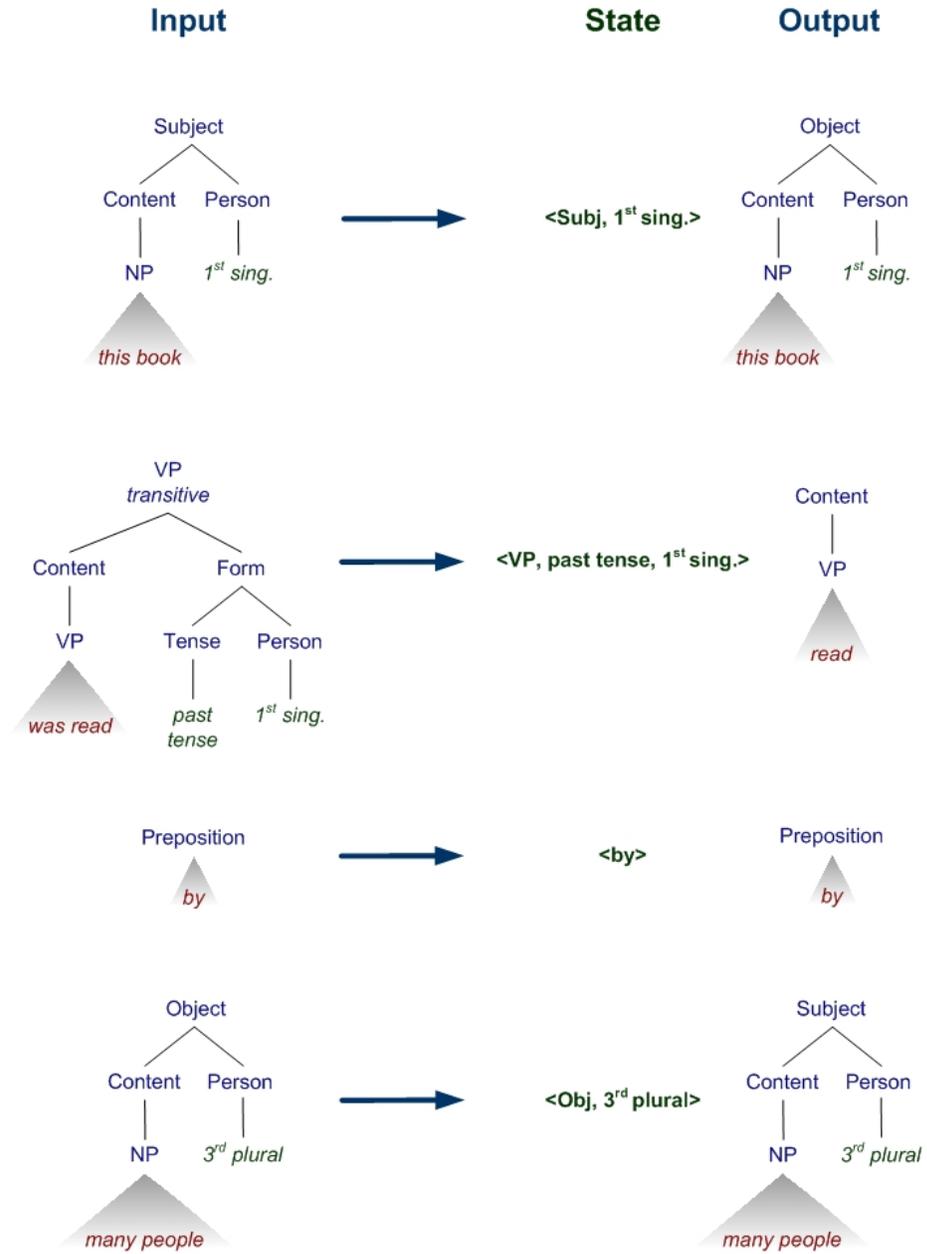


Figure 4.1: The outputs of selected subtrees and the reached states with further informations.

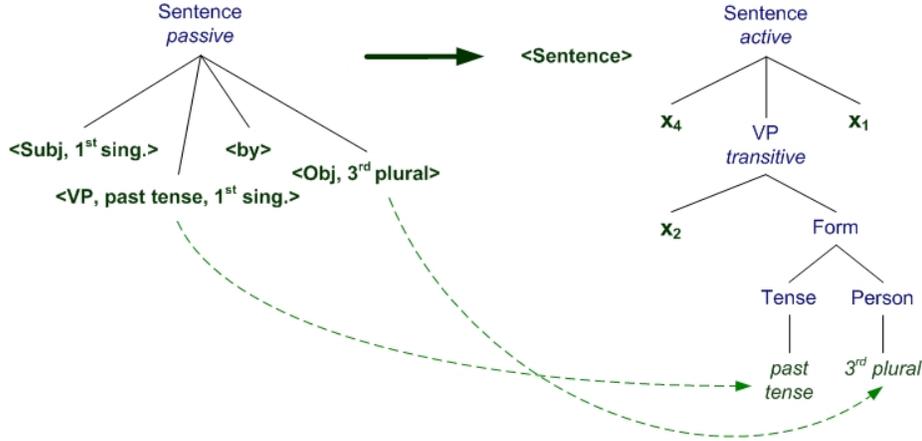


Figure 4.2: A transition to transform a root node of a syntax tree of a sentence in passive voice, depending on the states reached by the subtrees (dashed arrows); x_1 , x_2 , and x_4 mark where to insert the output of the 1st, 2nd and 4th subtree.

We only consider *deterministic* bottom-up tree transducers in this thesis. Thus, we use the term *bottom-up tree transducer* synonymously. Note that the right-hand side $q(z)$ is a tree over the ranked alphabet $\Delta \cup Q \cup \mathcal{X}$ where each state in Q has rank 1, and each variable in \mathcal{X} has rank 0. Similarly, the left-hand side $\mathbf{a}(q_1, \dots, q_m)$ can be considered as an abbreviation for the tree $\mathbf{a}(q_1(x_1), \dots, q_m(x_m))$ over $\Sigma \cup Q \cup \mathcal{X}$.

For every input symbol \mathbf{a} of rank m and sequence of states $q_1 \dots q_m$ of Q , let $R(\mathbf{a}, q_1 \dots q_m)$ be the right-hand side of the transition for \mathbf{a} and $q_1 \dots q_m$ if it is defined, and let $R(\mathbf{a}, q_1 \dots q_m)$ be undefined otherwise.

Example 4.1. In Figure 4.2, a transition of a DBTT for the transformation from sentences in passive voice into sentences in active voice is pictured. It is the transformation with the left-hand side

$$\text{Sentence passive}(\langle \text{Subj, 1}^{st} \text{ sing.} \rangle, \langle \text{VP, past tense, 1}^{st} \text{ sing.} \rangle, \langle \text{by} \rangle, \langle \text{Obj, 3}^{rd} \text{ plural} \rangle).$$

The dashed arrows illustrate how the right-hand side depends on the states on the left-hand side of the transition. It can be applied to the root of the tree in Figure 1.1(a), when the four subtrees are transformed as in Figure 4.1. It would reach state $\langle \text{Sentence} \rangle$ and produces an output tree with root label **Sentence active**. The transducer also has a transition with the left-hand side

$$\text{Sentence passive}(\langle \text{Subj, 1}^{st} \text{ sing.} \rangle, \langle \text{VP, past tense, 1}^{st} \text{ sing.} \rangle, \langle \text{by} \rangle, \langle \text{Obj, 2}^{nd} \text{ sing.} \rangle)$$

where the person of the object is 2nd singular. In this case, the node 2.2.2.1 of the right-hand side is labeled with 2nd sing., instead of 3rd plural. It does not

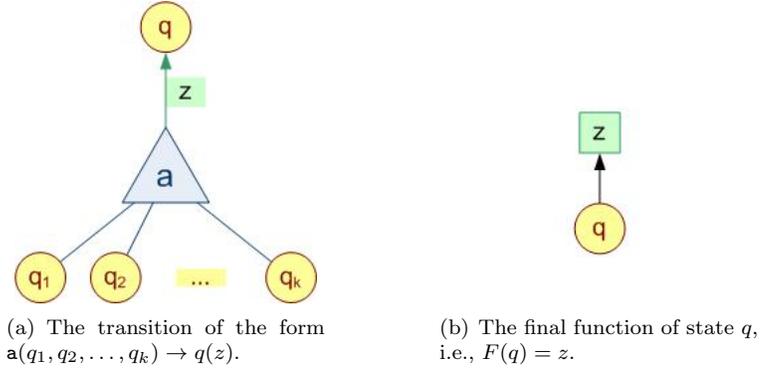


Figure 4.3: Sketch of a transition and a final function of a DBTT.

contain a transition with left-hand side

Sentence passive(\langle Subj, 1^{st} sing. \rangle , \langle VP, past tense, 3^{rd} plural \rangle ,
 \langle by \rangle , \langle Obj, 2^{nd} sing. \rangle)

where the persons of the subject and of the predicate differ. Because this does not occur in proper sentences. \triangleleft

The *size* of a DBTT T , denoted by $|T|$, is the sum of the sizes (= number of symbols) of its final outputs and the sizes of the left-hand sides and right-hand sides of its transitions.

To illustrate DBTTs, we design states by circles, present output trees by squares, and input labels by triangles. In Figure 4.3(a), a transition of the form $a(q_1, q_2, \dots, q_k) \rightarrow q(z)$ is drawn. The order of the states on the left-hand side is given by the sorting of the incoming edges at the bottom of the triangle representing the label (not the position of the states). The output is located at the arrow connecting the tip of the triangle with the state of the right-hand side of the transition. The final function is presented in Figure 4.3(b). In general, we draw transducers or parts of them only with labels of rank less than 2 (Figure 4.4).

Example 4.2. Consider a smaller example. The following transformation checks if the input tree contains an even or an odd number of **b**-labeled nodes. Depending on that, the root of the output tree is labeled **Even** or **Odd**, respectively. In the second case, the input tree is copied where lower case letters are transformed into capital letters.

$$\tau : \mathbf{b}^n(\mathbf{e}) \mapsto \begin{cases} \mathbf{Even} & n \in \mathbb{N}_0 \text{ even} \\ \mathbf{Odd}(\mathbf{B}^n(\mathbf{E})) & n \in \mathbb{N}_0 \text{ odd} \end{cases}$$

It maps trees of \mathcal{T}_Σ to trees of \mathcal{T}_Δ where the input alphabet is $\Sigma = \{\mathbf{b}^{(1)}, \mathbf{e}^{(0)}\}$ and the output alphabet is given by $\Delta = \{\mathbf{Odd}^{(1)}, \mathbf{B}^{(1)}, \mathbf{E}^{(0)}, \mathbf{Even}^{(0)}\}$, respectively. Both input and output trees are monadic. A DBTT defining this transformation is $T = (Q, \Sigma, \Delta, R, F)$ (Figure 4.4) with set of states $Q = \{q_{\text{even}}, q_{\text{odd}}\}$, the transition function R is given by

$$\begin{aligned} \mathbf{e} &\rightarrow q_{\text{even}}(\mathbf{E}) \\ \mathbf{b}(q_{\text{even}}) &\rightarrow q_{\text{odd}}(\mathbf{B}(x_1)) \\ \mathbf{b}(q_{\text{odd}}) &\rightarrow q_{\text{even}}(\mathbf{B}(x_1)). \end{aligned}$$

The final function F is defined as follows:

$$F(q_{\text{even}}) = \mathbf{Even} \quad \text{and} \quad F(q_{\text{odd}}) = \mathbf{Odd}(y)$$

The transducer memorizes in the state if it has processed an even or an odd number of \mathbf{b} -labeled nodes. In the end, it possibly needs the copy of the input tree. Therefore, it produces for every node the corresponding output node. \triangleleft

Let $T = (Q, \Sigma, \Delta, R, F)$ be a DBTT. Assume that $t \in \mathcal{T}_\Sigma(y)$ and $q \in Q$. The *result* $\llbracket t \rrbracket_q^T$ of a computation of T on input t when starting in state q at variable leaves y is defined by induction on the structure of t :

$$\begin{aligned} \llbracket y \rrbracket_q^T &= q(y) \\ \llbracket \mathbf{a}(t_1, \dots, t_m) \rrbracket_q^T &= q'(z[z_1, \dots, z_k]) \\ &\quad \text{if } \forall i \llbracket t_i \rrbracket_q^T = q_i(z_i) \text{ and } R(\mathbf{a}, q_1 \dots q_m) = q'(z) \end{aligned}$$

If $\llbracket t \rrbracket_q^T = q'(z')$, then z' is also called the *output* produced for t at q . Note that the function $\llbracket \cdot \rrbracket_q^T$ may not be defined for all trees t . The superscript T can be omitted if T is clear from the context. If $t \in \mathcal{T}_\Sigma$ we also omit the subscript q , i.e., we write $\llbracket t \rrbracket^T$ for $\llbracket t \rrbracket_q^T$. In this case if $\llbracket t \rrbracket^T = q'(z)$, we say that q' is *reached by* t . For a ground tree $t \in \mathcal{T}_\Sigma$, which reaches some state q' , we refer to the output by $\text{out}^T(t)$, i.e., $\text{out}^T(t) = z'$ iff $\llbracket t \rrbracket^T = q'(z')$.

The *image* $\tau_q^T(t)$ of the tree t is then defined by $\tau_q^T(t) = z' \cdot z$ iff $\llbracket t \rrbracket_q^T = q'(z)$ for some state q' with $F(q') = z'$. Again, we omit the subscript q if the tree t does not contain the variable y . The *transformation described by* DBTT T is the partial function

$$\tau^T : \mathcal{T}_\Sigma \dashrightarrow \mathcal{T}_\Delta.$$

Example 4.3. Consider again the transformation translating passive into active voice. In Figure 4.1 the results of the subtrees of the root of the input tree in Figure 1.1(a) are given. Since the four states correspond to the states in the transition given in Figure 4.2, this transition can be applied to the root of the input tree. By the definition, we get the result of the whole tree, by replacing

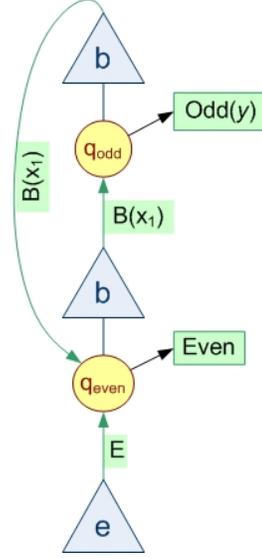


Figure 4.4: The DBTT T of Example 4.2.

the variables x_1, x_2 and x_4 by the corresponding outputs of the first, the second, and the forth subtree, respectively. The result is the state (Sentence) together with the output tree in Figure 1.1(b). \triangleleft

We say that two transducers T and T' are *equivalent* if they describe the same transformation, i.e., for all $u \in \mathcal{T}_\Sigma$, $\tau^T(u)$ is defined iff $\tau^{T'}(u)$ is defined and if they are defined $\tau^T(u)$ and $\tau^{T'}(u)$ are equal. We also use the following notations: The *language* $\mathcal{L}^T(q)$ of a state q is the set of all ground input trees by which q is reached, i.e.,

$$\mathcal{L}^T(q) = \{u \mid \exists v \in \mathcal{T}_\Delta : \llbracket u \rrbracket^T = q(v)\}.$$

A context $c \in \mathcal{C}_\Sigma$ (cf. definition on Page 21) is a *context* of a state q if $\tau_q^T(c)$ is defined. The set of all contexts of state q is denoted by $\mathcal{C}^T(q)$:

$$\mathcal{C}^T(q) = \{c \mid \exists z \in \mathcal{T}_\Delta(y) : \tau_q^T(c) = z\}$$

Example 4.4. The tree \mathbf{e} is an input tree of the transducer T in Example 4.2. The result of this tree is $\llbracket \mathbf{e} \rrbracket^T = q_{\text{even}}(\mathbf{E})$ and its image is $\tau^T(\mathbf{e}) = \mathbf{Even}$ because $F(q_{\text{even}}) = \mathbf{Even}$. Note that the produced output $\text{out}^T(\mathbf{e}) = \mathbf{E}$ disappears. However, the output is used to compute the result and output of other trees containing \mathbf{e} as subtree. Consider the tree $\mathbf{b}(\mathbf{e})$ here, for example. With the transition $\mathbf{b}(q_{\text{even}}) \rightarrow q_{\text{odd}}(\mathbf{B}(x_1))$ and the result of \mathbf{e} we get

$$\llbracket \mathbf{b}(\mathbf{e}) \rrbracket^T = q_{\text{odd}}(\mathbf{B}(x_1)[\mathbf{E}/x_1]) = q_{\text{odd}}(\mathbf{B}(\mathbf{E})).$$

For this tree, the output is also used in the image $\tau^T(\mathbf{b}(\mathbf{e}))$. The final function with $F(q_{\text{odd}}) = \mathbf{Odd}(y)$ leads to

$$\tau^T(\mathbf{b}(\mathbf{e})) = \mathbf{Odd}(y) \cdot \mathbf{B}(\mathbf{E}) = \mathbf{Odd}(\mathbf{B}(\mathbf{E})).$$

In the end, we get that T describes the transformation τ of Example 4.2, i.e., $\tau = \tau^T$. The languages of the states of T are

$$\begin{aligned} \mathcal{L}^T(q_{\text{even}}) &= \{\mathbf{b}^n(\mathbf{e}) \mid n \in \mathbb{N}_0 \text{ is even}\}, \\ \mathcal{L}^T(q_{\text{odd}}) &= \{\mathbf{b}^n(\mathbf{e}) \mid n \in \mathbb{N}_0 \text{ is odd}\}. \end{aligned}$$

Now consider the context $\mathbf{b}(\mathbf{b}(y))$. The result of this context starting at q_{odd} is deduced as follows

$$\begin{aligned} \llbracket y \rrbracket_{q_{\text{odd}}}^T &= q_{\text{odd}}(y) \\ \llbracket \mathbf{b}(y) \rrbracket_{q_{\text{odd}}}^T &= q_{\text{even}}(\mathbf{B}(x_1)[y]) = q_{\text{even}}(\mathbf{B}(y)) \quad \text{with } \mathbf{b}(q_{\text{odd}}) \rightarrow q_{\text{even}}(\mathbf{B}(x_1)) \\ \llbracket \mathbf{b}(\mathbf{b}(y)) \rrbracket_{q_{\text{odd}}}^T &= q_{\text{odd}}(\mathbf{B}(x_1)[\mathbf{B}(y)]) = q_{\text{odd}}(\mathbf{B}(\mathbf{B}(y))) \quad \text{with } \mathbf{b}(q_{\text{even}}) \rightarrow q_{\text{odd}}(\mathbf{B}(x_1)) \end{aligned}$$

Thus, the image of this context starting at q_{odd} is $\tau_{q_{\text{odd}}}^T(\mathbf{b}(\mathbf{b}(y))) = \mathbf{Odd}(\mathbf{B}(\mathbf{B}(y)))$. The context is element of $\mathcal{C}^T(q_{\text{odd}})$, but we also get that $\mathbf{b}(\mathbf{b}(y)) \in \mathcal{C}^T(q_{\text{even}})$. The sets of contexts of the two states are the same:

$$\mathcal{C}^T(q_{\text{even}}) = \{\mathbf{b}^n(y) \mid n \in \mathbb{N}_0\} = \mathcal{C}^T(q_{\text{odd}}) \quad \triangleleft$$

4.1 Trim Transducers

Transducers may contain useless transitions or states that we want to eliminate, while preserving the described transformation. A state q of a DBTT is *reachable* if the language $\mathcal{L}^T(q)$ is non-empty. A state q is *meaningful* if q has at least one context, i.e., $\mathcal{C}_T(q)$ is non-empty. Furthermore, the output at state q is *potentially useful* if there is a context c of q such that the image $\tau_q^T(c)$ contains the variable y . Otherwise, the output at q is called *useless*.

Definition 4.2 (Trim). A bottom-up tree transducer T is called *trim* if T has the following properties:

1. Every state is reachable.
2. Every state is meaningful.
3. If the output at a state q is useless, then for each transition of the form $\mathbf{a}(q_1, \dots, q_m) \rightarrow q(z)$ leading into state q , $z = *$.

In this definition, $*$ is a special output symbol, which does not occur in any image produced by T . It is easy to show that each deterministic bottom-up tree transducer is equivalent to a trim DBTT.

Proposition 4.1. *For every deterministic bottom-up tree transducer T a deterministic bottom-up tree transducer T' can be constructed in polynomial time with the following properties:*

1. T' is equivalent to T
2. $|T'| \leq |T|$
3. T' is trim.

Proof. For convenience, we recall the construction. Let $T = (Q, \Sigma, \Delta, R, F)$ be a DBTT. First, we compute the set **Reach** of reachable states of T as the least solution of the following inequation over the complete lattice of subsets of Q :

$$\text{Reach} \supseteq \{q \in Q \mid \exists m \in \mathbb{N}_0, \mathbf{a} \in \Sigma^{(m)}, q_1, \dots, q_m \in \text{Reach}, z \in \mathcal{T}_\Delta(\mathcal{X}_m) : R(\mathbf{a}, q_1 \dots q_m) = q(z)\}$$

Since the set Q is finite, the least solution can be computed by standard fixpoint iteration in polynomial time. Analogously, we define the set **Mean** of reachable and meaningful states of T as the least solution of the following inequations over the complete lattice of subsets of **Reach**:

$$\begin{aligned} \text{Mean} &\supseteq \{q \in \text{Reach} \mid F(q) \text{ is defined}\} \\ \text{Mean} &\supseteq \{q_1, \dots, q_m \in \text{Reach} \mid \exists m \in \mathbb{N}, \mathbf{a} \in \Sigma^{(m)}, q \in \text{Mean}, z \in \mathcal{T}_\Delta(\mathcal{X}_m) : R(\mathbf{a}, q_1 \dots q_m) = q(z)\} \end{aligned}$$

Again, since the set **Reach** is finite, the least solution can be computed by standard fixpoint iteration in polynomial time.

Additionally, we define the set $\text{Use} \subseteq \text{Mean}$ of reachable and meaningful states of T for which output is potentially useful as the least solution of the following inequations:

$$\begin{aligned} \text{Use} &\supseteq \{q \in \text{Mean} \mid F(q) \in \hat{\mathcal{T}}_{\Delta}(y)\} \\ \text{Use} &\supseteq \{q_i \in \text{Mean} \mid \exists m, i \in \mathbb{N}, \mathbf{a} \in \Sigma^{(m)}, q_1, \dots, q_m \in \text{Mean}, q \in \text{Use}, \\ &\quad z \in \mathcal{T}_{\Delta}(\mathcal{X}_m) : R(\mathbf{a}, q_1 \dots q_m) = q(z) \text{ and } x_i \text{ occurs in } z\} \end{aligned}$$

The DBTT $T' = (Q', \Sigma, \Delta, R', F')$ to be then constructed is given by:

- $Q' = \text{Mean}$
- $F' = F|_{Q'}$
- $R'(\mathbf{a}, q_1 \dots q_m) = q(s')$ iff $R(\mathbf{a}, q_1 \dots q_m) = q(s)$ and $q_1, \dots, q_m, q \in Q'$ where $s' = *$ if $q \notin \text{Use}$ and $s' = s$ otherwise.

By construction, $|T'| \leq |T|$. The proofs that T' is trim and equivalent to T are omitted. \square

In the remainder of this part we consider trim transducers only. For a trim transducer T with set Q of states, we denote by Q_* the set of states with useless output, i.e., for which the output is always $*$.

The transducer of Example 4.2 illustrated in Figure 4.4 is already trim. Both states are reachable and meaningful (cf. Example 4.4). Since every possible subtree is a subtree of an input tree with an odd number of \mathbf{b} -labeled nodes, its output is potentially useful.

Chapter 5

Partition

A deterministic bottom-up tree transducer describes a partial tree transformation. In general, this transformation is not subtree-closed, i.e., there may be a tree, which is not in the domain of the transformation τ^T , but it is a subtree of some other tree in the domain of τ^T . However, the transducer produces some output for such a subtree (but no image). For instance, a DBTT translating sentences in passive voice into their counterparts in active voice only accepts input trees with **Sentence passive** as root label. However, as Figure 3.1 illustrates, such a transducer produces some output for a subtree with root label **Object**. In Figure 4.1 some more output trees of input subtrees are given. Since a deterministic bottom-up tree transducer produces exactly one output tree for every input subtree, we get a function, which produces an output for every subtree of input trees. To get a characterization of the whole transformation, it remains to map the context of the subtree, too (Figure 3.1). We will show that this mapping is also implicitly given by a DBTT (Section 5.1). Motivated by these induced functions of DBTTs, we consider general tree functions and define *partitions* of such functions. A partition consists of a subtree-closed function and mappings of contexts. In the next section, we will show that every DBTT induces such a partition in a natural way. Moreover, every tree function can be defined by a partition.

In this chapter, we consider such function partitions independent of transducers. A partition of a function τ consists of a *subtree function* defined on all subtrees of input trees of τ and a *residual function*, which comprises the remaining information about possible context of a subtree. By that, we get a function on subtrees of input trees without losing information about the whole function. Such a pair of a subtree function and a residual function has to fulfill some properties to describe a tree transformation. The two partial functions of a partition have the same subtree-closed domain. A residual function maps trees to context functions. These context functions are partial functions from contexts to trees with arbitrarily many occurrences of the variable y . For instance, a context function φ maps, amongst others, the context of the **Object**-subtree in Figure 3.1 to the tree **Sentence active**(y, v_{VP}, v_{Obj}) on the right-hand side of this

figure (where we abbreviated the second and the third subtree by v_{VP} and v_{Obj} , respectively). The residual function then maps the input subtree $\text{Object}(\dots)$ to this context function φ .

Formally, a *context function* φ is a partial function from contexts (over Σ) to trees (over Δ), possibly containing the variable y : $\varphi : \mathcal{C}_\Sigma(y) \dashrightarrow \mathcal{T}_\Delta(y)$. A *residual function* r is, in addition, a partial function, which maps trees to context functions:

$$r : \mathcal{T}_\Sigma \dashrightarrow \mathcal{C}_\Sigma(y) \dashrightarrow \mathcal{T}_\Delta(y) .$$

The context function $r(u)$ of a tree $u \in \mathcal{T}_\Sigma$ is called *residual* of u . A residual function defines an equivalence relation.

Definition 5.1 (*r-Equivalence*). Assume r is a residual function. Two trees $u, u' \in \text{dom}(r)$ are *r-equivalent* if they have the same residual:

$$u \equiv_r u' \quad \text{iff} \quad r(u) = r(u') .$$

For a tree $u \in \text{dom}(r)$, we write $[u]_r$ for the equivalence class of u w.r.t. \equiv_r , i.e., $[u]_r = \{u' \mid u \equiv_r u', u' \in \text{dom}(r)\}$.

With these definitions, we get the definition of a function partition, i.e., a pair of a subtree function and a residual function:

Definition 5.2 (*Partition*). Let (σ, r) be a pair of partial functions

$$\sigma : \mathcal{T}_\Sigma \dashrightarrow \mathcal{T}_\Delta \quad \text{and} \quad r : \mathcal{T}_\Sigma \dashrightarrow \mathcal{C}_\Sigma(y) \dashrightarrow \mathcal{T}_\Delta(y) .$$

Then the pair (σ, r) is a *function partition* (or *partition*) if

- the functions have the same domain, i.e., $\text{dom}(\sigma) = \text{dom}(r)$,
- the domain is subtree-closed, and
- the pair is *consistent*, i.e., for all trees $u, u' \in \text{dom}(\sigma)$ and all contexts $c, c' \in \mathcal{C}_\Sigma(y)$ with $c \cdot u = c' \cdot u'$ holds: If $c \in \text{dom}(r(u))$ then $c' \in \text{dom}(r(u'))$ and

$$r(u)(c) \cdot \sigma(u) = r(u')(c') \cdot \sigma(u') . \tag{5.1}$$

The *domain* of a partition (σ, r) is the domain of σ and r . The *tree function defined by the partition* (σ, r) is the partial function $\tau : \mathcal{T}_\Sigma \dashrightarrow \mathcal{T}_\Delta$ given by

$$\tau(u) = \begin{cases} r(u)(y) \cdot \sigma(u) & u \in \text{dom}(\sigma), y \in \text{dom}(r(u)) \\ \text{undefined} & \text{otherwise.} \end{cases}$$

On the other hand, we say that (σ, r) is a *partition of the tree function* τ and σ is a *subtree function* of τ . Two partitions are *equivalent* if they define the same tree function. The *index of a partition* is the index of the *r*-equivalence, i.e., the number of equivalence classes. By Equation 5.1, we get a more general property of the tree function τ of a partition (σ, r) :

Lemma 5.1. *Let (σ, r) be a partition of tree transformation τ . For a context $c \in \mathcal{C}_\Sigma(y)$ and a tree $u \in \mathcal{T}_\Sigma$ with $c \cdot u \in \text{dom}(\tau)$ holds*

$$\tau(c \cdot u) = r(u)(c) \cdot \sigma(u) .$$

Proof. Let $c \in \mathcal{C}_\Sigma(y)$, $u \in \mathcal{T}_\Sigma$, and $c \cdot u \in \text{dom}(\tau)$. By definition of τ we get

$$\tau(c \cdot u) = r(c \cdot u)(y) \cdot \sigma(c \cdot u) .$$

Furthermore, since the partition is consistent and $y \cdot c \cdot u = c \cdot u$:

$$r(c \cdot u)(y) \cdot \sigma(c \cdot u) = r(u)(c) \cdot \sigma(u)$$

Thus, it follows $\tau(c \cdot u) = r(u)(c) \cdot \sigma(u)$. \square

In general, there are more than one partition for a tree function τ , and for a subtree function σ there are different residual functions r such that (σ, r) is a partition of τ . In particular, for every tree function τ , the *canonical partition* (σ_*, r_*) exists:

$$\sigma_*(u) = \begin{cases} * & \text{if } u \in \text{Subtrees}(\text{dom}(\tau)) \\ \text{undefined} & \text{otherwise} \end{cases}$$

$$r_*(u)(c) = \begin{cases} \tau(c \cdot u) & \text{if } c \cdot u \in \text{dom}(\tau) \\ \text{undefined} & \text{otherwise} \end{cases}$$

Here, $*$ is an arbitrary symbol of the output alphabet. Since the image of every residual is in \mathcal{T}_Δ (the trees do not contain the variable y), the output of σ_* is never used in τ . In general, the r_* -equivalence of this partition has infinitely many equivalence classes, i.e., its index of \equiv_{r_*} is infinite.

As already said, every deterministic bottom-up tree transducer defines a partition. This will be defined more precisely in the next section. Often, this partition is more significant than the canonical partition.

Example 5.1. Consider again the transducer $T = (Q, \Sigma, \Delta, R, F)$ of Example 4.2, which defines the transformation

$$\tau : \mathbf{b}^n(\mathbf{e}) \mapsto \begin{cases} \text{Odd}(\mathbf{B}^n(\mathbf{E})) & \text{for all odd } n \in \mathbb{N}_0 \\ \text{Even} & \text{for all even } n \in \mathbb{N}_0 \end{cases}$$

At each state, the transducer produces an output tree for every input subtree, which reaches this state. Thus, the states of T describe functions:

$$\begin{aligned} \mathbf{b}^n(\mathbf{e}) &\mapsto \mathbf{B}^n(\mathbf{E}) & \text{if } \mathbf{b}^n(\mathbf{e}) \in \mathcal{L}^T(q_{\text{even}}) \\ \mathbf{b}^n(\mathbf{e}) &\mapsto \mathbf{B}^n(\mathbf{E}) & \text{if } \mathbf{b}^n(\mathbf{e}) \in \mathcal{L}^T(q_{\text{odd}}) \end{aligned}$$

Since the state reached by a tree is unambiguous and the transducer is deterministic, we get the subtree-closed function

$$\text{out}^T : \mathbf{b}^n(\mathbf{e}) \mapsto \mathbf{B}^n(\mathbf{E}) \quad \text{for all } n \in \mathbb{N}_0 .$$

This function, however, differs from the given transformation $\tau^T = \tau$. To get the remaining parts of the output trees (or actually the proper output), we have to consider another function, which is given by every state. It is the context function, which maps every possible context of a state to an image:

$$\begin{aligned} \tau_{q_{\text{even}}}^T : \mathbf{b}^k(y) &\mapsto \begin{cases} \text{Odd}(\mathbf{B}^k(y)) & \text{for all odd } k \in \mathbb{N}_0 \\ \text{Even} & \text{for all even } k \in \mathbb{N}_0 \end{cases} \\ \tau_{q_{\text{odd}}}^T : \mathbf{b}^k(y) &\mapsto \begin{cases} \text{Even} & \text{for all odd } k \in \mathbb{N}_0 \\ \text{Odd}(\mathbf{B}^k(y)) & \text{for all even } k \in \mathbb{N}_0 \end{cases} \end{aligned}$$

Since every subtree of an input tree reaches exactly one state of the DBTT, we can define a function r^T , which maps every input subtree to such a function, i.e., we get

$$r^T(\mathbf{b}^n(\mathbf{e}))(\mathbf{b}^k(y)) = \begin{cases} \text{Odd}(\mathbf{B}^k(y)) & \text{for all } n, k \in \mathbb{N}_0 \text{ with } n+k \text{ is odd.} \\ \text{Even} & \text{for all } n, k \in \mathbb{N}_0 \text{ with } n+k \text{ is even.} \end{cases}$$

For the two functions out^T and r^T , we observe the following properties with respect to the function τ^T :

1. They have the same domain, $\text{dom}(\text{out}^T) = \text{dom}(r^T) = \{\mathbf{b}^n(\mathbf{e}) \mid n \in \mathbb{N}_0\}$, which is subtree-closed and contains exactly the subtrees of $\text{dom}(\tau^T)$.
2. Together, they describe τ^T : For every tree $\mathbf{b}^n(\mathbf{e})$ and every decomposition of the form $\mathbf{b}^{n_1}(y) \cdot \mathbf{b}^{n_2}(\mathbf{e})$, it holds:

$$\tau^T(\mathbf{b}^n(\mathbf{e})) = r^T(\mathbf{b}^{n_2}(\mathbf{e}))(\mathbf{b}^{n_1}(y)) \cdot \text{out}^T(\mathbf{b}^{n_2}(\mathbf{e}))$$

Thus, (out^T, r^T) is a partition of τ^T . ◁

In the following, let (σ, r) be a partition with input alphabet Σ and output alphabet Δ and the tree function of (σ, r) is denoted τ . We designate a special output symbol $*$ $\in \Delta$, which does not occur in any image of τ . Thus, $*$ may only occur in the output of σ .

5.1 Partitions of DBTTs

The pair (out^T, r^T) in Example 5.1 is a partition of the transformation τ^T defined by the DBTT T of Figure 4.4. As for this transducer, we will now define a partition for every DBTT.

Every transducer produces some output for every subtree of an input tree. A transducer T provides the subtree-closed function $\text{out}^T : \mathcal{T}_\Sigma \dashrightarrow \mathcal{T}_\Delta$. Additionally, for the remaining part of every input tree of the transformation, we get an image. It is given by the partial function τ_q^T depending on the state q , which is reached by the subtree. More formally, let u be a subtree of an input tree u' of the transformation τ^T of DBTT T , i.e., there is a context $c \in \mathcal{C}_\Sigma(y)$ with $u' = c \cdot u$. Then the state q reached by u induces the function τ_q^T , which

provides an image $\tau_q^T(c)$ of the context. In general, a transducer T defines a partial function

$$r^T : \mathcal{T}_\Sigma \dashrightarrow \mathcal{C}_\Sigma(y) \dashrightarrow \mathcal{T}_\Delta(y)$$

which assigns an image to every context depending on a subtree. Precisely, for a tree u , which reaches state q and a context $c \in \mathcal{C}^T(q)$, the function r^T is defined by

$$r^T(u)(c) = \tau_q^T(c) .$$

As for the example transducer, we observe for the two functions out^T and r^T of any transducer T the following properties:

Lemma 5.2. *Let T be a DBTT. Then for the functions out^T and r^T hold:*

1. *They have the same subtree-closed domain, which is the set of subtrees of all input trees of T , i.e.,*

$$\text{dom}(\text{out}^T) = \text{dom}(r^T) = \text{Subtrees}(\text{dom}(\tau^T)) .$$

2. *They are consistent and describe τ^T : For every tree $u' \in \text{dom}(\tau^T)$ and every decomposition of the form $u' = c \cdot u$ with $u \in \mathcal{T}_\Sigma$ and $c \in \mathcal{C}_\Sigma(y)$ holds:*

$$\tau^T(u') = r^T(u)(c) \cdot \text{out}^T(u)$$

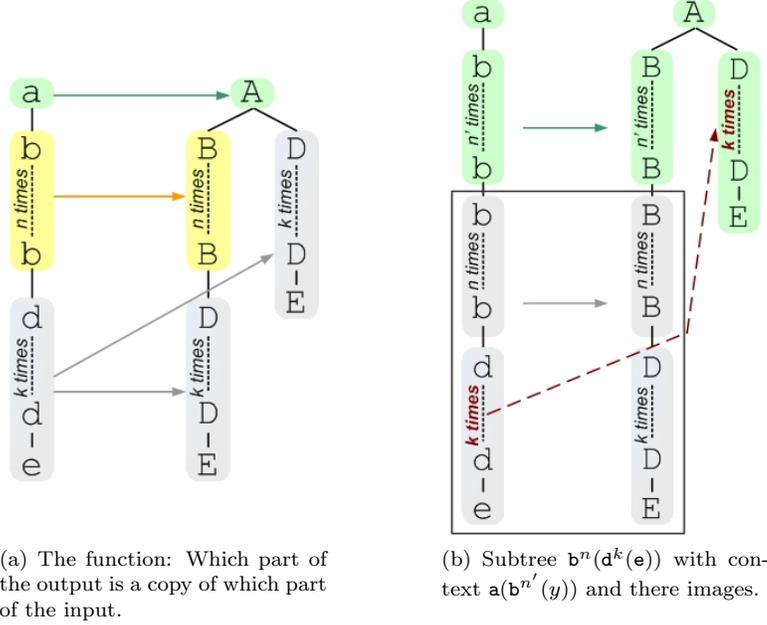
With that, we conclude

Corollary 5.3. *The pair (out^T, r^T) of a transducer T is a partition of τ^T .*

For a DBTT T , we call the pair (out^T, r^T) the *partition of T* . All partitions of DBTTs have some properties in common. In particular, we will observe that their equivalence relations are congruences of finite index and that they are path-finite. The last property describes a restriction of the form of the images of similar contexts. It will be defined in Section 5.1.2. Before that, we consider the index of the equivalence relation \equiv_{r^T} for a DBTT T .

5.1.1 Finite Index

Let T be a DBTT with partition (out^T, r^T) . If we consider two trees, which reach the same state in T , we observe that they have the same residual. Continuing the translation at a state, the transducer does not distinguish with which tree it reaches this state. Thus, two trees, which are in the language of the same state of T are r^T -equivalent. It follows that the number of equivalence classes of \equiv_{r^T} is bounded by the number of states of T . Consider for instance the DBTT T in Example 4.2. All trees with an even number of \mathfrak{b} -labeled nodes reach the state q_{even} . All these trees in $\mathcal{L}^T(q_{\text{even}})$ have the same residual (cf. the partition in Example 5.1). Both states define different residuals. Thus, this partition has index 2. Since a deterministic bottom-up tree transducer has a finite number of states, every partition of a DBTT has finite index.

Figure 5.1: The function τ of Example 5.2.

Lemma 5.4. Let $T = (Q, \Sigma, \Delta, R, F)$ be a DBTT and (out^T, r^T) its partition. Then $|\equiv_{r^T}| \leq |Q|$.

Proof. Let u and u' be two trees, which reach the same state q in T , i.e., $u, u' \in \mathcal{L}^T(q)$. By definition of the residual function r^T , it follows

$$r^T(u) = \tau_q^T = r^T(u').$$

That means $u \equiv_{r^T} u'$. Consequently, $|\equiv_{r^T}| \leq |Q|$. \square

Thus, if we want to get a partition, which defines a DBTT, it is a necessary property that the r -equivalence relation of the partition has finite index. The following example presents a tree transformation, which has no partition with equivalence relation of finite index. Hence, this transformation is not definable by a DBTT.

Example 5.2. Let us consider the tree function (Figure 5.1(a))

$$\tau : a(b^n(d^k(e))) \mapsto A(B^n(D^k(E)), D^k(E)) \quad \text{for all } n, k \geq 0. \quad (5.2)$$

The input tree has two variable parts. Under the root labeled a , the tree has first a chain of b -labeled nodes, followed by a chain of nodes with label d . The leaf has label e . The transformation copies the whole input tree once (by replacing small letters by capital letters). Additionally, the root of the output tree has

a further subtree, which is a copy of the chain of \mathbf{d} 's (Figure 5.1(a)). We can specify the following partition (σ, r) for this transformation (Figure 5.1(b)):

$$\begin{aligned} \sigma : \begin{cases} \mathbf{b}^n(\mathbf{d}^k(\mathbf{e})) \mapsto \mathbf{B}^n(\mathbf{D}^k(\mathbf{E})) & \text{for all } n, k \geq 0 \\ \mathbf{a}(\mathbf{b}^n(\mathbf{d}^k(\mathbf{e}))) \mapsto \mathbf{A}(\mathbf{B}^n(\mathbf{D}^k(\mathbf{E})), \mathbf{D}^k(\mathbf{E})) \end{cases} \\ r(\mathbf{d}^k(\mathbf{e})) : \mathbf{a}(\mathbf{b}^n(\mathbf{d}^l(y))) \mapsto \mathbf{A}(\mathbf{B}^n(\mathbf{D}^l(y)), \mathbf{D}^l(y)) & \text{for all } n, k, l \geq 0 \\ r(\mathbf{b}^n(\mathbf{d}^k(\mathbf{e}))) : \mathbf{a}(\mathbf{b}^m(y)) \mapsto \mathbf{A}(\mathbf{B}^m(y), \mathbf{D}^k(\mathbf{E})) & \text{for all } n, k, m \geq 0 \\ r(\mathbf{a}(\mathbf{b}^n(\mathbf{d}^k(\mathbf{e})))) : y \mapsto y & \text{for all } n, k \geq 0 \end{aligned}$$

However, this partition has infinitely many equivalence classes. For every $k \neq k'$ the trees $\mathbf{b}^n(\mathbf{d}^k(\mathbf{e}))$ and $\mathbf{b}^n(\mathbf{d}^{k'}(\mathbf{e}))$ have different residuals (Figure 5.1(b)).

Unfortunately, one can prove that there is no partition with equivalence relation of finite index that describes this transformation. The idea is the following: Every partition can only produce one output tree for an input tree $\mathbf{b}^n(\mathbf{d}^k(\mathbf{e}))$ (for fixed $n, k \geq 1$). However, for the context $\mathbf{a}(y)$, we need two different subtrees dependent of the input tree. Hence, the residual has to produce one of these subtrees. More precisely, the residual has to produce the part of the left subtree, which contains all n \mathbf{B} -labeled nodes, or the whole right subtree (with all k \mathbf{D} -labeled nodes). For arbitrary n and k , there are infinitely many different such output subtrees, both for the left and the right child of the root. Hence, there are infinitely many different residuals. \triangleleft

This example shows that not every function has a partition with finite index, but we have observed that this is a necessary property to be a function, which is definable by a DBTT. Thus, there is no deterministic bottom-up tree transducer, which defines the tree function τ of Example 5.2.

5.1.2 Path-Finite Partition

Another necessary property is path-finiteness. The idea is the following. A deterministic bottom-up tree transducer applies one rule to every node of an input tree. Thus, for every path p to some node of an input tree, we get a sequence of rules. Although p may belong to infinitely many trees, there are only finitely many different such sequences of rules for p . This property restricts the form of images according to contexts in which p refers to the y node. We will introduce this property as *path-finiteness*, which is a necessary property of a partition to be a partition of a DBTT. Consider for instance the partition illustrated in Figure 3.1. The path $p = (\mathbf{Sentence\ passive}, 4)$ belongs to infinitely many syntax trees of sentences in passive voice. Even though, the images of those contexts with p as y -path only have the y -path $(\mathbf{Sentence\ active}, 1)$.

We give an example, which shows that not every partition is path-finite, nor if its equivalence relation has finite index. Furthermore, path-finiteness is a proper restriction in reference to partitions with equivalence relations of finite index. We will give an example for a tree transformation, which is not definable by a path-finite partition with equivalence relation of finite index.

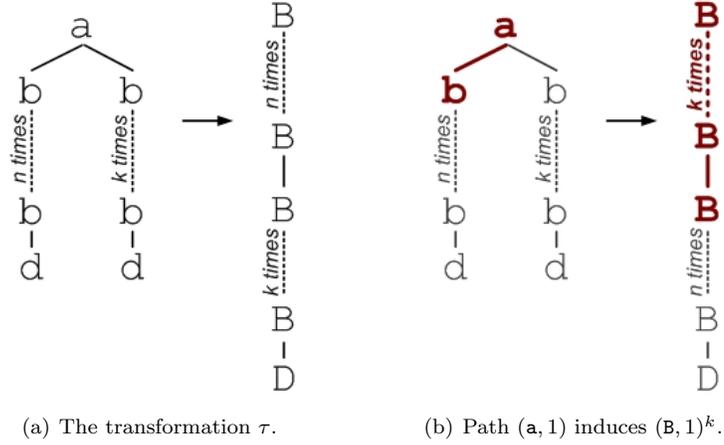


Figure 5.2: The transformation of Example 5.3, which is not path-finite.

Let (σ, r) be an arbitrary partition. For an input tree u of σ and a path p , let $C_{u,p}$ be the set of contexts in the domain of $r(u)$ in which p refers to the y -labeled node. The partition is path-finite if, for every pair (u, p) of an input tree and a path, the number of y -paths in images of $C_{u,p}$ under $r(u)$ is finite.

Definition 5.3 (Path-finite). A partition (σ, r) is called *path-finite* if for all $u \in \text{dom}(\sigma)$ and each path $p \in \text{Paths}_\Sigma$ holds

$$\left| \bigcup \{ \text{Paths}_y(r(u)(c)) \mid c \in \text{dom}(r(u)) \wedge p \in \text{Paths}_y(c) \} \right| < \infty. \quad (5.3)$$

That means, for each y -path in input contexts there are only finitely many paths in corresponding output contexts.

Example 5.3. Consider the tree function

$$\tau : \mathbf{a}(\mathbf{b}^n(\mathbf{d}), \mathbf{b}^k(\mathbf{d})) \mapsto \mathbf{B}^{n+k}(\mathbf{D}) \quad \text{for all } n, k \geq 0 \quad (5.4)$$

which maps trees with two subtrees of depth $n+1$ and $k+1$, respectively, to one tree of depth $n+k+1$ (Figure 5.2(a)). We will show that this transformation has no path-finite partition with finite index (Lemma 5.5). Before that, let us consider some partitions for τ . The canonical partition (σ_*, r_*) (cf. Page 39) is path-finite because the images of contexts in the residuals do not contain any y -paths. However, this partition has infinite index. On the other hand, we can specify the following partition (σ, r) with the domain $\text{Subtrees}(\text{dom}(\tau))$:

$$\sigma = \begin{cases} \mathbf{b}^n(\mathbf{d}) \mapsto \mathbf{B}^n(\mathbf{D}) \\ \mathbf{a}(\mathbf{b}^n(\mathbf{d}), \mathbf{b}^k(\mathbf{d})) \mapsto \mathbf{B}^{n+k}(\mathbf{D}) \end{cases} \quad \text{for all } n, k \geq 0$$

$$r(\mathbf{b}^n(\mathbf{d})) = \begin{cases} \mathbf{a}(\mathbf{b}^m(y), \mathbf{b}^k(\mathbf{d})) \mapsto \mathbf{B}^{m+k}(y) & \text{for all } n, k, m \geq 0 \\ \mathbf{a}(\mathbf{b}^m(\mathbf{d}), \mathbf{b}^k(y)) \mapsto \mathbf{B}^{m+k}(y) & \end{cases}$$

$$r(\mathbf{a}(\mathbf{b}^n(\mathbf{d}), \mathbf{b}^k(\mathbf{d}))) = y \mapsto y . \quad \text{for all } n, k \geq 0$$

In contrast to the canonical partition, this partition has finite index. It has two equivalence classes: The one of all input trees of τ and the one of all proper subtrees, i.e.,

$$\{\mathbf{a}(\mathbf{b}^n(\mathbf{d}), \mathbf{b}^k(\mathbf{d})) \mid n, k \geq 0\} \quad \text{and} \quad \{\mathbf{b}^n(\mathbf{d}) \mid n \geq 0\} .$$

However, the partition (σ, r) is not path-finite: For any tree of the form $\mathbf{b}^n(\mathbf{d})$ and the path $p = (\mathbf{a}, 1)$ we get infinitely many y -paths in the images of all contexts containing p as y -path under $r(\mathbf{b}^n(\mathbf{d}))$ (cf. Figure 5.2(b)). More precisely, let $u_n = \mathbf{b}^n(\mathbf{d})$. Then, we get $C_{u_n, p} = \{\mathbf{a}(y, \mathbf{b}^k(\mathbf{d})) \mid k \geq 0\}$. For the set in Definition 5.3, it follows

$$\begin{aligned} \text{Paths}_y(r(u_n)(C_{u_n, p})) &= \bigcup \{\text{Paths}_y(\mathbf{B}^k(y)) \mid k \geq 0\} \\ &= \{(\mathbf{B}, 1)^k \mid k \geq 0\} , \end{aligned}$$

which is an infinite set. As we prove in the following lemma, for this tree function there is no partition, which combines both properties, i.e., a path-finite partition with an equivalence relation of finite index. \triangleleft

The transformation in the Example 5.3 is not definable by a DBTT. The following lemma shows that every partition for this transformation is not path-finite or has an equivalence of infinite index.

Lemma 5.5. *There is a tree function τ , which is not definable by a path-finite partition with finite index.*

Proof. Consider the function τ of Example 5.3 and let (σ, r) be a partition of τ . We show that if (σ, r) has finitely many r -equivalence classes, the partition is not path-finite. For every subtree $\mathbf{b}^n(\mathbf{d}) \in \text{dom}(\sigma)$ and every context $\mathbf{a}(\mathbf{b}^m(y), \mathbf{b}^k(\mathbf{d})) \in \text{dom}(r(\mathbf{b}^n(\mathbf{d})))$ it holds by Lemma 5.1

$$\begin{aligned} &r(\mathbf{b}^n(\mathbf{d}))(\mathbf{a}(\mathbf{b}^m(y), \mathbf{b}^k(\mathbf{d}))) \cdot \sigma(\mathbf{b}^n(\mathbf{d})) \\ &= \tau(\mathbf{a}(\mathbf{b}^{m+n}, \mathbf{b}^k(\mathbf{d}))) \\ &= \mathbf{B}^{n+m+k}(\mathbf{D}) . \end{aligned}$$

Assume \equiv_r has finite index. Then there are infinitely many input trees of the form $\mathbf{b}^n(\mathbf{d})$ with the same residual. Let $\mathbf{b}^n(\mathbf{d})$ and $\mathbf{b}^m(\mathbf{d})$ be two trees with the same residual (for some $n \neq m$). For every $k \geq 0$, there are two possibilities for the residual:

$$r(\mathbf{b}^n(\mathbf{d}))(\mathbf{a}(y, \mathbf{b}^k(\mathbf{d}))) = \mathbf{B}^{k+n}(\mathbf{D}) \quad (5.5)$$

$$r(\mathbf{b}^m(\mathbf{d}))(\mathbf{a}(y, \mathbf{b}^k(\mathbf{d}))) = \mathbf{B}^{i_k}(y) \quad \text{for some } i_k \geq 0 \quad (5.6)$$

First, assume that for some k the image is ground (Equation 5.5). Then, we get for the output trees

$$\begin{aligned}
\mathbf{B}^{k+m}(\mathbf{D}) &= \tau(\mathbf{a}(\mathbf{b}^m(\mathbf{d}), \mathbf{b}^k(\mathbf{d}))) \\
&= r(\mathbf{b}^m(\mathbf{d}))(\mathbf{a}(y, \mathbf{b}^k(\mathbf{d}))) \cdot \sigma(\mathbf{b}^m(\mathbf{d})) \\
&= r(\mathbf{b}^n(\mathbf{d}))(\mathbf{a}(y, \mathbf{b}^k(\mathbf{d}))) \cdot \sigma(\mathbf{b}^m(\mathbf{d})) \\
&= \mathbf{B}^{k+n}(\mathbf{D}) \cdot \sigma(\mathbf{b}^m(\mathbf{d})) \\
&= \mathbf{B}^{k+n}(\mathbf{D}),
\end{aligned}$$

which is not possible because $n \neq m$. Thus, for every k , the residual is defined by Equation 5.6. With that, it follows

$$\begin{aligned}
\mathbf{B}^{i_k}(y) \cdot \mathbf{B}^{k+m-i_k}(\mathbf{D}) &= \mathbf{B}^{k+m}(\mathbf{D}) \\
&= \tau(\mathbf{a}(\mathbf{b}^m(\mathbf{d}), \mathbf{b}^k(\mathbf{d}))) \\
&= r(\mathbf{b}^m(\mathbf{d}))(\mathbf{a}(y, \mathbf{b}^k(\mathbf{d}))) \cdot \sigma(\mathbf{b}^m(\mathbf{d})) \\
&= r(\mathbf{b}^n(\mathbf{d}))(\mathbf{a}(y, \mathbf{b}^k(\mathbf{d}))) \cdot \sigma(\mathbf{b}^m(\mathbf{d})) \\
&= \mathbf{B}^{i_k}(y) \cdot \sigma(\mathbf{b}^m(\mathbf{d}))
\end{aligned}$$

and with that, $\sigma(\mathbf{b}^m(\mathbf{d})) = \mathbf{B}^{k+m-i_k}(\mathbf{D})$. Since k is arbitrary, but $k+m-i_k$ must be the same for all k , i_k must be variable and dependent of k . In particular, the set $\{i_k \mid k \geq 0\}$ is infinite, but then, every image $r(\mathbf{b}^n(\mathbf{d}))(\mathbf{a}(y, \mathbf{b}^k(\mathbf{d}))) = \mathbf{B}^{i_k}(y)$ has another y -paths. More precisely, $(\mathbf{B}, 1)^{i_k} \in \text{Paths}_y(r(\mathbf{b}^n(\mathbf{d}))(C_{\mathbf{b}^n(\mathbf{a}), (\mathbf{a}, 1)}))$ for all i_k . Thus, (σ, r) is not path-finite. \square

It remains to prove that path-finiteness is a necessary property of a partition to be a partition of a deterministic bottom-up tree transducer, i.e., every partition of a DBTT is path-finite:

Theorem 5.6. *The partition (out^T, r^T) of a deterministic bottom-up tree transducer T is path-finite.*

Proof. The idea is given in the beginning of this section: For a node in an input tree the transducer applies one rule. There are finitely many possibilities for each node depending on its subtrees. Thus, for a finite y -path in the input context, there are finitely many different y -paths in its image.

Let $T = (Q, \Sigma, \Delta, R, F)$, $u \in \text{dom}(r^T)$ and $c_u \in \text{dom}(r^T(u))$. We consider the path $p \in \text{Paths}_y(c_u)$ and show by induction over the length of p that there are only finitely many y -paths in the image of the residual of u such that p is a y -path in the context, i.e., that Equation 5.3 holds.

First, assume $p = \epsilon$. Thus, $c_u = y$ and there is no other context $c' \in \mathcal{C}_\Sigma(y)$ with $p \in \text{Paths}_y(c')$. The image $r^T(u)(c_u)$ is defined by the final function F of T , i.e., $r^T(u)(c_u) = F(q)$ for some state $q \in Q$. Thus, the image has only finitely many different paths. Consequently, the set $\text{Paths}_y(r^T(u)(c_u))$ of all y -paths in the image is finite.

Now assume $p = p_1 \cdot (\mathbf{a}, i)$ for some label $\mathbf{a} \in \Sigma$ of rank k and $i \leq k$. We have to prove that the set

$$\bigcup \{ \text{Paths}_y(r^T(u)(c)) \mid c \in \text{dom}(r^T(u)) \wedge p \in \text{Paths}_y(c) \}$$

is finite. Consider a context $c \in \mathcal{C}_\Sigma(y)$ with $p \in \text{Paths}_y(c)$, then the context c can be written as $c = c_1 \cdot c_a$ with $p_1 \in \text{Paths}_y(c_1)$ and $(\mathbf{a}, i) \in \text{Paths}_y(c_a)$. If $c \in \text{dom}(r^T(u))$, it follows by the definition of r^T (cf. Page 41) that u reaches a state $q_u \in Q$ and $\tau_{q_u}^T = r^T(u)$. Thus, $c \in \mathcal{C}^T(q_u)$ and $r^T(u)(c) = \tau_{q_u}^T(c)$. The transducer processes the context bottom-up. Thus, there is a state q_a and some image z_a with $\llbracket c_a \rrbracket_{q_u}^T = q_a(z_a)$. The remaining context is also translated to some image: $\tau_{q_a}^T(c_1) = z_1$. These two images build the image of the context c , i.e., $\tau_{q_u}^T(c) = z_1 \cdot z_a$. Consequently,

$$r^T(u)(c) = z_1 \cdot z_a.$$

Furthermore, the tree $c_a \cdot u$ reaches state q_a with $c_1 \in \mathcal{C}^T(q_a)$. Thus, we have $c_a \cdot u \in \text{dom}(r^T)$ and $c_1 \in \text{dom}(r^T(c_a \cdot u))$. By induction, the set

$$\bigcup \{ \text{Paths}_y(r^T(c_a \cdot u)(c_1)) \mid c_1 \in \text{dom}(r^T(c_a \cdot u)) \wedge p_1 \in \text{Paths}_y(c_1) \}$$

is finite for $c_a \cdot u$ and p_1 . It remains to prove that there are only finitely many y -paths in images of c_a . Since $(\mathbf{a}, 1)$ is the y -path in c_a , the context has the form

$$c_a = \mathbf{a}(u_1, \dots, u_{i-1}, y, u_{i+1}, \dots, u_k)$$

for some subtrees $u_1, \dots, u_{i-1}, u_{i+1}, \dots, u_k \in \mathcal{T}_\Sigma$. Let $u_i = u$. By the definition of $r^T(u)$ we get: There are $q_1, \dots, q_k \in Q$ and output trees $v_1, \dots, v_k \in \mathcal{T}_\Delta$ with $\llbracket u_j \rrbracket^T = q_j(v_j)$ for all $j \in [k]$. In particular, $q_i = q_u$. In addition, the right-hand side of $(\mathbf{a}, q_1 \dots q_k)$ is $R(\mathbf{a}, q_1 \dots q_k) = q_a(z'_a)$ for some $z'_a \in \mathcal{T}_\Delta(\mathcal{X}_k)$. The image of c_a is then given by $z_a = z'_a[v_1, \dots, v_{i-1}, y, v_{i+1}, \dots, v_k]$. Thus, the y -paths in this image are the x_i -paths in z'_a . However, the number of different images z'_a is bounded by the number of different sequences of states $q_1 \dots q_k$. Since Q is finite, there are only finitely many sequences. For each of these images, the number of x_i -paths is finite. Thus, the set of y -paths

$$\bigcup \{ \text{Paths}_y(z_a) \mid u \in \mathcal{L}^T(q_u) \wedge \llbracket c_a \rrbracket_{q_u}^T = q_a(z_a) \wedge (\mathbf{a}, 1) \in \text{Paths}_y(c_a) \}$$

is finite. Together, we get for $p = p_1 \cdot (\mathbf{a}, 1)$

$$\begin{aligned} & \left| \bigcup \{ \text{Paths}_y(r^T(u)(c)) \mid c \in \text{dom}(r^T(u)) \wedge p \in \text{Paths}_y(c) \} \right| \\ & \leq \left| \bigcup \{ \text{Paths}_y(r^T(c_a \cdot u)(c_1)) \mid c_1 \in \text{dom}(r^T(c_a \cdot u)) \right. \\ & \quad \left. \wedge p_1 \in \text{Paths}_y(c_1) \wedge (\mathbf{a}, 1) \in \text{Paths}_y(c_a) \} \right| \\ & \quad \cdot \left| \bigcup \{ \text{Paths}_y(z_a) \mid u \in \mathcal{L}^T(q_u) \wedge \llbracket c_a \rrbracket_{q_u}^T = q_a(z_a) \wedge (\mathbf{a}, 1) \in \text{Paths}_y(c_a) \} \right| \\ & < \infty. \end{aligned} \quad \square$$

Summarizing, we get by Lemma 5.4 and Theorem 5.6 that for every DBTT T , its partition (out^T, r^T) is path-finite and has finite index.

5.2 Bottom-Up Partition

For every transformation, infinitely many different partitions are possible. First of all, there is the canonical partition (cf. Page 39), which realizes the output in the residual and produces $*$ for every subtree in the subtree function. We can define arbitrarily many partitions for the same transformation by only changing the subtree function of the canonical partition. There are also partitions with different residual functions for one tree function. We figured out that for every transformation, which is definable by a DBTT, there exists a partition (the partition of the transducer), which is path-finite and has an equivalence relation of finite index. However, if a transformation has a path-finite partition with finite index, in general, the transformation has more than one partition with these constraints. For instance, we can separate one input tree and define its output under the subtree function and its residual as in the canonical partition.

In this chapter, we specify a unique partition for every transformation. We are interested only in transformations, which are definable by DBTTs. Thus, if there is no path-finite partition with finite index, we fix the canonical partition as normal form. Otherwise, we present a normal form, which is path-finite and has finite index. This normal form is a unique partition. In Section 6.3, we show that this partition is already the partition of a unique DBTT, describing the given transformation. The normal form has six intrinsic properties:

trim Every subtree produces significant output only if an image of some context uses it, i.e., contains a y . Otherwise, it produces $*$.

proper The subtrees in an equivalence class are mapped to infinitely many different output trees. Otherwise, all trees in this class are mapped to $*$.

earliest The output of such an equivalence class with infinitely many different output trees is produced as early as possible, i.e., there is no common suffix in all images of the residual of a subtree. This common part would be produced by the subtree function.

unified earliest The residual of trees in different equivalence classes cannot be unified by ground trees.

congruent The equivalence relation is a congruence relation.

minimal We consider the partition with the smallest index, which is trim, proper, earliest, unified earliest, and path-finite.

We show that for every path-finite partition with finite index, there exists exactly one partition, which fulfills all these properties and describes the same transformation. We will call this unique partition the *bottom-up partition* of the transformation if it exists (cf. Definition 5.10).

5.2.1 Trim Partition

Similar to trim transducers, we only consider *trim* partitions. Meaning, for an input subtree u , σ outputs a tree $v \neq *$ if and only if it is used in some image of a context of u . Otherwise, it maps u to $*$. Furthermore, every subtree u in $\text{dom}(\sigma)$ has a non-empty residual. Consider, for instance, the transformation of sentences in passive voice into sentences in active voice (Figure 1.1). Figure 4.1 shows parts of a possible subtree function, but if the preposition “by” only occurs in the third subtree of the root node, its output is never used. Changing the output of the subtree **Preposition**(...by...) does not change the transformation defined by the suggested partition. Thus, there are infinitely many different partitions for this transformation. To get a unique partition, we fix the output for such subtrees to $*$.

We say that the output of $u \in \text{dom}(\sigma)$ is *useful* in a partition (σ, r) if a context of u exists, whose image under $r(u)$ contains the variable y , i.e., $r(u)(\mathcal{C}_\Sigma(y)) \cap \hat{\mathcal{T}}_\Delta(y) \neq \emptyset$. Otherwise, the output is called *useless*. A trim partition restricts the useless outputs of the subtree function.

Definition 5.4 (Trim). A partition (σ, r) is called *trim* if for each $u \in \text{dom}(\sigma)$ holds $\text{dom}(r(u)) \neq \emptyset$ and

$$r(u)(\mathcal{C}_\Sigma(y)) \subseteq \mathcal{T}_\Delta \Leftrightarrow \sigma(u) = * .$$

The non-emptiness of the residuals implies that the domain of a trim partition (σ, r) contains exactly the subtrees of input trees of its tree function τ , i.e., $\text{dom}(\sigma) = \mathbf{Subtrees}(\text{dom}(\tau))$. The transformations, which are definable by path-finite partitions with finite index are also definable by trim and path-finite partitions with finite index. If we restrict the domain and replace useless output of the subtree function by $*$, we get an equivalent trim and path-finite partition with the same equivalence relation (restricted to the domain).

Theorem 5.7. *For every path-finite partition (σ, r) with finite index exists an equivalent partition (σ', r') , which is*

- *trim,*
- *path-finite, and*
- *it holds $|\equiv_{r'}| \leq |\equiv_r|$.*

Proof. Let τ be the tree function of the partition (σ, r) . First, let

$$D = \{u \mid u \in \text{dom}(\sigma), \text{dom}(r(u)) \neq \emptyset\} .$$

This set equals the set of subtrees of input trees of the transformation τ , i.e., $D = \mathbf{Subtrees}(\text{dom}(\tau))$. We define a subtree function σ' on D , which differs from σ (restricted to D) only on useless outputs. It is defined for every $u \in D$ as follows:

$$\sigma'(u) = \begin{cases} * & r(u)(\mathcal{C}_\Sigma(y)) \subseteq \mathcal{T}_\Delta \\ \sigma(u) & \text{otherwise} \end{cases}$$

Also r' is restricted to D . But apart from that, it equals r on this domain. Now (σ', r') is trim and it describes the same tree function τ as (σ, r) :

Let $c \cdot u \in \text{dom}(\tau)$. If $r(u)(c) \in \mathcal{T}_\Delta$ is ground then

$$r'(u)(c) \cdot \sigma'(u) = r(u)(c) \cdot \sigma'(u) = r(u)(c) = r(u)(c) \cdot \sigma(u) = \tau(c \cdot u) .$$

Otherwise, $\sigma(u)$ is useful and we get

$$r'(u)(c) \cdot \sigma'(u) = r(u)(c) \cdot \sigma(u) = \tau(c \cdot u) .$$

Since the residuals on the new domain D do not change, the number of equivalence classes is again finite, more precise, it holds $|\equiv_{r'}| \leq |\equiv_r|$. Furthermore, if (σ, r) is path-finite, also (σ', r') is path-finite. \square

This theorem can also be enhanced to partitions with infinite index. The canonical partition is trim (every output of the subtree function is $*$) and path-finite (there are no y -paths in images). If the domain is finite, it has finite index. With this theorem, we assume in the following that every partition is already trim.

5.2.2 Proper Partition

If there exists a path-finite partition with finite index for a transformation, then there may exist more than one. In particular, these different partitions may have the same index. The next example shows a transformation where two partitions (both are trim and path-finite) have the same index and there is no path-finite partition with a smaller index. Thus, these properties do not suffice to get a unique minimal partition.

Example 5.4. Let us consider the quotient ring of integers modulo 3, i.e. $\mathbb{Z}_3 = \{0, 1, 2\}$, with addition $+_3$ and multiplication \times_3 (modulo 3). We use prefix notation. Since the operations are commutative, there are six different (ordered) input pairs. We get the following transformation τ :

$$\begin{array}{ll} +_3(\langle 0, 0 \rangle) \mapsto 0 & \times_3(\langle 0, 0 \rangle) \mapsto 0 \\ +_3(\langle 0, 1 \rangle) \mapsto 1 & \times_3(\langle 0, 1 \rangle) \mapsto 0 \\ +_3(\langle 0, 2 \rangle) \mapsto 2 & \times_3(\langle 0, 2 \rangle) \mapsto 0 \\ +_3(\langle 1, 1 \rangle) \mapsto 2 & \times_3(\langle 1, 1 \rangle) \mapsto 1 \\ +_3(\langle 1, 2 \rangle) \mapsto 0 & \times_3(\langle 1, 2 \rangle) \mapsto 2 \\ +_3(\langle 2, 2 \rangle) \mapsto 1 & \times_3(\langle 2, 2 \rangle) \mapsto 1 \end{array}$$

Seen as a tree transformation, we get input trees with root $+_3$ or \times_3 (of rank 1) and leaves (of rank 0) in the set

$$\Sigma^{(0)} = \{\langle 0, 0 \rangle, \langle 0, 1 \rangle, \langle 0, 2 \rangle, \langle 1, 1 \rangle, \langle 1, 2 \rangle, \langle 2, 2 \rangle\}$$

To minimize the number of equivalence classes, we have to subsume input subtrees such that they have the same residual. One possibility is to subsume all

subtrees with the same result with respect to the operation \times_3 . It is done by the following partition (σ_+, r_+) :

$$\begin{aligned} \sigma_+(p) = \tau(+_3(p)) \quad r_+(p) &= \begin{cases} +_3(y) & \mapsto y \\ \times_3(y) & \mapsto \tau(\times_3(p)) \end{cases} & \text{for all pairs } p \in \Sigma^{(0)} \\ \sigma_+(u) = \tau(u) \quad r_+(u)(y) &= y & \text{for all } u \in \text{dom}(\tau) \end{aligned}$$

The subtree function maps every pair at a leaf to the output if $+_3$ is applied, for instance $\sigma_+(\langle 0, 1 \rangle) = 1$. Thus, the residuals of such pairs map $+_3(y)$ to y and use the pre-produced output. For the context $\times_3(y)$, they have to produce the correct output by the residual, for example $r_+(\langle 0, 1 \rangle)(\times_3(y)) = 0$. Thus, depending on the output of $\tau(\times_3(p))$, there are three different equivalence classes of pairs $p \in \Sigma^{(0)}$. Additionally, all whole input trees have the same residual. They build one further equivalence class.

On the other hand, we can define the partition $(\sigma_\times, r_\times)$ in the same way. It subsumes the pairs with the same result w.r.t. $+_3$ because this result is realized by the residual function. Whereas the output of $\times_3(p)$ is produced at the leaf p :

$$\begin{aligned} \sigma_\times(p) = \tau(\times_3(p)) \quad r_\times(p) &= \begin{cases} +_3(y) & \mapsto \tau(+_3(p)) \\ \times_3(y) & \mapsto y \end{cases} & \text{for all pairs } p \in \Sigma^{(0)} \\ \sigma_\times(u) = \tau(u) \quad r_\times(u)(y) &= y & \text{for all } u \in \text{dom}(\tau) \end{aligned}$$

For the pair $\langle 0, 1 \rangle$, we get $\sigma_\times(\langle 0, 1 \rangle) = 0$ and $r_\times(\langle 0, 1 \rangle)(+_3(y)) = 1$. Again, we get four equivalence classes (depending on the output of $\tau(+_3(p))$).

One can prove that there is no partition for this transformation with less than four equivalence classes. Thereto consider the possible different residuals. If a residual maps some context to y , this determines the output of the corresponding subtree function and influences the images of the other contexts. Thus, for every leaf there are only a certain number of possible residuals. Hence, we do not have any unique partition with minimal index for this transformation. \triangleleft

In this example, we have seen that there is a transformation with different trim and path-finite partitions. We have to define a new property, which restricts the number of partitions further. The example shows that it is not sufficient to take the partition with the minimal index because it is not unique. The problem in this example is how to segment a finite set of input subtrees in equivalence classes. To solve that, we postpone finite sets of output trees to the residual. If there remain no finite sets of output trees of the subtree function for sets of equivalent input trees except for $\{*\}$, then we have a *proper* partition.

Definition 5.5 (Proper). A trim partition (σ, r) is called *proper* if for each $u \in \text{dom}(\sigma)$ holds

$$\sigma([u]_r) = \{*\} \quad \text{or} \quad |\sigma([u]_r)| = \infty .$$

That means, if there is an input tree u with output tree $\sigma(u)$ different to $*$, there are infinitely many output trees of σ , for which the input tree has the same residual $r(u)$. The residual $r(u)$ of a tree u with $|\sigma([u]_r)| = \infty$ then is called *essential*. Otherwise, if $\sigma(u) = *$, the residual $r(u)$ is called *inessential*.

As the property trim, also this property does not limit the set of transformations which are definable by proper partitions with finite index. For every finite set of output trees of r -equivalent input trees of σ , we postpone the output to the residual. The number of equivalence classes of the new partition is linear in the index of \equiv_r .

Theorem 5.8. *For every path-finite partition (σ, r) with finite index, there exists an equivalent partition (σ', r') , which is*

- proper
- path-finite
- and has finite index.

Proof. Let τ be the tree function of (σ, r) . We define a partition (σ', r') on the domain of (σ, r) . The new subtree function is given by:

$$\sigma'(u) = \begin{cases} \sigma(u) & \text{if } |\sigma([u]_r)| = \infty \\ * & \text{otherwise.} \end{cases}$$

In order of the new partition defining the same tree function, we have to change the residual of subtrees u with $|\sigma([u]_r)| < \infty$, and we have to postpone the output of σ .

$$r'(u)(c) = \begin{cases} r(u)(c) & \text{if } |\sigma([u]_r)| = \infty \\ r(u)(c) \cdot \sigma(u) & \text{otherwise.} \end{cases}$$

Now we prove that this pair (σ', r') is a proper and path-finite partition with finite index, which describe τ :

- The pair (σ', r') describes the same tree function τ as (σ, r) and it is consistent: Let $c \cdot u \in \text{dom}(\tau)$. Then $u \in \text{dom}(r)$ and $c \in \text{dom}(r(u))$. If the set of outputs of r -equivalent trees of u is infinite, i.e., $|\sigma([u]_r)| = \infty$, then

$$r'(u)(c) \cdot \sigma'(u) = r(u)(c) \cdot \sigma(u) = \tau(c \cdot u) .$$

Otherwise, if there are only finitely many outputs, $|\sigma([u]_r)| < \infty$, then

$$r'(u)(c) \cdot \sigma'(u) = r(u)(c) \cdot \sigma(u) \cdot * = r(u)(c) \cdot \sigma(u) = \tau(c \cdot u) .$$

Thus, (σ', r') is a partition for the transformation τ .

- The partition is trim: The partition (σ, r) is trim. Thus, for every subtree $u \in \text{dom}(r)$ with $|\sigma([u]_r)| = \infty$, the trim condition

$$r(u)(\mathcal{C}_\Sigma(y)) \subseteq \mathcal{T}_\Delta \Leftrightarrow \sigma(u) = *$$

holds also for (σ', r') on these trees. Otherwise, the output of σ' is $*$ and the range of the residual is in \mathcal{T}_Δ because $\sigma(u) \in \mathcal{T}_\Delta$ and with that also $r'(u)(c) = r(u)(c) \cdot \sigma(u)$ for all $c \in \text{dom}(r'(u))$.

- It is proper: Assume it is not proper, then there is a tree $u \in \text{dom}(r')$ with

$$|\sigma'([u]_{r'})| < \infty \quad \text{and} \quad \sigma'([u]_{r'}) \neq \{*\}.$$

It follows that there must be a r' -equivalent tree $u' \equiv_{r'} u$ with $\sigma'(u') \neq *$. Then, by the construction of σ' , the tree u' gets its output by σ , i.e., $\sigma'(u') = \sigma(u')$, which includes $|\sigma([u']_r)| = \infty$. Additionally, every tree u'' , which is r -equivalent to u' , i.e., $u'' \equiv_r u'$, is in the same infinite equivalence class: $|\sigma([u'']_r)| = \infty$. This implies that for all these trees, σ' and r' are defined as σ and r , respectively, i.e., $\sigma'(u'') = \sigma(u'')$ and $r'(u'') = r(u'')$. Since all these trees are r -equivalent, we get

$$r'(u'') = r(u'') = r(u') = r'(u').$$

With that, the trees are also equivalent w.r.t r' . Thus, the r' -equivalence class of u is infinite:

$$[u]_{r'} = [u']_{r'} \supseteq [u']_r.$$

It causes a contradiction because the set of output trees is infinite, too:

$$|\sigma'([u]_{r'})| = |\sigma'([u']_{r'})| > |\sigma([u']_r)| = \infty.$$

- The partition is path-finite: Since the residual under r' of a tree u equals the residual under r or it is a subset of the set \mathcal{T}_Δ of ground terms. Thus, for the set of y -paths in the images we get for every $c \in \text{dom}(r'(u))$:

$$\text{Paths}_y(r'(u)(c)) \subseteq \text{Paths}_y(r(u)(c))$$

If (σ, r) is path-finite, also (σ', r') is path-finite.

- It remains to proof that the index of the equivalence relation $\equiv_{r'}$ is finite. Consider two r -equivalent trees $u \equiv_r u'$. If $|\sigma([u]_r)| = \infty$ then they have the same residual under r' : $r'(u) = r(u) = r(u') = r'(u')$. Hence, u and u' are equivalent w.r.t. r' , too: $u \equiv_{r'} u'$. For every equivalence class $[u]_r$ with $|\sigma([u]_r)| = \infty$, there is one equivalence class $[u]_{r'}$. Thus, there are only finitely many classes of this form.

Assume that $|\sigma([u]_r)| < \infty$ and $\sigma(u) = \sigma(u')$. For every $c \in \text{dom}(r(u))$, we get

$$r'(u)(c) = r(u)(c) \cdot \sigma(u) = r(u')(c) \cdot \sigma(u') = r'(u')(c)$$

Thus, the two trees u and u' are r' -equivalent, i.e., $u \equiv_{r'} u'$. For every equivalence class $[u]_r$ with $|\sigma([u]_r)| < \infty$, there are at most $|\sigma([u]_r)|$ equivalence classes

$$[u]_{r'} = \{u' \mid \sigma(u) = \sigma(u'), u \equiv_r u'\}.$$

Thus, the number of equivalence classes of $\equiv_{r'}$ is finite. \square

Example 5.5. For the transformation τ of Example 5.4 we get the following proper partition (σ, r) , which is also path-finite and has finite index. This partition produces the output as late as possible because every output set (of equivalent input trees) is finite. Starting with an arbitrary (path-finite and trim) partition of τ , we get the following:

$$\begin{aligned} \sigma(p) = * \quad r(p) &= \begin{cases} +_3(y) & \mapsto \tau(+_3(p)) \\ \times_3(y) & \mapsto \tau(\times_3(p)) \end{cases} & \text{for all pairs } p \in \Sigma^{(0)} \\ \sigma(u) = * \quad r(u)(y) &= \tau(u) & \text{for all } u \in \text{dom}(\tau) \end{aligned}$$

This partition is the canonical partition (σ_*, r_*) (Page 39). It has twelve different equivalence classes, but for this transformation, it is the only partition (for a fix symbol $*$), which is proper. \triangleleft

5.2.3 Earliest Partition

In the previous section, we introduced a partition where finite sets of outputs are postponed to the residual, i.e., trivial outputs are produced as late as possible. If we have infinitely many outputs for an equivalence class, there also may be different partitions with the same index, which fulfill the properties mentioned until here. If we postpone infinitely many outputs to the residuals, we get an infinite index. To preserve the finiteness, we contrarily produce the output as early as possible, i.e., by the subtree function instead of the residual.

Example 5.6. Consider the following transformation (Figure 5.3):

$$\tau : \mathbf{b}^n(\mathbf{e}) \mapsto \mathbf{B}^{n+1}(\mathbf{E}) \quad \text{for all } n \in \mathbb{N}_0$$

This function copies the input tree (only changing the labels to capital letters), and adds an additional \mathbf{B} -labeled node. If we want to define a partition, we have to decide whether we produce this additional node by the subtree function or by the residual. If we add the \mathbf{B} -node in the residual (for every tree), we get the partition (σ, r) with

$$\begin{aligned} \sigma(\mathbf{b}^n(\mathbf{e})) &= \mathbf{B}^n(\mathbf{E}) & \text{for all } n \in \mathbb{N}_0 \\ r(\mathbf{b}^n(\mathbf{e}))(\mathbf{b}^k(y)) &= \mathbf{B}^{k+1}(y) & \text{for all } n, k \in \mathbb{N}_0 \end{aligned}$$

The other way around, we get the partition (σ', r') with

$$\begin{aligned} \sigma'(\mathbf{b}^n(\mathbf{e})) &= \mathbf{B}^{n+1}(\mathbf{E}) & \text{for all } n \in \mathbb{N}_0 \\ r'(\mathbf{b}^n(\mathbf{e}))(\mathbf{b}^k(y)) &= \mathbf{B}^k(y) & \text{for all } n, k \in \mathbb{N}_0 \end{aligned}$$

which produces the additional node in the subtree function (Figure 5.3). Both partitions fulfill all the former conditions, i.e., they are path-finite, proper and have finite index. In particular, the index is 1. They have both the minimal index. \triangleleft

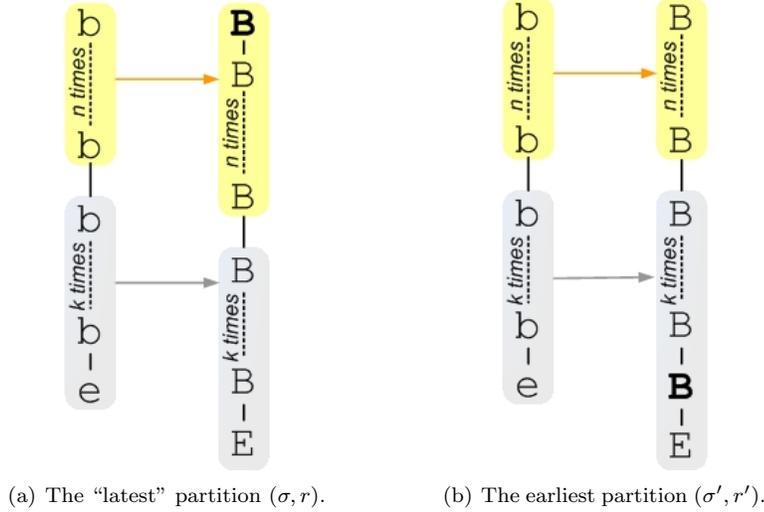


Figure 5.3: The transformation of Example 5.6 with two different partitions.

To get a unique partition in the end, we have to indicate one partition. In the previous section, we have postponed finite sets of outputs at the cost of a minimal index. If we now postpone infinite sets of outputs, we may get an infinite index. Which contradicts the goal of a partition, which is definable by a DBTT. Instead, we will move common parts of images to the subtree function, without increasing the number of equivalence classes. In the given example, we would choose the second partition (σ', r') where the additional B-node is produced as early as possible (Figure 5.3(b)).

Formally, for a context function $\varphi : \mathcal{C}_\Sigma(y) \dashrightarrow \mathcal{T}_\Delta(y)$ let $\hat{\mathcal{Z}}(\varphi)$ denote the set of images $z \in \hat{\mathcal{T}}_\Delta(y)$ in the range of φ . Thus, every tree in $\hat{\mathcal{Z}}(\varphi)$ contains an occurrence of the variable y . Then we define the *greatest common suffix* of φ as the greatest common suffix (Page 19) of all trees in $\hat{\mathcal{Z}}(\varphi)$, i.e.,

$$\text{gcs}(\varphi) = \text{gcs}(\hat{\mathcal{Z}}(\varphi)) = \bigsqcup \hat{\mathcal{Z}}(\varphi)$$

w.r.t. the order \sqsubseteq on $\hat{\mathcal{T}}_\Delta(y)_\perp$ from Section 2.3. Note that the least upper bound for the empty set in this order is \perp , i.e., $\bigsqcup \emptyset = \perp$. Thus, the greatest common suffices of inessential residuals are defined and equal \perp .

Let (σ, r) be a partition of a function τ . The *greatest common suffix* of an input tree $u \in \text{dom}(\sigma)$ with respect to a residual function r is defined as the greatest common suffix of its residual, i.e.,

$$\text{gcs}(u) = \text{gcs}(r(u)).$$

Analogically, we use $\hat{\mathcal{Z}}(u) = \hat{\mathcal{Z}}(r(u))$.

Definition 5.6 (Earliest). A proper partition (σ, r) is called *earliest* if for all trees $u \in \text{dom}(r)$ holds $\text{gcs}(u) \in \{y, \perp\}$.

The greatest common suffix of a tree u in an earliest partition is \perp if all images of contexts under $r(u)$ are ground, i.e., $r(u)$ is inessential. Otherwise, if there exists an image $z = r(u)(c)$ for some context c , then $\text{gcs}(u) = y$. For every subset of $\hat{\mathcal{T}}_\Delta(y)$ the least upper bound w.r.t. \sqsubseteq exists (cf. Section 2.3). Note that the canonical partition is earliest, because all residuals are inessential and, thus, have the greatest common suffix \perp . Assume that for a transformation τ a path-finite partition with finite index exists. We show that for such a transformation τ an earliest partition exists, which, in addition, is also path-finite and has finite index:

Theorem 5.9. *For every proper and path-finite partition (σ, r) with finite index, there exists an equivalent partition (σ', r') , which is*

- earliest,
- path-finite and
- has finite index.

Proof. Let τ be the tree function of (σ, r) . Since $(\hat{\mathcal{T}}_\Delta(y)_\perp, \sqsubseteq)$ is a complete lattice satisfying the ascending chain condition, the greatest common suffix $\text{gcs}(u)$ (w.r.t. r) exists for every tree $u \in \text{dom}(\sigma)$.

We define a partition (σ', r') on the domain of (σ, r) . The new subtree function is given by:

$$\sigma'(u) = \begin{cases} * & \text{if } \text{gcs}(u) = \perp \\ \text{gcs}(u) \cdot \sigma(u) & \text{otherwise.} \end{cases}$$

and the new residual $r'(u)$ of a tree $u \in \text{dom}(r)$ is defined by

$$r'(u)(c) = \begin{cases} r(u)(c) & \text{if } \text{gcs}(u) = \perp \\ z & \text{with } r(u)(c) = z \cdot \text{gcs}(u), \text{ otherwise.} \end{cases}$$

Let $\text{gcs}(u)$ and $\text{gcs}'(u)$ denote the greatest common suffix of trees $u \in \text{dom}(\sigma)$ w.r.t. r and r' , respectively. First, we observe the following:

1. For all $u \in \text{dom}(r)$ holds $\text{gcs}(u) = \perp \Leftrightarrow \text{gcs}'(u) = \perp$. This follows by the definition of (σ', r') : If $\text{gcs}(u) = \perp$, the residual is the same for r and r' . Thus, the greatest common suffix is the same. If $\text{gcs}(u) \neq \perp$, then there is a context c such that its image under $r(u)$ is not ground. But then also $r'(u)(c)$, which is determined by $r(u)(c) = r'(u)(c) \cdot \text{gcs}(u)$, is in $\hat{\mathcal{T}}_\Delta(y)$, too. With that, $\text{gcs}'(u) \neq \perp$.
2. For all $u \in \text{dom}(r)$ holds $\sigma(u) = * \Leftrightarrow \sigma'(u) = *$ because

$$\begin{aligned} \sigma(u) = * & \xLeftrightarrow{\text{trim}} r(u)(\mathcal{C}_\Sigma) \subseteq \mathcal{T}_\Delta \\ & \Leftrightarrow \text{gcs}(u) = \perp \\ & \Leftrightarrow \sigma'(u) = * . \end{aligned}$$

3. For all $u, u' \in \text{dom}(r)$ holds $u \equiv_r u' \implies u \equiv_{r'} u'$. It holds because u and u' have the same residual w.r.t. r if they are r -equivalent. Furthermore, if their $\text{gcs}(r(u))$ is \perp , then their residual under r' is the same as under r . Thus, they are also r' -equivalent. Otherwise, if $\text{gcs}(r(u)) \neq \perp$, the new image for a context c equals for $r'(u)$ and $r'(u')$:

$$r'(u)(c) \cdot \text{gcs}(u) = r(u)(c) = r(u')(c) = r'(u')(c) \cdot \text{gcs}(u') = r'(u')(c) \cdot \text{gcs}(u)$$

With Proposition 2.1 follows $r'(u)(c) = r'(u')(c)$ because $\text{gcs}(u) \in \hat{\mathcal{T}}_\Delta(y)$.

Now we show that (σ', r') is a trim, proper, path-finite and earliest partition with finite index, which describes τ .

- It is consistent and defines τ : Let $c \cdot u \in \text{dom}(\tau)$. Then $u \in \text{dom}(r)$ and $c \in \text{dom}(r(u))$. If the greatest common suffix of u is \perp , i.e., $\text{gcs}(u) = \perp$, then $\hat{\mathcal{Z}}(u) = \emptyset$. That means, all images of contexts of u are ground and since (σ, r) is trim, it follows $\sigma(u) = *$. We get

$$r'(u)(c) \cdot \sigma'(u) = r(u)(c) \cdot * = r(u)(c) \cdot \sigma(u) = \tau(c \cdot u) .$$

Otherwise, if $\text{gcs}(u) \in \hat{\mathcal{T}}_\Delta(y)$, we get

$$r'(u)(c) \cdot \sigma'(u) = r'(u)(c) \cdot \text{gcs}(u) \cdot \sigma(u) = r(u)(c) \cdot \sigma(u) = \tau(c \cdot u) .$$

Consequently, (σ', r') is a partition for the transformation τ .

- The new partition is trim: Let $u \in \text{dom}(\sigma')$. Then it follows:

$$\begin{aligned} \sigma'(u) = * &\Leftrightarrow \sigma(u) = * && \text{Property 2} \\ &\Leftrightarrow r(u)(\mathcal{C}_\Sigma) \subseteq \mathcal{T}_\Delta && (\sigma, r) \text{ is trim} \\ &\Leftrightarrow \text{gcs}(u) = \perp && \text{Def. gcs} \\ &\Leftrightarrow r'(u)(\mathcal{C}_\Sigma) = r(u)(\mathcal{C}_\Sigma) \subseteq \mathcal{T}_\Delta && \text{Def. } r' \end{aligned}$$

- It is proper: Let $u \in \text{dom}(\sigma')$ be a tree with $\sigma'(u) \neq *$. Then it follows with Property 2 that also $\sigma(u) \neq *$. This implies that there are infinitely many different output trees of r -equivalent input trees, i.e., $|\sigma([u]_r)| = \infty$. All trees $u' \equiv_r u$ have the same $\text{gcs}(u') = \text{gcs}(u) \in \hat{\mathcal{T}}_\Delta(y)$. With Property 3 we know that $[u]_r \subseteq [u]_{r'}$. It follows:

$$\sigma'([u]_{r'}) \supseteq \{\text{gcs}(u) \cdot v \mid v \in \sigma([u]_r)\}$$

Thus, $|\sigma'([u]_{r'})| = \infty$

- The partition is earliest: Let $u \in \text{dom}(r)$. If $\text{gcs}(u) = \perp$, with Property 1 it follows that also $\text{gcs}'(u) = \perp$. Otherwise, assume $\text{gcs}(u) \neq \perp$. Then, we know for every context c with $r'(u)(c) \in \hat{\mathcal{T}}_\Delta(y)$ that $r'(u)(c) \sqsubseteq \text{gcs}'(u)$. Furthermore,

$$r(u)(c) = r'(u)(c) \cdot \text{gcs}(u) \sqsubseteq \text{gcs}'(u) \cdot \text{gcs}(u) .$$

Since this holds for every context c with non-ground image $r(u)(c)$, we also get $\text{gcs}(u) \sqsubseteq \text{gcs}'(u) \cdot \text{gcs}(u)$. Thus, $\text{gcs}'(u) = y$.

- It is path-finite: Let $u \in \text{dom}(\sigma')$, $c \in \text{dom}(r'(u))$ and p a y -path in $\text{Paths}_y(r'(u)(c))$. Since p is a y -path in $r'(u)(c)$, this image is in $\hat{\mathcal{T}}_\Delta(y)$. Thus, it fulfills the equation

$$r'(u)(c) \cdot \text{gcs}(u) = r(u)(c) .$$

Also $\text{gcs}(u)$ is in $\hat{\mathcal{T}}_\Delta(y)$ and there is a y -path $p' \in \text{Paths}_y(\text{gcs}(u))$ such that $p \cdot p'$ is a y -path in $r(u)(c)$, i.e., $p \cdot p' \in \text{Paths}_y(r(u)(c))$. With

$$|\text{Paths}_y(r'(u)(\mathcal{C}_\Sigma))| \leq |\text{Paths}_y(r(u)(\mathcal{C}_\Sigma))| ,$$

it follows that (σ', r') is path-finite because (σ, r) is path-finite.

- The equivalence relation $\equiv_{r'}$ has finite index: With Property 3, we get that $\equiv_{r'}$ is a coarsening of \equiv_r . Thus, the index of (σ', r') is less or equal the index of (σ, r) , which is finite. \square

Returning to the example in the beginning of this section, we see that we already have found an earliest partition:

Example 5.7. Let us consider the greatest common suffices of the different residuals in the two partitions of Example 5.6. For the first partition (σ, r) we get the same residual for every input subtree $\mathbf{b}^n(\mathbf{e})$ (for every $n \in \mathbb{N}_0$). The set of images is

$$\hat{\mathcal{Z}}(\mathbf{b}^n(\mathbf{e})) = r(\mathbf{b}^n(\mathbf{e}))(\mathcal{C}_\Sigma) = \{\mathbf{B}^{k+1}(y) \mid k \in \mathbb{N}_0\} .$$

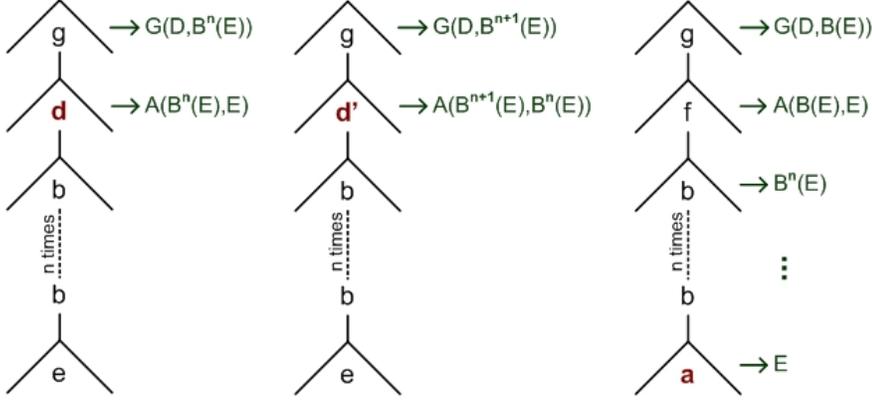
Thus, $\text{gcs}(\mathbf{b}^n(\mathbf{e})) = \bigsqcup \hat{\mathcal{Z}}(\mathbf{b}^n(\mathbf{e})) = \mathbf{B}(y)$. This partition is not earliest, but if we apply the construction of the proof, we get exactly the second partition (σ', r') with

$$\text{gcs}(r'(\mathbf{b}^n(\mathbf{e}))) = \bigsqcup r'(\mathbf{b}^n(\mathbf{e}))(\mathcal{C}_\Sigma) = \bigsqcup \{\mathbf{B}^k(y) \mid k \in \mathbb{N}_0\} = y$$

for every $n \in \mathbb{N}_0$. Hence, this partition is earliest. \triangleleft

5.2.4 Unified Partition

In the Example 5.4, we considered a tree function with different partitions, which are not proper. We had to decide how to subsume the leaves to equivalence classes. The obvious possibilities were either by the result of $+_3$ or by the result of \times_3 . This gave us two different partitions of the same index. We solved that by taking a proper partition. In that case, the involved equivalence classes have finitely many different outputs. Now there is a similar case for equivalence classes with infinitely many outputs. There may be an input subtree u for which we can choose two different output trees v and v' , which cause different essential residuals for u . Nevertheless, we want to define the same transformation. Hence, the residuals only differ somehow in the subtrees v and v' in the images, i.e., the two residuals are unifiable by $\langle v, v' \rangle$ (cf. Page 20). Consider the following example, which presents a partition with two unifiable essential residuals.

Figure 5.4: The function τ_{mgu} of Example 5.8.

The subtree function σ is defined (for all $n \in \mathbb{N}_0$) by:

$$\begin{array}{ll} \mathbf{g}(\mathbf{d}(\mathbf{b}^n(\mathbf{e}))) \mapsto \mathbf{G}(\mathbf{D}, \mathbf{B}^n(\mathbf{E})) & \mathbf{g}(\mathbf{d}'(\mathbf{b}^n(\mathbf{e}))) \mapsto \mathbf{G}(\mathbf{D}, \mathbf{B}^{n+1}(\mathbf{E})) \\ \mathbf{d}(\mathbf{b}^n(\mathbf{e})) \mapsto \mathbf{B}^n(\mathbf{E}) & \mathbf{d}'(\mathbf{b}^n(\mathbf{e})) \mapsto \mathbf{B}^n(\mathbf{E}) \\ \mathbf{b}^n(\mathbf{e}) \mapsto \mathbf{B}^n(\mathbf{E}) & \mathbf{g}(\mathbf{f}(\mathbf{b}^n(\mathbf{a}))) \mapsto * \\ \mathbf{b}^n(\mathbf{a}) \mapsto \mathbf{B}^n(\mathbf{E}) & \mathbf{f}(\mathbf{b}^n(\mathbf{a})) \mapsto \mathbf{E} \end{array}$$

The residual of $\mathbf{b}^n(\mathbf{e})$ (for all $n \in \mathbb{N}_0$) is given by (with $k \in \mathbb{N}_0$):

$$\begin{array}{ll} \mathbf{d}(\mathbf{b}^k(\mathbf{y})) \mapsto \mathbf{A}(\mathbf{B}^k(\mathbf{y}), \mathbf{E}) & \mathbf{d}'(\mathbf{b}^k(\mathbf{y})) \mapsto \mathbf{A}(\mathbf{B}^{k+1}(\mathbf{y}), \mathbf{B}^k(\mathbf{y})) \\ \mathbf{g}(\mathbf{d}(\mathbf{b}^k(\mathbf{y}))) \mapsto \mathbf{G}(\mathbf{D}, \mathbf{B}^k(\mathbf{y})) & \mathbf{g}(\mathbf{d}'(\mathbf{b}^k(\mathbf{y}))) \mapsto \mathbf{G}(\mathbf{D}, \mathbf{B}^{k+1}(\mathbf{y})) . \end{array}$$

For $\mathbf{b}^n(\mathbf{a})$ (for all $n \in \mathbb{N}_0$) we have the residual (with $k \in \mathbb{N}_0$)

$$\mathbf{b}^k(\mathbf{y}) \mapsto \mathbf{B}^k(\mathbf{y}) \quad \mathbf{f}(\mathbf{b}^k(\mathbf{y})) \mapsto \mathbf{A}(\mathbf{B}(\mathbf{E}), \mathbf{E}) \quad \mathbf{g}(\mathbf{f}(\mathbf{b}^k(\mathbf{y}))) \mapsto \mathbf{G}(\mathbf{D}, \mathbf{B}(\mathbf{E})) .$$

For all other input subtrees, we have the following residual for all $n \in \mathbb{N}_0$:

u	$\sigma(u)$	$r(u)$	
		y	$\mathbf{g}(y)$
$\mathbf{g}(\mathbf{d}(\mathbf{b}^n(\mathbf{e})))$	$\mathbf{G}(\mathbf{D}, \mathbf{B}^n(\mathbf{E}))$	y	
$\mathbf{g}(\mathbf{d}'(\mathbf{b}^n(\mathbf{e})))$	$\mathbf{G}(\mathbf{D}, \mathbf{B}^{n+1}(\mathbf{E}))$	y	
$\mathbf{g}(\mathbf{f}(\mathbf{b}^n(\mathbf{a})))$	$*$	$\mathbf{G}(\mathbf{D}, \mathbf{B}(\mathbf{E}))$	
$\mathbf{d}(\mathbf{b}^n(\mathbf{e}))$	$\mathbf{B}^n(\mathbf{E})$	$\mathbf{A}(y, \mathbf{E})$	$\mathbf{G}(\mathbf{D}, y)$
$\mathbf{d}'(\mathbf{b}^n(\mathbf{e}))$	$\mathbf{B}^n(\mathbf{E})$	$\mathbf{A}(\mathbf{B}(y), y)$	$\mathbf{G}(\mathbf{D}, \mathbf{B}(y))$
$\mathbf{f}(\mathbf{b}^n(\mathbf{a}))$	\mathbf{E}	$\mathbf{A}(\mathbf{B}(y), y)$	$\mathbf{G}(\mathbf{D}, \mathbf{B}(y))$

Figure 5.5: The partition (σ, r) for the tree function τ_{mgu} of Example 5.8.

Example 5.8. The Figure 5.4 illustrates the following function τ_{mgu} ($\forall n \in \mathbb{N}_0$):

$$\begin{aligned} \mathbf{g}(\mathbf{d}(\mathbf{b}^n(\mathbf{e}))) &\mapsto \mathbf{G}(\mathbf{D}, \mathbf{B}^n(\mathbf{E})) & \mathbf{g}(\mathbf{f}(\mathbf{b}^n(\mathbf{a}))) &\mapsto \mathbf{G}(\mathbf{D}, \mathbf{B}(\mathbf{E})) \\ \mathbf{g}(\mathbf{d}'(\mathbf{b}^n(\mathbf{e}))) &\mapsto \mathbf{G}(\mathbf{D}, \mathbf{B}^{n+1}(\mathbf{E})) & \mathbf{f}(\mathbf{b}^n(\mathbf{a})) &\mapsto \mathbf{A}(\mathbf{B}(\mathbf{E}), \mathbf{E}) \\ \mathbf{d}(\mathbf{b}^n(\mathbf{e})) &\mapsto \mathbf{A}(\mathbf{B}^n(\mathbf{E}), \mathbf{E}) & \mathbf{b}^n(\mathbf{a}) &\mapsto \mathbf{B}^n(\mathbf{E}) \\ \mathbf{d}'(\mathbf{b}^n(\mathbf{e})) &\mapsto \mathbf{A}(\mathbf{B}^{n+1}(\mathbf{E}), \mathbf{B}^n(\mathbf{E})) \end{aligned}$$

All input trees are monadic. Note that the input trees on the left-hand side differ pairwise in the node with label \mathbf{d} and \mathbf{d}' . On the right-hand side, we have another leaf with label \mathbf{a} ; the \mathbf{d} -labeled node is replaced by a node with label \mathbf{f} . The subtrees of the form $\mathbf{b}^n(\mathbf{e})$ on the left-hand side do not have an image. In contrast to that, all subtrees on the right-hand side have an image.

Now let us define a partition, which is earliest, path-finite and has finite index. We first consider the left-hand side. Let $u_n = \mathbf{b}^n(\mathbf{e})$ for all $n \in \mathbb{N}_0$. For this subtrees u_n (for every $n \in \mathbb{N}_0$), the subtree function has to output $\mathbf{B}^n(\mathbf{E})$. This is necessary because this is the common part of the images of $\mathbf{g}(\mathbf{d}(\mathbf{b}^n(\mathbf{e})))$ and $\mathbf{d}(\mathbf{b}^n(\mathbf{e}))$, which depends on the number n of \mathbf{b} -labeled nodes in u_n . Otherwise, it contradicts one of the mentioned properties. We get the following subtree function σ for the input trees on the left-hand side ($\forall n \in \mathbb{N}_0$):

$$\begin{aligned} \mathbf{g}(\mathbf{d}(u_n)) &\mapsto \mathbf{G}(\mathbf{D}, \mathbf{B}^n(\mathbf{E})) & \mathbf{g}(\mathbf{d}'(u_n)) &\mapsto \mathbf{G}(\mathbf{D}, \mathbf{B}^{n+1}(\mathbf{E})) \\ \mathbf{d}(u_n) &\mapsto \mathbf{B}^n(\mathbf{E}) & \mathbf{d}'(u_n) &\mapsto \mathbf{B}^n(\mathbf{E}) \\ u_n &\mapsto \mathbf{B}^n(\mathbf{E}) \end{aligned}$$

With the restriction of the mentioned properties, the corresponding residual function r is defined as follows. For the input subtrees u_n for every $n \in \mathbb{N}_0$ we get the residual (for all $k \in \mathbb{N}_0$)

$$\begin{aligned} \mathbf{d}(\mathbf{b}^k(y)) &\mapsto \mathbf{A}(\mathbf{B}^k(y), \mathbf{E}) & \mathbf{d}'(\mathbf{b}^k(y)) &\mapsto \mathbf{A}(\mathbf{B}^{k+1}(y), \mathbf{B}^k(y)) \\ \mathbf{g}(\mathbf{d}(\mathbf{b}^k(y))) &\mapsto \mathbf{G}(\mathbf{D}, \mathbf{B}^k(y)) & \mathbf{g}(\mathbf{d}'(\mathbf{b}^k(y))) &\mapsto \mathbf{G}(\mathbf{D}, \mathbf{B}^{k+1}(y)) . \end{aligned}$$

For the other elements in $\text{dom}(\sigma)$ we have for all $n \in \mathbb{N}_0$:

u	$r(u)$	
	y	$\mathbf{g}(y)$
$\mathbf{g}(\mathbf{d}(\mathbf{b}^n(\mathbf{e})))$	y	
$\mathbf{g}(\mathbf{d}'(\mathbf{b}^n(\mathbf{e})))$	y	
$\mathbf{d}(\mathbf{b}^n(\mathbf{e}))$	$\mathbf{A}(y, \mathbf{E})$	$\mathbf{G}(\mathbf{D}, y)$
$\mathbf{d}'(\mathbf{b}^n(\mathbf{e}))$	$\mathbf{A}(\mathbf{B}(y), y)$	$\mathbf{G}(\mathbf{D}, \mathbf{B}(y))$

We get four different residuals and there is no possibility to cut it down to three or less. Now consider the trees on the right-hand side. Again, we have to produce $\mathbf{B}^n(\mathbf{E})$ for the subtrees $\mathbf{b}^n(\mathbf{a})$. The trees $\mathbf{g}(\mathbf{f}(\mathbf{b}^n(\mathbf{a})))$ have only the context y . Thus, we produce $\mathbf{G}(\mathbf{D}, \mathbf{B}(\mathbf{E}))$ by the residual and output $*$ by the subtree function, to get a proper partition.

However, for the tree $\mathbf{f}(\mathbf{b}^n(\mathbf{a}))$, there are different possibilities. First, let us consider the residual if

$$\sigma(\mathbf{f}(\mathbf{b}^n(\mathbf{a}))) = \mathbf{B}(\mathbf{E}) .$$

We get

	$r(u)$	
u	y	$\mathbf{g}(y)$
$\mathbf{f}(\mathbf{b}^n(\mathbf{a}))$	$\mathbf{A}(y, \mathbf{E})$	$\mathbf{G}(\mathbf{D}, y)$

This residual equals the residual of $\mathbf{d}(\mathbf{b}^n(\mathbf{e}))$. Thus, we do not get a new equivalence class. This partition (σ, r) is completely given in Figure 5.5. On the other hand, one can choose

$$\sigma(\mathbf{f}(\mathbf{b}^n(\mathbf{a}))) = \mathbf{E} .$$

Then the residual equals the residual of $\mathbf{d}'(\mathbf{b}^n(\mathbf{e}))$:

	$r(u)$	
u	y	$\mathbf{g}(y)$
$\mathbf{f}(\mathbf{b}^n(\mathbf{a}))$	$\mathbf{A}(\mathbf{B}(y), y)$	$\mathbf{G}(\mathbf{D}, \mathbf{B}(y))$

We have the same number of equivalence classes as before. Both of these partitions fulfill the former properties. \triangleleft

This example shows that the properties path-finite, earliest, and having finite index do not suffice to get a unique partition. The problem for the transformation of the example is that the two residuals of $u = \mathbf{d}(\mathbf{b}(\mathbf{e}))$ and $u' = \mathbf{d}'(\mathbf{b}(\mathbf{e}))$ have a unifier. That means, the pair $\langle \mathbf{B}(\mathbf{E}), \mathbf{E} \rangle$ fulfills the following condition. For all contexts $c \in \text{dom}(r(u)) = \text{dom}(r(u'))$ holds:

$$r(u)(c) \cdot \mathbf{B}(\mathbf{E}) = r(u')(c) \cdot \mathbf{E} \quad (5.7)$$

In Section 2.3, the most-general unifier of two trees is introduced and it is extended to the most-general unifier of a set of pairs of trees (Page 21). Here, we consider the most-general unifier of two residuals. It is the most-general unifier of the set of pairs of images corresponding to the same context.

Formally, for context functions $\varphi, \varphi' : \mathcal{C}_\Sigma(y) \dashrightarrow \mathcal{T}_\Delta(y)$ let

$$C(\varphi, \varphi') = \{(\varphi(c), \varphi'(c)) \mid c \in \text{dom}(\varphi)\}$$

if $\text{dom}(\varphi) = \text{dom}(\varphi')$ and $C(\varphi, \varphi')$ is undefined, otherwise. Then the most-general unifier of two context functions φ and φ' is

$$\text{mgu}(\varphi, \varphi') = \begin{cases} \perp & \text{if } C(\varphi, \varphi') \text{ is undefined} \\ \text{mgu}(C(\varphi, \varphi')) & \text{otherwise.} \end{cases}$$

As for trees, we say that two context functions are *unifiable* if the most-general unifier is not \perp . For example, Equation 5.7 shows that the residuals $r(\mathbf{d}(\mathbf{b}(\mathbf{e})))$ and $r(\mathbf{d}'(\mathbf{b}(\mathbf{e})))$ are unifiable with the unifier $\langle \mathbf{B}(\mathbf{E}), \mathbf{E} \rangle$. This is already the most-general unifier of $r(\mathbf{d}(\mathbf{b}(\mathbf{e})))$ and $r(\mathbf{d}'(\mathbf{b}(\mathbf{e})))$, i.e.,

$$\text{mgu}(r(\mathbf{d}(\mathbf{b}(\mathbf{e}))), r(\mathbf{d}'(\mathbf{b}(\mathbf{e})))) = \langle \mathbf{B}(\mathbf{E}), \mathbf{E} \rangle .$$

Let (σ, r) be a path-finite and earliest partition with finite index. The most-general unifier $\text{mgu}(r(u_1), r(u_2)) = \langle z_1, z_2 \rangle$ for unifiable residuals $r(u_1), r(u_2)$ of input trees $u_1, u_2 \in \text{dom}(r)$ has the following properties:

- If $r(u_i)$ is inessential, then $r(u_i)(\mathcal{C}_\Sigma) \in \mathcal{T}_\Delta$. Therefore, $z_i = \top$.
- Moreover, z_1 contains y iff z_2 contains y . If both z_1 and z_2 contain y , the mgu must equal $\langle y, y \rangle$, otherwise the partition would not be earliest.

If a ground term v is in the set of outputs $\sigma([u]_r)$ of an equivalence class with residual $r(u)$, then v is called *realizable* in $r(u)$. Note that the ground terms z occurring in most-general unifiers of residuals are, however, not necessarily realizable.

Definition 5.7 (Unified Earliest). An earliest partition (σ, r) is called *unified earliest* if no ground term in most-general unifiers of residuals of r is realizable.

Assume for a path-finite and earliest partition with finite index, the most-general unifiers for pairs of residuals are known. Then an equivalent unified earliest partition can be constructed, which is again path-finite and has finite index.

Theorem 5.10. *For each earliest and path-finite partition (σ, r) with finite index, there exists an equivalent partition (σ', r') , which is*

- *unified earliest,*
- *path-finite and*
- *has finite index.*

Proof. Let τ be the tree function of (σ, r) and Z the set of all subtrees of ground most-general unifiers, i.e.,

$$Z = \text{Subtrees}\{z \in \mathcal{T}_\Delta \mid u, u' \in \text{dom}(r), \text{mgu}(r(u), r(u')) = \langle z, z' \rangle \text{ for some } z'\}.$$

Then we define the new partition (σ', r') for the same domain as follows.

$$\sigma'(u) = * \quad \text{and} \quad r'(u)(c) = r(u)(c) \cdot \sigma(u) \quad \text{if } \sigma(u) \in Z \quad (5.8)$$

$$\sigma'(u) = \sigma(u) \quad \text{and} \quad r'(u)(c) = r(u)(c) \quad \text{if } \sigma(u) \notin Z \quad (5.9)$$

Since the new definition of $r'(u)$ only depends on $r(u)$ and $\sigma(u)$, it is a well-defined partition. First, we observe:

1. If the new residual under r' of a tree u do not equals the former residual $r(u)$, then it is inessential (because all $\sigma(u) \in Z$ are ground).
2. The number of new residuals is finite. This follows because \equiv_r has finite index. Hence, the number of different most-general unifiers is finite, too. A new residual $r'(u)$ for a tree u is composed by an old residual $r(u)$ and a ground subterm of a most-general unifier.

Now we have to prove that (σ', r') is unified earliest and path-finite, has finite index, and is equivalent to (σ, r) .

- The partition (σ', r') defines τ : Let $u \in \text{dom}(\sigma)$. Assume the output of u is in Z , i.e., $\sigma(u) \in Z$. Then we get:

$$\tau(c \cdot u) = r(u)(c) \cdot \sigma(u) = r'(u)(c) = r'(u)(c) \cdot * = r'(u)(c) \cdot \sigma'(u)$$

Otherwise, if $\sigma(u) \notin Z$, it is the same as for (σ, r) , i.e.,

$$\tau(c \cdot u) = r(u)(c) \cdot \sigma(u) = r'(u)(c) \cdot \sigma'(u).$$

Thus, the tree function of (σ', r') is τ .

- (σ', r') is trim: If $\sigma'(u) \neq *$ then $r'(u)(\mathcal{C}_\Sigma(y)) = r(u)(\mathcal{C}_\Sigma(y))$, which is no subset of \mathcal{T}_Δ because (σ, r) is trim. If $\sigma'(u) = *$ then

$$r'(u)(\mathcal{C}_\Sigma(y)) = r(u)(\mathcal{C}_\Sigma(y)) \cdot \sigma(u) \subseteq \mathcal{T}_\Delta.$$

- It is proper: This follows directly by Property 2. New outputs under σ' are $*$ only. In addition, the number of outputs that are changed is finite. Thus, the infinite sets of equivalence classes with essential residual are only reduced by a finite number.
- The partition (σ', r') is earliest and path-finite: For essential residuals nothing changes (Property 1). Thus, their greatest common suffices remain y and the number of y -paths is the same as in r . The greatest common suffix of an inessential residual is \perp . The set of y -paths in the images is empty.
- The partition is unified earliest. Assume there exists a tree $u \in \text{dom}(\sigma')$ with $\text{mgu}(r'(u), r'(u')) = \langle \sigma'(u), z \rangle$ for some z . Then $r'(u)$ is not inessential and with Property 1, it equals $r(u)$ and $\sigma'(u) = \sigma(u)$. Assume $r'(u') = r(u')$, then $\langle \sigma'(u), z \rangle$ is also the most-general unifier of $r(u)$ and $r(u')$. However, then $\sigma'(u) = \sigma(u)$ is in Z and $\sigma'(u) = *$ cannot be part of the most-general unifier. Otherwise, if $r'(u') \neq r(u')$, it is inessential and we get for every $c \in \text{dom}(r(u'))$:

$$r(u)(c) \cdot \sigma(u) = r'(u)(c) \cdot \sigma'(u) = r'(u')(c) = r(u')(c) \cdot \sigma(u')$$

Thus, $\text{mgu}(r(u), r(u')) = \langle \sigma(u), \sigma(u') \rangle$ and with that, $\sigma(u)$ is in Z . This contradicts the assumption that $r'(u)$ is essential.

- The equivalence relation $\equiv_{r'}$ has finite index: Since r has finite index and we only get finitely many new residuals (Property 2), the index of r' is finite, too. \square

We have seen on Page 61 that the partition (σ, r) in Example 5.8 is not unified earliest. In the following example, we apply the construction of the proof to that partition.

Example 5.9. Consider again the tree function τ_{mgu} of the Example 5.8. We want to define a unified earliest partition for this function. Continuing along the construction in the proof (Equations 5.8 and 5.9), we start with the partition (σ, r) of Example 5.8. The complete partition (σ, r) , which is not unified earliest, was given in Figure 5.5.

The residual function r maps for arbitrary $n, m \in \mathbb{N}_0$ the trees $u_n = \mathbf{d}(\mathbf{b}^n(\mathbf{e}))$ and $u'_m = \mathbf{d}'(\mathbf{b}^m(\mathbf{e}))$ to unifiable residuals. As seen on Page 61, for those trees holds for every context $c \in \text{dom}(r(u_n))$: $r(u_n)(c) \cdot \mathbf{B}(\mathbf{E}) = r(u'_m)(c) \cdot \mathbf{E}$. It follows

$$\text{mgu}(r(u_n), r(u'_m)) = \langle \mathbf{B}(\mathbf{E}), \mathbf{E} \rangle .$$

Additionally, there is another non-trivial most-general unifier (of the residuals in the first and third lines of the table on Page 59):

$$\text{mgu}(r(\mathbf{g}(\mathbf{d}(\mathbf{b}^n(\mathbf{e})))), r(\mathbf{g}(\mathbf{f}(\mathbf{b}^n(\mathbf{a})))) = \langle \mathbf{G}(\mathbf{D}, \mathbf{B}(\mathbf{E})), \top \rangle$$

Thus, the set of subtrees of ground unifiers as described in the proof is

$$Z = \{ \mathbf{G}(\mathbf{D}, \mathbf{B}(\mathbf{E})), \mathbf{D}, \mathbf{B}(\mathbf{E}), \mathbf{E} \} .$$

To get a unified earliest partition, we have to change the output and the residual of every tree u with $\sigma(u) \in Z$. For instance, we have $\sigma(\mathbf{e}) = \mathbf{E}$. Thus, we get $\sigma'(\mathbf{e}) = *$ and the residual $r'(\mathbf{e})$ is defined by

$$\begin{aligned} \mathbf{d}(\mathbf{b}^k(y)) &\mapsto \mathbf{A}(\mathbf{B}^k(\mathbf{E}), \mathbf{E}) & \mathbf{d}'(\mathbf{b}^k(y)) &\mapsto \mathbf{A}(\mathbf{B}^{k+1}(\mathbf{E}), \mathbf{B}^k(\mathbf{E})) \\ \mathbf{g}(\mathbf{d}(\mathbf{b}^k(y))) &\mapsto \mathbf{G}(\mathbf{D}, \mathbf{B}^k(\mathbf{E})) & \mathbf{g}(\mathbf{d}'(\mathbf{b}^k(y))) &\mapsto \mathbf{G}(\mathbf{D}, \mathbf{B}^{k+1}(\mathbf{E})) . \end{aligned}$$

For the sake of completeness, we present the unified earliest partition in Figure 5.6. The proof provides that this partition is also path-finite and has finite index. \triangleleft

5.2.5 Congruence Relation

Having a closer look at the partition constructed in the proof of Theorem 5.10, we observe that the equivalence relation $\equiv_{r'}$ is a congruence relation, i.e., for equivalent trees the same context results in equivalent trees (cf. Page 21). In general, this is not implicitly given if a partition is unified earliest. The following Example 5.10 presents a further partition for the tree function τ_{mgu} of Example 5.8. Even though this partition is unified earliest, path-finite, and has finite index, the partition's equivalence relation is not a congruence.

Example 5.10. Consider the tree function τ_{mgu} in Figure 5.4. The following partition $(\hat{\sigma}, \hat{r})$ is unified earliest, path-finite, and has finite index. It differs from the partition (σ, r) of Example 5.8 only for those trees, which realize a ground most-general unifier. In contrast to that, the partition (σ', r') of Example 5.9 never outputs any ground subtree of a most-general unifier, although if it is not involved.

u	$\sigma'(u)$	y	$r'(u)$	$g(y)$
$g(d(b^n(e)))$	$G(D, B^n(E))$	y		
$g(d(b(e)))$	*	$G(D, B(E))$		
$g(d(e))$	$G(D, E)$	y		
$g(d'(b^n(e)))$	$G(D, B^{n+1}(E))$	y		
$g(d'(b(e)))$	$G(D, B^2(E))$	y		
$g(d'(e))$	*	$G(D, B(E))$		
$g(f(b^m(a)))$	*	$G(D, B(E))$		
$d(b^n(e))$	$B^n(E)$	$A(y, E)$		$G(D, y)$
$d(b(e))$	*	$A(B(E), E)$		$G(D, B(E))$
$d(e)$	*	$A(E, E)$		$G(D, E)$
$d'(b^n(e))$	$B^n(E)$	$A(B(y), y)$		$G(D, B(y))$
$d'(b(e))$	*	$A(B^2(E), B(E))$		$G(D, B^2(E))$
$d'(e)$	*	$A(B(E), E)$		$G(D, B(E))$
$f(b^m(a))$	*	$A(B(E), E)$		$G(D, B(E))$

u	$\sigma'(u)$	$d(b^k(y))$	$r'(u)$	$g(d(b^k(y)))$
$b^n(e)$	$B^n(E)$	$A(B^k(y), E)$	$A(B^{k+1}(y), B^k(y))$	$G(D, B^{k+1}(y))$
$b(e)$	*	$A(B^{k+1}(E), E)$	$A(B^{k+2}(E), B^{k+1}(E))$	$G(D, B^{k+2}(E))$
e	*	$A(B^k(E), E)$	$A(B^{k+1}(E), B^k(E))$	$G(D, B^{k+1}(E))$

u	$\sigma'(u)$	$b^k(y)$	$r'(u)$	$g(f(b^k(y)))$
$b^n(a)$	$B^n(E)$	$B^k(y)$	$A(B(E), E)$	$G(D, B(E))$
$b(a)$	*	$B^{k+1}(E)$	$A(B(E), E)$	$G(D, B(E))$
a	*	$B^k(E)$	$A(B(E), E)$	$G(D, B(E))$

Figure 5.6: The unified earliest partition (σ', r') of Example 5.9. ($m \in \mathbb{N}_0, n \geq 2$)

In comparison to the partition (σ, r) , the residual and the output of $\mathbf{b}^n(\mathbf{e})$ and $\mathbf{b}^n(\mathbf{a})$ (for all $n \in \mathbb{N}_0$) does not change (Figure 5.5), i.e.,

$$\begin{aligned} \hat{r}(\mathbf{b}^n(\mathbf{e})) &= r(\mathbf{b}^n(\mathbf{e})) & \hat{\sigma}(\mathbf{b}^n(\mathbf{e})) &= \mathbf{B}^n(\mathbf{E}) \\ \hat{r}(\mathbf{b}^n(\mathbf{a})) &= r(\mathbf{b}^n(\mathbf{a})) & \hat{\sigma}(\mathbf{b}^n(\mathbf{a})) &= \mathbf{B}^n(\mathbf{E}) . \end{aligned}$$

Apart from that, the subtree function $\hat{\sigma}$ and the residual function \hat{r} are defined by

u	$\hat{\sigma}(u)$	$\hat{r}(u)$		
		y	$\mathbf{g}(y)$	
$\mathbf{g}(\mathbf{d}(\mathbf{b}^n(\mathbf{e})))$	$\mathbf{G}(\mathbf{D}, \mathbf{B}^n(\mathbf{E}))$	y		$n \in \mathbb{N}_0$
$\mathbf{g}(\mathbf{d}'(\mathbf{b}^n(\mathbf{e})))$	$\mathbf{G}(\mathbf{D}, \mathbf{B}^{n+1}(\mathbf{E}))$	y		$n \in \mathbb{N}_0$
$\mathbf{g}(\mathbf{f}(\mathbf{b}^n(\mathbf{a})))$	*	$\mathbf{G}(\mathbf{D}, \mathbf{B}(\mathbf{E}))$		$n \in \mathbb{N}_0$
$\mathbf{d}(\mathbf{b}^n(\mathbf{e}))$	$\mathbf{B}^n(\mathbf{E})$	$\mathbf{A}(y, \mathbf{E})$	$\mathbf{G}(\mathbf{D}, y)$	$n \neq 1$
$\mathbf{d}'(\mathbf{b}^n(\mathbf{e}))$	$\mathbf{B}^n(\mathbf{E})$	$\mathbf{A}(\mathbf{B}(y), y)$	$\mathbf{G}(\mathbf{D}, \mathbf{B}(y))$	$n \neq 0$
$\mathbf{d}(\mathbf{b}(\mathbf{e}))$	*	$\mathbf{A}(\mathbf{B}(\mathbf{E}), \mathbf{E})$	$\mathbf{G}(\mathbf{D}, \mathbf{B}(\mathbf{E}))$	
$\mathbf{d}'(\mathbf{e})$	*	$\mathbf{A}(\mathbf{B}(\mathbf{E}), \mathbf{E})$	$\mathbf{G}(\mathbf{D}, \mathbf{B}(\mathbf{E}))$	
$\mathbf{f}(\mathbf{b}^n(\mathbf{a}))$	*	$\mathbf{A}(\mathbf{B}(\mathbf{E}), \mathbf{E})$	$\mathbf{G}(\mathbf{D}, \mathbf{B}(\mathbf{E}))$	$n \in \mathbb{N}_0$

The two trees \mathbf{e} and $\mathbf{b}(\mathbf{e})$ are equivalent. However, if we enhance the trees by the context $\mathbf{d}(y)$, we get non-equivalent trees $\mathbf{d}(\mathbf{e})$ and $\mathbf{d}(\mathbf{b}(\mathbf{e}))$. Thus, this partition is not congruent, but it is unified earliest, path-finite and has finite index. \triangleleft

This example shows that congruence is not implied by the previous \hat{r} properties. However, the proof of Theorem 5.10 provides an equivalent partition with congruence relation for every path-finite partition with finite index. To prove this, we first define the image of a subcontext c of a context $c' \cdot c$ of a tree u . This (unique) image does not exist in every partition. We will show for every unified earliest, path-finite partition with finite index holds: If for every pair c, u there is a unique image of c , depending on c and u only, the equivalence relation is already a congruence relation (and vice versa).

Definition 5.8. Let (σ, r) be a partition, $u \in \text{dom}(\sigma)$, and $c \in \mathcal{C}_\Sigma(y)$ with $c \cdot u \in \text{dom}(\sigma)$. Let $Z_{u,c} \subseteq \mathcal{T}_\Delta(y)$ be the set of possible images of c w.r.t. u , i.e., $z \in Z_{u,c}$ iff for all contexts $c' \in \text{dom}(r(c \cdot u))$ holds

$$r(c \cdot u)(c') \cdot z = r(u)(c' \cdot c) .$$

The *dedicated image* $\text{di}(u)(c) \in \mathcal{T}_\Delta(y)$ of c w.r.t. u is then defined by:

- If $r(c \cdot u)$ is essential and $Z_{u,c}$ is a singleton, it is the unique image $z \in Z_{u,c}$.
- If $r(c \cdot u)$ is inessential and $Z_{u,c} \neq \emptyset$, it equals $*$.
- Otherwise, it is undefined.

Note that in this definition, it is not required that c is a context of u w.r.t. r , i.e., c is not necessarily in $\text{dom}(r(u))$. Furthermore, $c \in \text{dom}(r(u))$ does not imply $\text{di}(u)(c) = r(u)(c)$.

If the dedicated image of $c \in \mathcal{C}_\Sigma(y)$ w.r.t. $u \in \text{dom}(\sigma)$ is defined, the following holds for all contexts $c' \in \text{dom}(r(c \cdot u))$:

$$r(c \cdot u)(c') \cdot \text{di}(u)(c) = r(u)(c' \cdot c) \quad (5.10)$$

Example 5.11. Consider again the tree function in Figure 5.4 and the partition (σ', r') in Example 5.9: For all trees $u \in \text{dom}(\sigma')$, $c \in \mathcal{C}_\Sigma(y)$ with $c \cdot u \in \text{dom}(\sigma')$ the dedicated image $\text{di}(u)(c)$ with respect to the partition (σ', r') in Example 5.9 exists. Consider for instance $u = \mathbf{b}(\mathbf{e})$ and $c = \mathbf{d}(y)$. The residual $r'(\mathbf{d}(\mathbf{b}(\mathbf{e})))$ is inessential and for every $c' \in \text{dom}(r'(\mathbf{b}(\mathbf{e})))$ holds the Equation 5.10:

$$\begin{aligned} r'(\mathbf{d}(\mathbf{b}(\mathbf{e}))) (y) \cdot * &= \mathbf{A}(\mathbf{B}(\mathbf{E}), \mathbf{E}) = r'(\mathbf{b}(\mathbf{e}))(\mathbf{d}(y)) \\ r'(\mathbf{d}(\mathbf{b}(\mathbf{e}))) (\mathbf{g}(y)) \cdot * &= \mathbf{G}(\mathbf{D}, \mathbf{B}(\mathbf{E})) = r'(\mathbf{b}(\mathbf{e}))(\mathbf{g}(\mathbf{d}(y))) \end{aligned}$$

Thus, the dedicated image $\text{di}(\mathbf{e})(\mathbf{b}(y))$ equals $*$.

With respect to the partition $(\hat{\sigma}, \hat{r})$ of Example 5.10, the dedicated image does not exist for every pair u, c . Consider the same pair: $u = \mathbf{b}(\mathbf{e})$ and $c = \mathbf{d}(y)$. For every z holds:

$$\begin{aligned} \hat{r}(\mathbf{d}(\mathbf{b}(\mathbf{e}))) (y) \cdot z &= \mathbf{A}(\mathbf{B}(\mathbf{E}), \mathbf{E}) \neq \mathbf{A}(y, \mathbf{E}) = \hat{r}(\mathbf{b}(\mathbf{e}))(\mathbf{d}(y)) \\ \hat{r}(\mathbf{d}(\mathbf{b}(\mathbf{e}))) (\mathbf{g}(y)) \cdot z &= \mathbf{G}(\mathbf{D}, \mathbf{B}(\mathbf{E})) \neq \mathbf{G}(\mathbf{D}, y) = \hat{r}(\mathbf{b}(\mathbf{e}))(\mathbf{g}(\mathbf{d}(y))) \quad \triangleleft \end{aligned}$$

This example shows that not for every unified earliest, path-finite partition with finite index the dedicated images are defined. However, another property of the counterexample partition $(\hat{\sigma}, \hat{r})$ is that its equivalence relation is not a congruence (cf. Example 5.10). We will show that we can generalize this observation, i.e., for every unified earliest, path-finite partition with finite index, these properties are equivalent.

Lemma 5.11. *For every unified earliest and path-finite partition (σ, r) with finite index, the following are equivalent*

1. *The equivalence \equiv_r is a congruence relation.*
2. *For all $u \in \text{dom}(\sigma)$ and every $c \in \mathcal{C}_\Sigma(y)$ with $c \cdot u \in \text{dom}(\sigma)$, the dedicated image $\text{di}(u)(c)$ is defined.*

Proof. Let τ be the tree function defined by (σ, r) . First, assume that \equiv_r is a congruence, i.e., for all $c \in \mathcal{C}_\Sigma(y)$ with $c \cdot u \in \text{dom}(\sigma)$ holds:

$$u \equiv_r u' \implies c \cdot u \equiv_r c \cdot u' .$$

Let $u \in \text{dom}(\sigma)$ and $c \in \mathcal{C}_\Sigma(y)$ with $c \cdot u \in \text{dom}(\sigma)$. For every $c' \in \text{dom}(r(c \cdot u))$ holds

$$r(c \cdot u)(c') \cdot \sigma(c \cdot u) = r(u)(c' \cdot c) \cdot \sigma(u) .$$

Assume $r(c \cdot u)$ is inessential but $r(u)$ is essential. Then there is a tree $u' \equiv_r u$ with $\sigma(u') \neq \sigma(u)$. Furthermore, it holds $c \cdot u \equiv_r c \cdot u'$. Thus, we get for all $c' \in \text{dom}(r(c \cdot u))$

$$\begin{aligned} r(u)(c' \cdot c) \cdot \sigma(u) &= r(c \cdot u)(c') \\ &= r(c \cdot u')(c') \\ &= r(u')(c' \cdot c) \cdot \sigma(u') \\ &= r(u)(c' \cdot c) \cdot \sigma(u'). \end{aligned}$$

Since $\sigma(u) \neq \sigma(u')$, for all $c' \in \text{dom}(r(c \cdot u))$ holds $r(u')(c' \cdot c) = r(u)(c' \cdot c) \in \mathcal{T}_\Delta$. It follows that for all $c' \in \text{dom}(r(c \cdot u))$ holds $r(c \cdot u)(c') \cdot * = r(u)(c' \cdot c)$. Thus, the dedicated image $\text{di}(u)(c)$ equals $*$. Analogously, $\text{di}(u)(c) = *$ follows also if $r(c \cdot u)$ and $r(u)$ are inessential.

Now assume $r(c \cdot u)$ is essential. If $r(u)$ is inessential, it follows for all $c' \in \text{dom}(r(c \cdot u))$

$$r(c \cdot u)(c') \cdot \sigma(c \cdot u) = r(u)(c' \cdot c).$$

The dedicated image is defined, i.e., $\text{di}(u)(c) = \sigma(c \cdot u)$. Thus, assume $r(u)$ and $r(c \cdot u)$ are essential. It follows that there is a tree $u' \equiv_r u$ with $\sigma(u) \neq \sigma(u')$. Since the relation is congruent, also $c \cdot u' \equiv_r c \cdot u$. We get

$$\begin{aligned} r(c \cdot u)(c') \cdot \sigma(c \cdot u) &= r(u)(c' \cdot c) \cdot \sigma(u) \\ r(c \cdot u)(c') \cdot \sigma(c \cdot u') &= r(u)(c' \cdot c) \cdot \sigma(u') \end{aligned}$$

By Proposition 2.1, there are three possibilities. Since $r(c \cdot u)$ is essential, there is a context c' with $r(c \cdot u)(c') \notin \mathcal{T}_\Delta$. Thus, there is a tree $s_{c'} \in \hat{\mathcal{T}}_\Delta(y)$ with

$$r(c \cdot u)(c') \cdot s_{c'} = r(u)(c' \cdot c) \quad (5.11)$$

$$\text{or} \quad r(c \cdot u)(c') = r(u)(c' \cdot c) \cdot s_{c'} \quad (5.12)$$

If the first equation holds, it follows that $s_{c'} \cdot \sigma(u) = \sigma(c \cdot u)$ and with that, for all context \bar{c} holds the equivalent equation. Thus, $s_{c'} = \text{di}(u)(c)$ is unique. If Equation 5.12 holds, then it holds consequently for all contexts \bar{c} of $c \cdot u$ for the same tree $s_{c'}$. Since the partition is earliest, $s_{c'} = y$ and with that, we can state

$$r(c \cdot u)(c') \cdot y = r(u)(c' \cdot c).$$

The dedicated image then is $\text{di}(u)(c) = y$.

Now assume that Property 2. of the lemma holds, i.e., the dedicated images are defined. We show that \equiv_r is a congruence. Assume the opposite. Then there are subtrees $u, u' \in \text{dom}(\sigma)$ with $u \equiv_r u'$ and $c \in \mathcal{C}_\Sigma(y)$ with $c \cdot u \in \text{dom}(r)$ but

1. $c \cdot u' \notin \text{dom}(r)$ or
2. $c \cdot u \not\equiv_r c \cdot u'$.

First, assume there exists such a triple u, u', c with $c \cdot u' \notin \text{dom}(r)$. Since the partition is trim, there is a context $c' \in \text{dom}(r(c \cdot u))$. Thus, $c' \cdot c$ is a context

of u , i.e., $c' \cdot c \in \text{dom}(r(u)) = \text{dom}(r(u'))$. However, then $c' \cdot c \cdot u'$ is an input tree of τ . Thus, if $u \equiv_r u'$ and $c \cdot u \in \text{dom}(r)$ then $c \cdot u' \in \text{dom}(r)$, too.

Now assume there is a triple u, u', c , which contradicts the second condition, i.e., it exists a context $c' \in \text{dom}(r(c \cdot u)) = \text{dom}(r(c \cdot u'))$ with different images $r(c \cdot u)(c') \neq r(c \cdot u')(c')$. Since for every context c and every tree u with $c \cdot u \in \text{dom}(\sigma)$, the dedicated image $\text{di}(u)(c)$ exists, we get for u, u' that for all $\bar{c} \in \text{dom}(r(c \cdot u))$ holds

$$\begin{aligned} r(c \cdot u)(\bar{c}) \cdot \text{di}(u)(c) &= r(u)(\bar{c} \cdot c) \\ r(c \cdot u')(\bar{c}) \cdot \text{di}(u')(c) &= r(u')(\bar{c} \cdot c) \end{aligned}$$

Since $u \equiv_r u'$ we get for all $\bar{c} \in \text{dom}(r(c \cdot u))$

$$r(c \cdot u)(\bar{c}) \cdot \text{di}(u)(c) = r(c \cdot u')(\bar{c}) \cdot \text{di}(u')(c)$$

Thus, $\langle \text{di}(u)(c), \text{di}(u')(c) \rangle$ is a unifier of $r(c \cdot u), r(c \cdot u')$. The most-general unifier of these residuals then is either

- $\langle y, y \rangle, \langle \top, \top \rangle$, or
- $\langle \text{di}(u)(c), \text{di}(u')(c) \rangle, \langle \text{di}(u)(c), \top \rangle$, or $\langle \top, \text{di}(u')(c) \rangle$ with ground dedicated images $\text{di}(u)(c) \in \mathcal{T}_\Delta$ and $\text{di}(u')(c) \in \mathcal{T}_\Delta$, respectively.

In the first case, the residuals are equivalent, i.e., $c \cdot u \equiv_r c \cdot u'$. In the second case, assume (w.l.o.g.) the dedicated image $\text{di}(u)(c) \in \mathcal{T}_\Delta$ is a ground term and $\text{mgu}(r(c \cdot u), r(c \cdot u')) = \langle \text{di}(u)(c), z' \rangle$ with $z' \in \{\text{di}(u')(c), \top\}$. It follows:

$$r(c \cdot u)(\bar{c}) \cdot \text{di}(u)(c) = r(u)(\bar{c} \cdot c) \cdot \sigma(u) = r(c \cdot u)(\bar{c}) \cdot \sigma(c \cdot u)$$

Thus, $\text{di}(u)(c) = \sigma(c \cdot u)$, which is realizable in $r(c \cdot u)$. This contradicts the fact that (σ, r) is unified earliest. — The equivalence relation is a congruence. \square

Finally, by this lemma and Theorem 5.10 we get for every earliest, path-finite partition with finite index an equivalent partition, which is unified earliest, path-finite and defines a *congruence* relation of finite index. This partition is the one given in the proof of Theorem 5.10. We will show that for this partition the dedicated images exist.

Theorem 5.12. *For every earliest and path-finite partition (σ, r) with finite index, there exists a partition (σ', r') such that*

- (σ', r') is unified earliest and path-finite
- $\equiv_{r'}$ is a congruence of finite index.

Proof. Let τ be the tree function defined by (σ, r) and let (σ', r') be the partition defined in Equations 5.8 and 5.9 in the proof of Theorem 5.10. In that proof, we have already shown that this partition is unified earliest and path-finite. It is also proven that the index of the equivalence relation $\equiv_{r'}$ is finite. It remains to proof that the equivalence relation is already a congruence. Thereto, we show

that for all $u \in \text{dom}(\sigma')$ and every $c \in \mathcal{C}_\Sigma(y)$ with $c \cdot u \in \text{dom}(\sigma')$ the dedicated image $\text{di}(u)(c)$ is defined. Then it follows with Lemma 5.11 that the equivalence relation $\equiv_{r'}$ is a congruence.

For every $u \in \text{dom}(\sigma')$, $c \in \mathcal{C}_\Sigma(y)$ with $c \cdot u \in \text{dom}(\sigma')$, and every context $c' \in \text{dom}(r'(c \cdot u))$, we know

$$r'(c \cdot u)(c') \cdot \sigma'(c \cdot u) = r'(u)(c' \cdot c) \cdot \sigma'(u). \quad (5.13)$$

Thus, let $u \in \text{dom}(\sigma')$ and $c \in \mathcal{C}_\Sigma(y)$ with $c \cdot u \in \text{dom}(\sigma')$. We consider the different possibilities whether the residuals $r'(c \cdot u)$ and $r'(u)$ are each essential or inessential.

- First, assume for all $c' \in \text{dom}(r'(c \cdot u))$ the image of $c' \cdot c$ under $r'(u)$ is ground, i.e., $r'(u)(c' \cdot c) \in \mathcal{T}_\Delta$. Then $r'(c \cdot u)(c') \cdot \sigma'(c \cdot u) = r'(u)(c' \cdot c)$. If $r'(c \cdot u)$ is essential, we get with $\text{di}(u)(c) = \sigma'(c \cdot u)$ the dedicated image of c w.r.t. u . Otherwise, if the residual $r'(c \cdot u)$ is inessential, we get $r'(c \cdot u)(c') = r'(u)(c' \cdot c)$ for all $c' \in \text{dom}(r'(c \cdot u))$ and define the dedicated image by $\text{di}(u)(c) = *$.
- Now assume $r'(c \cdot u)$ is essential and it exists $c' \in \text{dom}(r'(c \cdot u))$ with $r'(u)(c' \cdot c) \in \hat{\mathcal{T}}_\Delta(y)$. Since the partition is proper, there are infinitely many different outputs of equivalent trees, i.e., $|\sigma'([u]_{r'})| = \infty$. The number of equivalence classes is finite. Thus, there are two trees $u_1, u_2 \in [u]_{r'}$ with different outputs for which $c \cdot u_1$ and $c \cdot u_2$ are equivalent. Summarizing, we have

- $u_1 \equiv_{r'} u_2 \equiv_{r'} u$
- $\sigma'(u_1) \neq \sigma'(u_2)$
- $c \cdot u_1 \equiv_{r'} c \cdot u_2$.

We get with Equation 5.13

$$\begin{aligned} r'(c \cdot u_1)(c') \cdot \sigma'(c \cdot u_1) &= r'(u_1)(c' \cdot c) \cdot \sigma'(u_1) \\ r'(c \cdot u_2)(c') \cdot \sigma'(c \cdot u_2) &= r'(u_2)(c' \cdot c) \cdot \sigma'(u_2). \end{aligned}$$

Since the residuals of u_1 and u_2 are equivalent and so are the residuals of $c \cdot u_1$ and $c \cdot u_2$, it follows:

- $\sigma'(c \cdot u_1) \neq \sigma'(c \cdot u_2)$

By Proposition 2.1, there are two possibilities:

$$r'(c \cdot u_1)(c') \cdot z = r'(u_1)(c' \cdot c) \quad (5.14)$$

$$\text{or} \quad r'(c \cdot u_1)(c') = r'(u_1)(c' \cdot c) \cdot z \quad (5.15)$$

for some $z \in \hat{\mathcal{T}}_\Delta(y)$. Case (a) of the Proposition 2.1, i.e., $r'(u_1)(c' \cdot c)$ and $r'(u)(c' \cdot c)$ are equal and ground, is impossible because we assumed $r'(u_1)(c' \cdot c) = r'(u)(c' \cdot c) \in \hat{\mathcal{T}}_\Delta(y)$.

If Equation 5.14 holds, for both $i \in \{1, 2\}$ follows $z \cdot \sigma'(u_i) = \sigma'(c \cdot u_i)$. Then Proposition 2.1 provides that for all $\bar{c} \in \text{dom}(r'(c \cdot u_i))$ the same equation as for c' holds:

$$r'(c \cdot u_1)(\bar{c}) \cdot z = r'(u_1)(\bar{c} \cdot c)$$

If Equation 5.15 holds, it follows analogically for both $i \in [2]$ and all $\bar{c} \in \text{dom}(r'(c \cdot u_i))$:

$$\begin{aligned} - \sigma'(u_i) &= z \cdot \sigma'(c \cdot u_i) \\ - r'(c \cdot u_1)(\bar{c}) &= r'(u_1)(\bar{c} \cdot c) \cdot z. \end{aligned}$$

First, assume that Equation 5.14 holds. The trees u_1 and u are equivalent. This implies, $r'(c \cdot u_1)(\bar{c}) \cdot z = r'(u)(\bar{c} \cdot c)$ holds for all $\bar{c} \in \text{dom}(r'(c \cdot u_1))$. Thus, for all contexts $\bar{c} \in \text{dom}(r'(c \cdot u_1))$ holds

$$r'(c \cdot u)(\bar{c}) \cdot \sigma'(c \cdot u) = r'(u)(\bar{c} \cdot c) \cdot \sigma'(u) = r'(c \cdot u_1)(\bar{c}) \cdot z \cdot \sigma'(u).$$

Hence, $r'(c \cdot u)$ and $r'(c \cdot u_1)$ are unifiable. Since $r'(c \cdot u)$ is essential and $\sigma'(c \cdot u)$ is realizable in $r'(c \cdot u)$, the most-general unifier is

$$\text{mgu}(r'(c \cdot u), r'(c \cdot u_1)) = \langle y, y \rangle.$$

However, that means that the two trees $c \cdot u$ and $c \cdot u_1$ are equivalent and it follows for all $\bar{c} \in \text{dom}(r'(c \cdot u))$

$$r'(c \cdot u)(\bar{c}) \cdot z = r'(c \cdot u_1)(\bar{c}) \cdot z = r'(u_1)(\bar{c} \cdot c) = r'(u_1)(\bar{c} \cdot c).$$

With $\text{di}(u)(c) = z$ we get the dedicated image of c under u .

On the other hand, assume that the second equation (Equation 5.15) holds, i.e., for all $\bar{c} \in \text{dom}(r'(c \cdot u_i))$ holds

$$r'(c \cdot u_1)(\bar{c}) = r'(u_1)(\bar{c} \cdot c) \cdot z.$$

Since the (σ', r') is earliest and z is a common suffix of all images in $r'(c \cdot u_1)$, it follows that $z = y$. Again, for all contexts $\bar{c} \in \text{dom}(r'(c \cdot u_1))$ holds

$$r'(c \cdot u)(\bar{c}) \cdot \sigma'(c \cdot u) = r'(u)(\bar{c} \cdot c) \cdot \sigma'(u) = r'(c \cdot u_1)(\bar{c}) \cdot \sigma'(u).$$

Thus, $r'(c \cdot u)$ and $r'(c \cdot u_1)$ are unifiable. Analogically as in the first case, it follows that $r'(c \cdot u)(\bar{c}) \cdot \text{di}(u)(c) = r'(u)(\bar{c} \cdot c)$ with $\text{di}(u)(c) = y$.

- Finally, assume $r'(c \cdot u)$ is inessential and it exists $c' \in \text{dom}(r'(c \cdot u))$ with $r'(u)(c' \cdot c) \in \hat{\mathcal{T}}_{\Delta}(y)$. Then $r'(u)$ is essential and there are infinitely many trees $u' \equiv_{r'} u$ with $\sigma'(u') \neq \sigma'(u)$.

Assume the residual $r'(c \cdot u')$ is inessential for all these u' . With Equation 5.13 follows

$$\begin{aligned} r'(c \cdot u)(c') &= r'(u)(c' \cdot c) \cdot \sigma'(u) \\ &\neq r'(u')(c' \cdot c) \cdot \sigma'(u') \\ &= r'(c \cdot u')(c'). \end{aligned}$$

Thus, $c \cdot u \not\equiv_{r'} c \cdot u'$. Since the partition has finite index, there is a tree $u' \in [u]_{r'}$ with essential $r'(c \cdot u')$ and $\sigma'(u') \neq \sigma'(u)$. It follows by the previous parts of this proof that the dedicated images $\text{di}(u')(c)$ exist with

$$r'(c \cdot u')(c') \cdot \text{di}(u')(c) = r'(u')(c' \cdot c)$$

for all $c' \in \text{dom}(r'(c \cdot u'))$. Now we get with Equation 5.13 and the equivalence of u and u' that for all $c' \in \text{dom}(r'(c \cdot u))$ holds:

$$r'(c \cdot u')(c') \cdot \text{di}(u')(c) \cdot \sigma'(u) = r'(u')(c' \cdot c) \cdot \sigma'(u') = r'(c \cdot u)(c')$$

Thus, the most-general unifier of the residuals of $c \cdot u'$ and $c \cdot u$ is

$$\text{mgu}(r'(c \cdot u'), r'(c \cdot u)) = \langle \text{di}(u')(c) \cdot \sigma'(u), \top \rangle.$$

Then, a subtree of a ground most-general unifier is realizable in some tree (either $\sigma'(u)$ or $\text{di}(u')(c)$ equals $\sigma'(u')$). This does not contradict the property earliest unified, but the partition (σ', r') in the proof of Theorem 5.10 has no such most-general unifier by construction. \square

5.2.6 Minimal Partition

Let τ be a tree function with a partition, which is path-finite and has finite index and let (σ, r) be a partition of τ , which is path-finite and has finite index, but additionally it is unified earliest and \equiv_r is a congruence. We have seen that such a partition exists. Unfortunately, this partition is not necessarily unique. For instance, for the partition in Figure 5.6, which fulfills these properties, we can increase the index without suffering the loss of one of the previous properties. If we change the definition such that \mathbf{a} , $\mathbf{b}(\mathbf{a})$, and $\mathbf{b}^n(\mathbf{a})$ for all $n \in \mathbb{N}$ are equivalent, the properties are preserved, but the index is smaller.

In this section, we show that the partition for τ with minimal index, which preserves all afore introduced properties, is unique. To prove that, let \sim'_r denote the smallest equivalence relation over $\text{dom}(r)$ with the following properties:

- If $\text{mgu}(r(u), r(u')) = \langle y, y \rangle$ or $\text{mgu}(r(u), r(u')) = \langle \top, \top \rangle$ then $u \sim'_r u'$;
- Assume that $\text{mgu}(r(u), r(u_1)) = \langle \top, v_1 \rangle$ for some ground term $v_1 \in \mathcal{T}_\Delta$. If for all u_2 with $\text{mgu}(r(u), r(u_2)) = \langle \top, v_2 \rangle$ for some ground term $v_2 \in \mathcal{T}_\Delta$, $\text{mgu}(r(u_1), r(u_2)) = \langle y, y \rangle$ holds then $u \sim'_r u_1$.

The relation \sim_r is the greatest congruence relation, which is a refinement of \sim'_r , i.e., $u_1 \sim_r u_2$ whenever for every $c \in \mathcal{C}_\Sigma(y)$ holds $c \cdot u_1 \sim'_r c \cdot u_2$. Two trees $u_1, u_2 \in \mathcal{T}_\Sigma$ are called *similar* w.r.t. r if $u_1 \sim_r u_2$.

Definition 5.9 (Minimal). A unified earliest partition (σ, r) is called *minimal* if similar trees are equivalent, i.e., \equiv_r and \sim_r define the same congruence relation.

If the most-general unifier is known for every pair of residuals, a minimal partition for τ can be defined.

Theorem 5.13. *For every unified earliest and path-finite partition (σ, r) with finite index and \equiv_r is a congruence, there exists an equivalent partition (σ', r') , which is*

- *minimal,*
- *path-finite, and*
- *$\equiv_{r'}$ is a congruence relation of finite index.*

Additionally, the index of (σ', r') is less or equal the index of (σ, r) .

Proof. Let τ be the tree function of (σ, r) . We define the partition (σ', r') on the same domain as follows. Let $u \in \text{dom}(r)$. Assume there is a tree $u' \in \text{dom}(r)$ with $u' \sim_r u$ and $\text{mgu}(r(u'), r(u)) = \langle v', \top \rangle$ for some $v' \in \mathcal{T}_\Delta$. Note that this implies that $r(u)$ is inessential. We define

$$\sigma'(u) = v' \quad \text{and} \quad r'(u)(c) = r(u')(c) \quad (5.16)$$

Otherwise, if no such $u' \sim_r u$ with non trivial most-general unifier exists, the output and the residual of u do not change, i.e.,

$$\sigma'(u) = \sigma(u) \quad \text{and} \quad r'(u)(c) = r(u)(c) \quad (5.17)$$

Let $u \in \text{dom}(r)$ and $c \in \text{dom}(r(u))$. If there is a tree $u' \in \text{dom}(r)$ with $u' \sim_r u$ and $\text{mgu}(r(u'), r(u)) = \langle v', \top \rangle$ for some $v' \in \mathcal{T}_\Delta$, then Equation 5.16 provides

$$r'(u)(c) \cdot \sigma'(u) = r(u')(c) \cdot v' = r(u)(c) = r(u)(c) \cdot \sigma(u) = \tau(c \cdot u) .$$

Otherwise, we get by Equation 5.17

$$r'(u)(c) \cdot \sigma'(u) = r(u)(c) \cdot \sigma(u) = \tau(c \cdot u) .$$

Thus, the pair is a partition of τ . We observe the following properties for all trees $u, u' \in \text{dom}(r)$:

1. If $r(u)$ is essential, then $r'(u)$ is essential and equals $r(u)$.
2. Contrapositive, if $r'(u)$ is inessential, then $r(u)$ is inessential.
3. If $u' \sim_r u$, then u becomes r' -equivalent to u' , i.e., $u \equiv_{r'} u'$.
4. No new residuals are generated, i.e., there exists a tree $u'' \in \text{dom}(r)$ such that $r'(u) = r(u'')$. In particular, $u \sim_r u''$.

First, we show that the partition is minimal. Thereto, we prove, on the one hand, that $\equiv_{r'}$ is a congruence. Then, $\equiv_{r'}$ is a refinement of $\sim'_{r'}$ (cf. Page 72) and with that $\equiv_{r'} \subseteq \sim_{r'}$.

Let $u_1 \equiv_{r'} u_2$. The most-general unifier of $r'(u_1)$ and $r'(u_2)$ is either $\langle \top, \top \rangle$ or $\langle y, y \rangle$. In the first case, both residuals are inessential and with Property 2, we get also $\text{mgu}(r(u_1), r(u_2)) = \langle \top, \top \rangle$. That implies $u_1 \equiv_r u_2$. The equivalence \equiv_r is a congruence. Hence, $u_1 \sim_r u_2$. Now assume $\text{mgu}(r'(u_1), r'(u_2)) = \langle y, y \rangle$. Property 4 provides that there are trees u'_1, u'_2 with $r(u'_i) = r'(u_i)$ for $i \in [2]$. Thus, we get

$$u_1 \sim_r u'_1 \equiv_r u'_2 \sim_r u_2 .$$

The equivalence \equiv_r is a congruence. Hence, $u'_1 \sim_r u'_2$ and with that $u_1 \sim_r u_2$. Thus, in both cases, we get that u_1 and u_2 are similar w.r.t. r . Since \sim_r is a congruence, it follows that for all $c \in \mathcal{C}_\Sigma(y)$ holds $c \cdot u_1 \sim_r c \cdot u_2$. In addition, Property 3 implies that for all $c \in \mathcal{C}_\Sigma(y)$ holds $c \cdot u_1 \equiv_{r'} c \cdot u_2$.

On the other hand, we have to show that $\sim_{r'} \subseteq \equiv_{r'}$. First, we show that for all trees u_1, u_2 , which are similar in (σ', r') , holds $u_1 \sim'_{r'} u_2$, i.e:

$$u_1 \sim_{r'} u_2 \implies u_1 \sim'_{r'} u_2 . \quad (5.18)$$

Thereto, let $u_1 \sim_{r'} u_2$. The most-general unifier is

1. $\text{mgu}(r'(u_1), r'(u_2)) = \langle \top, \top \rangle$,
2. $\text{mgu}(r'(u_1), r'(u_2)) = \langle y, y \rangle$,
3. $\text{mgu}(r'(u_1), r'(u_2)) = \langle v_1, \top \rangle$ for some $v_1 \in \mathcal{T}_\Delta$, or
4. $\text{mgu}(r'(u_1), r'(u_2)) = \langle \top, v_2 \rangle$ for some $v_2 \in \mathcal{T}_\Delta$.

In the first case, also $\text{mgu}(r(u_1), r(u_2)) = \langle \top, \top \rangle$ (by Property 2) and with that $u_1 \sim'_{r'} u_2$. In the second case, we know by Property 3 that there is a tree u_3 such that $r(u_3) = r'(u_1) = r'(u_2)$ and $u_1 \sim_r u_3 \sim_r u_2$.

Let us now assume the most-general unifier equals $\langle v_1, \top \rangle$ for some ground term $v_1 \in \mathcal{T}_\Delta$. Additionally, for every tree u_3 holds (by definition of $\sim'_{r'}$)

$$\text{mgu}(r'(u_3), r'(u_2)) = \langle v_3, \top \rangle \text{ with } v_3 \in \mathcal{T}_\Delta \implies u_1 \sim_{r'} u_3 . \quad (5.19)$$

Since u_1 and u_3 have both essential residuals under r' , they must be r' -equivalent. For u_2 holds $r'(u_2) = r(u_2)$ is inessential (Property 2).

First, assume $r(u_1) = r'(u_1)$ is essential. Then $\text{mgu}(r(u_1), r(u_2)) = \langle v_1, \top \rangle$ and for all u_3 with $\text{mgu}(r(u_3), r(u_2)) = \langle v_3, \top \rangle$ and $v_3 \in \mathcal{T}_\Delta$ holds with Property 1 also $\text{mgu}(r'(u_3), r'(u_2)) = \langle v_3, \top \rangle$. It follows with Equation 5.19 that

$$r(u_3) = r'(u_3) = r'(u_1) = r(u_1) .$$

Thus, $u_1 \sim'_{r'} u_2$ (by definition of $\sim'_{r'}$).

Otherwise, assume $r(u_1)$ is inessential and different to $r'(u_1)$. There is some tree u_4 with $r'(u_1) = r(u_4)$ (Property 4) and $\text{mgu}(r(u_1), r(u_4)) = \langle \top, v_4 \rangle$ for some $v_4 \in \mathcal{T}_\Delta$ and $u_4 \sim_r u_1$. With Property 1 we get $r'(u_4) = r'(u_1)$. Then

$$\text{mgu}(r'(u_4), r'(u_2)) = \langle v_1, \top \rangle$$

for u_4 with essential residual $r(u_4)$. Analog to the last assumption, we get $u_4 \sim'_r u_2$. It follows $u_1 \sim'_r u_4 \sim'_r u_2$.

The last case that the most-general unifier equals $\langle \top, v_2 \rangle$ for some $v_2 \in \mathcal{T}_\Delta$ is analog. Summarizing, we get Implication 5.18 for all trees u_1, u_2 . Since $\sim_{r'}$ is a congruence, we get:

$$\begin{aligned} u_1 \sim_{r'} u_2 &\implies \text{for all } c \in \mathcal{C}_\Sigma(y) : c \cdot u_1 \sim_{r'} c \cdot u_2 && \text{(Congruence)} \\ &\implies \text{for all } c \in \mathcal{C}_\Sigma(y) : c \cdot u_1 \sim'_r c \cdot u_2 && \text{(Implication 5.18)} \\ &\implies u_1 \sim_r u_2 \\ &\implies u_1 \equiv_{r'} u_2 && \text{(Property 3)} \end{aligned}$$

It remains to prove that the partition is unified earliest. Then, we get that the partition is minimal. Thus, we have to show that (σ', r') is a trim, proper, earliest, unified earliest, and path-finite partition, and has finite index.

- It is trim and proper: We only have to consider trees where the residual under r' differs from the residual under r . In that case, it equals a residual of a tree u' , i.e., $r'(u) = r(u') = r'(u')$ (Property 1), which is essential. Thus, $\sigma'(u) = s' \neq *$ implies that $r'(u)$ is essential. — It is trim.

In addition, for every essential residual, the set of outputs only increases, i.e., $\sigma([u]_r) \subseteq \sigma'([u]_{r'})$. — It is proper.

- The partition is earliest, path-finite, and has finite index: With Property 4, these properties are directly inherited from (σ, r) .
- It is unified earliest: Assume there is a unifiable pair of residuals in the range of r' with realizable ground unifier, i.e., there are trees u_1, u_2 with $\text{mgu}(r'(u_1), r'(u_2)) = \langle v_1, v_2 \rangle$ where $\sigma'(u_1) = v_1 \in \mathcal{T}_\Delta$ is realizable in $r'(u_1)$ and $v_2 \in \mathcal{T}_\Delta \cup \{\top\}$. Property 4 implies that there are trees u'_1, u'_2 in $\text{dom}(r)$ with $r(u'_i) = r'(u_i)$ and $u_i \sim_r u'_i$ (for $i \in \{1, 2\}$). Thus, for the most-general unifier we get

$$\text{mgu}(r(u'_1), r(u'_2)) = \langle v_1, v_2 \rangle .$$

Since (σ, r) is unified earliest, $u_1 \not\equiv_r u'_1$. Otherwise, $\sigma(u_1) = \sigma'(u_1) = v_1$ would be realizable in $r(u_1)$. Hence, the residual $r(u'_1)$ is inessential and we get $\text{mgu}(r(u_1), r(u'_1)) = \langle \top, v_1 \rangle$. Additionally, we have for the residuals of u_1 and u'_2 : $\text{mgu}(r(u_1), r(u'_2)) = \langle \top, v_2 \rangle$. If $v_2 = \top$, $u_1 \sim_r u'_2 \sim_r u_2$ and with Property 3: $u_1 \equiv_{r'} u_2$. This contradicts the assumption that the most-general unifier $\text{mgu}(r'(u_1), r'(u_2))$ is different from $\langle y, y \rangle$ and $\langle \top, \top \rangle$. Otherwise, if $v_2 \neq \top$, we get that $u'_1 \sim_r u'_2$ (Definition of \sim_r). It follows

with Property 3: $u'_1 \equiv_r u'_2$. Since $r(u'_1)$ and $r(u'_2)$ are essential, we also have (with Property 1) $u'_1 \equiv_{r'} u'_2$. Furthermore, we know for both $i \in [2]$ that u'_i is r' -equivalent to u_i , i.e., we get

$$u_1 \equiv_{r'} u'_1 \equiv_{r'} u'_2 \equiv_{r'} u_2 .$$

Again, this contradicts the assumption for $\text{mgu}(r'(u_1), r'(u_2))$. \square

In Example 5.9 we have defined a unified earliest partition for the example transformation τ_{mgu} . Since we get this partition by the construction of the proof of Theorem 5.10, we know by the proof of Theorem 5.12 that the relation of this partition is a congruence relation. In the following example, we build the minimal partition for the transformation τ_{mgu} as described in the proof of Theorem 5.13:

Example 5.12. Let τ_{mgu} be the tree function of Example 5.8 and (σ, r) the unified earliest partition with congruence relation of Example 5.9 (Figure 5.6). Now we apply the construction of the previous proof. Thereto, we have to consider the most-general unifiers to deduce the congruence relation \sim_r . We are only interested in the unifiers of an inessential and an essential residual. For every residual, we choose the first tree occurring in the table on Page 65 as representative. Let $u = \mathbf{b}^2(\mathbf{e})$. We get the following:

$$\text{mgu}(r(\mathbf{g}(\mathbf{d}(u))), r(\mathbf{g}(\mathbf{d}(\mathbf{b}(\mathbf{e})))) = \langle \mathbf{G}(\mathbf{D}, \mathbf{B}(\mathbf{E})), \top \rangle \quad (5.20)$$

$$\text{mgu}(r(\mathbf{d}(u)), r(\mathbf{d}(\mathbf{b}(\mathbf{e})))) = \langle \mathbf{B}(\mathbf{E}), \top \rangle \quad (5.21)$$

$$\text{mgu}(r(\mathbf{d}'(u)), r(\mathbf{d}(\mathbf{b}(\mathbf{e})))) = \langle \mathbf{E}, \top \rangle \quad (5.22)$$

$$\text{mgu}(r(\mathbf{d}(u)), r(\mathbf{d}(\mathbf{e})))) = \langle \mathbf{E}, \top \rangle \quad (5.23)$$

$$\text{mgu}(r(\mathbf{d}'(u)), r(\mathbf{d}'(\mathbf{b}(\mathbf{e})))) = \langle \mathbf{B}(\mathbf{E}), \top \rangle \quad (5.24)$$

$$\text{mgu}(r(u), r(\mathbf{b}(\mathbf{e})))) = \langle \mathbf{B}(\mathbf{E}), \top \rangle \quad (5.25)$$

$$\text{mgu}(r(u), r(\mathbf{e})))) = \langle \mathbf{E}, \top \rangle \quad (5.26)$$

$$\text{mgu}(r(\mathbf{b}^2(\mathbf{a})), r(\mathbf{b}(\mathbf{a})))) = \langle \mathbf{B}(\mathbf{E}), \top \rangle \quad (5.27)$$

$$\text{mgu}(r(\mathbf{b}^2(\mathbf{a})), r(\mathbf{a})))) = \langle \mathbf{E}, \top \rangle \quad (5.28)$$

Now we see in the first equation that tree $\mathbf{g}(\mathbf{d}(\mathbf{b}(\mathbf{e})))$ with inessential residual is potential similar to $\mathbf{g}(\mathbf{d}(u))$, i.e., $\mathbf{g}(\mathbf{d}(\mathbf{b}(\mathbf{e}))) \sim_r' \mathbf{g}(\mathbf{d}(u))$. Their residuals have a most-general unifier of the form $\langle v, \top \rangle$ with ground tree $v \in \mathcal{T}_\Delta$ and there is no other essential residual in the range of r , which has such a most-general unifier with $r(\mathbf{g}(\mathbf{d}(\mathbf{b}(\mathbf{e}))))$. Since y is the only context in $\text{dom}(r(\mathbf{g}(\mathbf{d}(u))))$, the trees are similar:

$$\mathbf{g}(\mathbf{d}(\mathbf{b}(\mathbf{e}))) \sim_r \mathbf{g}(\mathbf{d}(u)) .$$

Let us consider the tree $t = \mathbf{d}(\mathbf{b}(\mathbf{e}))$. Also the residuals of $\mathbf{d}(u)$ and t have a most-general unifier of the form $\langle s, \top \rangle$ with ground tree $s \in \mathcal{T}_\Delta$ (Equation 5.21). On the other hand, also the residuals of $\mathbf{d}'(u)$ and t have such a unifier (Equation 5.22). The trees $\mathbf{d}(u)$ and $\mathbf{d}'(u)$ are not equivalent, hence, $\mathbf{d}(\mathbf{b}(\mathbf{e}))$ is similar

to none of the trees with essential residual. In particular,

$$\mathbf{d}(\mathbf{b}(\mathbf{e})) \not\sim_r \mathbf{d}(u) .$$

This difference causes that $\mathbf{b}(\mathbf{e})$ and u are not similar, as well as \mathbf{e} and $\mathbf{b}(\mathbf{e})$. Also u and \mathbf{e} are not similar. Altogether, we observe the following similarities (of representatives of the equivalence classes):

$$\begin{aligned} \mathbf{g}(\mathbf{d}(u)) &\sim_r \mathbf{g}(\mathbf{d}(\mathbf{b}(\mathbf{e}))) \\ \mathbf{d}(u) &\sim_r \mathbf{d}(\mathbf{e}) \\ \mathbf{d}'(u) &\sim_r \mathbf{d}'(\mathbf{b}(\mathbf{e})) \\ \mathbf{b}^2(\mathbf{a}) &\sim_r \mathbf{b}(\mathbf{a}) \sim_r \mathbf{a} \end{aligned}$$

Additionally, for all trees $u, u' \in \text{dom}(r)$ with $u \equiv_r u'$ holds $u \sim_r u'$. We have to adapt the output under σ of these trees with inessential residuals to the essential residuals of the similar trees. The bottom line is the partition (σ', r') given in Figure 5.7. This congruence relation $\equiv_{r'}$ equals $\sim_{r'}$. It has eight equivalence classes. \triangleleft

It follows that for every path-finite partition of finite index, an equivalent minimal and path-finite partition with finite index exists. Moreover, we will see that this partition is unique for its transformation.

Theorem 5.14. *Let τ be a tree function with path-finite partition of finite index, then the equivalent minimal and path-finite partition with finite index is unique.*

Proof. Let (σ, r) and (σ', r') be two minimal and path-finite partitions with finite index of the same tree function τ . We show that $\sigma = \sigma'$ and $r = r'$.

If there is a tree $u \in \text{dom}(r)$ and context $c \in \text{dom}(r(u))$, then $c \cdot u$ is in the domain of τ and with that both $u \in \text{dom}(r')$ and $c \in \text{dom}(r'(u))$. The partitions have the same domain and also the residuals of a tree $u \in \text{dom}(r)$ under r and r' have the same domain, i.e., $\text{dom}(r(u)) = \text{dom}(r'(u))$. We prove that for all trees $u \in \text{dom}(r)$ holds

$$[u]_r = [u]_{r'} \quad \text{and} \quad r(u) = r'(u) .$$

It follows immediately that $\sigma = \sigma'$. Let $u \in \text{dom}(r)$. For every $c \in \text{dom}(r(u))$ holds:

$$r(u)(c) \cdot \sigma(u) = \tau(c \cdot u) = r'(u)(c) \cdot \sigma'(u) \quad (5.29)$$

First assume that $r(u)$ and $r'(u)$ are inessential. It follows that $r(u) = r'(u)$. For every tree $u' \in [u]_r$ with $r'(u')$ is inessential, too, the residual under r' equals the residual of u , i.e., $r'(u') = r'(u)$. It follows $u' \in [u]_{r'}$.

Now assume that $r(u)$ is essential. We will show the following:

$$\exists u_1 \equiv_r u \text{ with } r(u) = r(u_1) = r'(u_1) \quad (5.30)$$

u	$\sigma'(u)$	$r'(u)$	$g(y)$
$g(d(b^{mn}(e)))$	$G(D, B^{mn}(E))$	y	$G(D, y)$
$g(d'(b^m(e)))$	$G(D, B^{m+1}(E))$	y	$G(D, B(E))$
$g(f(b^{mn}(a)))$	$G(D, B(E))$	y	$G(D, y)$
$d(b^n(e))$	$B^n(E)$	$A(y, E)$	$G(D, y)$
$d(b(e))$	$*$	$A(B(E), E)$	$G(D, B(E))$
$d'(b^n(e))$	E	$A(y, E)$	$G(D, y)$
$d'(b(e))$	$B^n(E)$	$A(B(E), E)$	$G(D, B(E))$
$d(b(e))$	$B(E)$	$A(B(y), y)$	$G(D, B(y))$
$d'(e)$	$*$	$A(B(E), E)$	$G(D, B(E))$
$f(b^{mn}(a))$	$*$	$A(B(E), E)$	$G(D, B(E))$

u	$\sigma'(u)$	$r'(u)$		$g(d(b^k(y)))$	$g(d'(b^k(y)))$
$b^n(e)$	$B^n(E)$	$d(b^k(y))$	$d'(b^k(y))$	$G(D, B^k(y))$	$G(D, B^{k+1}(y))$
$b(e)$	$*$	$A(B^k(y), E)$	$A(B^{k+1}(y), B^k(y))$	$G(D, B^{k+1}(E))$	$G(D, B^{k+2}(E))$
e	$*$	$A(B^{k+1}(E), E)$	$A(B^{k+2}(E), B^{k+1}(E))$	$G(D, B^k(E))$	$G(D, B^{k+1}(E))$

u	$\sigma'(u)$	$r'(u)$		$g(f(b^k(y)))$
$b^{mn}(a)$	$B^{mn}(E)$	$b^k(y)$	$f(b^k(y))$	$G(D, B(E))$
		$B^k(y)$	$A(B(E), E)$	$G(D, B(E))$

Figure 5.7: The minimal partition (σ', r') of Example 5.12. ($m \in \mathbb{N}_0, n \geq 2$)

There are infinitely many different trees in $[u]_r$ with different outputs under σ , i.e., $|\sigma([u]_r)| = \infty$. Since the index of $\equiv_{r'}$ is finite, there is an infinite set of such trees, which are r' -equivalent, i.e., there is a tree $u' \equiv_r u$ and the set $S = [u]_r \cap [u']_{r'}$ is infinite and additionally $|\sigma(S)| = \infty$. Let $u_1, u_2 \in S$ with $\sigma(u_1) \neq \sigma(u_2)$. Note that the two trees u_1 and u_2 are equivalent with respect to both r and r' . Assume the residual of u_1 under r' is inessential. There is a context $c \in \text{dom}(r(u_1))$ with $r(u_1)(c) \in \hat{\mathcal{T}}_\Delta(y)$. We get

$$r(u_1)(c) \cdot \sigma(u_1) = r'(u_1)(c) = r'(u_2)(c) = r(u_1)(c) \cdot \sigma(u_2)$$

which is a contradiction to $\sigma(u_1) \neq \sigma(u_2)$. Thus, the residual $r'(u_1)$ is essential. With Proposition 2.1 and Equation 5.29 we will show that $r(u_1) = r'(u_1)$: For every context $c \in \text{dom}(r(u_1))$ holds

$$\begin{aligned} r(u_1)(c) \cdot \sigma(u_1) &= r'(u_1)(c) \cdot \sigma'(u_1) \\ r(u_2)(c) \cdot \sigma(u_2) &= r'(u_2)(c) \cdot \sigma'(u_2) \end{aligned}$$

For every context c , either $r(u_1)(c) = r'(u_1)(c) \in \mathcal{T}_\Delta$ is ground or there is a tree $z_c \in \hat{\mathcal{T}}_\Delta(y)$ with $r(u_1)(c) \cdot z_c = r'(u_1)(c)$ or $r(u_1)(c) = r'(u_1)(c) \cdot z_c$ (Proposition 2.1). Assume there are c and c' with different trees $z_c \neq z_{c'}$ and $r(u_1)(c) \cdot z_c = r'(u_1)(c)$ and $r(u_1)(c') \cdot z_{c'} = r'(u_1)(c')$, respectively. It follows that $\sigma(u_i) = z_c \cdot \sigma'(u_i) = z_{c'} \cdot \sigma'(u_i)$ (for both $i \in [2]$). It follows, again with Proposition 2.1 that $z_c = z_{c'}$. Otherwise, assume there are c and c' with different trees z_c and $z_{c'}$ and

$$\begin{aligned} r(u_1)(c) \cdot z_c &= r'(u_1)(c) \\ r(u_1)(c') &= r'(u_1)(c') \cdot z_{c'} \end{aligned}$$

Then $\sigma(u_i) = z_c \cdot \sigma'(u_i)$ and $z_{c'} \cdot \sigma(u_i) = \sigma'(u_i)$ for both $i \in [2]$. It follows that $z_c = z_{c'} = y$. Summarizing, there is a tree $z \in \hat{\mathcal{T}}_\Delta(y)$ with $r(u_1)(c) \cdot z = r'(u_1)(c)$ for every context c or $r(u_1)(c) = r'(u_1)(c) \cdot z$ for every context c . Since r and r' are earliest, z must equal y . Consequently, $r(u_1) = r'(u_1)$. This gives us Equation 5.30.

Now assume $r(u)$ is essential and $r'(u)$ is inessential. There is a tree $u_1 \equiv_r u$ with essential residual $r(u) = r(u_1) = r'(u_1)$ (cf. Equation 5.30). For all context $c \in \text{dom}(r(u))$, we get

$$r'(u_1)(c) \cdot \sigma(u) = r(u)(c) \cdot \sigma(u) = r'(u)(c) .$$

Thus, these residuals have the following most-general unifier:

$$\text{mgu}(r'(u_1), r'(u)) = \langle \sigma(u), \top \rangle$$

Since (σ', r') is minimal, the trees are not similar, i.e., $u_1 \not\sim_{r'} u$. There is some context $c \in \mathcal{C}_\Sigma(y)$ (possibly y) such that $c \cdot u_1 \not\sim_{r'} c \cdot u$.

Let us consider the possible most-general unifiers for such contexts c : Since the relation $\equiv_{r'}$ is a congruence, we know (by Lemma 5.11) that for every

context c with $c \cdot u_1 \in \text{dom}(r')$ the dedicated image $\text{di}(c)(u_1) \in \mathcal{T}_\Delta(y)$ exist. For every context $c' \in \text{dom}(r'(c \cdot u_1))$ holds

$$r'(c \cdot u_1)(c') \cdot \text{di}(c)(u_1) = r'(u_1)(c' \cdot c) .$$

Furthermore, it follows for every such context c and every $c' \in \text{dom}(r'(c \cdot u_1))$:

$$\begin{aligned} r'(c \cdot u_1)(c') \cdot \text{di}(c)(u_1) \cdot \sigma(u) &= r'(u_1)(c' \cdot c) \cdot \sigma(u) \\ &= r(u)(c' \cdot c) \cdot \sigma(u) \\ &= r'(u)(c' \cdot c) && (r'(u) \text{ is inessential}) \\ &= r'(c \cdot u)(c') \cdot \sigma'(c \cdot u) \end{aligned}$$

Thus, we get for the most-general unifier (w.r.t. the ordering \geq on Page 20):

$$\text{mgu}(r'(c \cdot u_1), r'(c \cdot u)) \geq \langle \text{di}(c)(u_1) \cdot \sigma(u), \sigma'(c \cdot u) \rangle$$

However, (σ', r') is unified earliest and $\sigma'(c \cdot u)$ is realizable in $r'(c \cdot u)$. Hence, the most-general unifier is in $\{\langle \text{di}(c)(u_1) \cdot \sigma(u), \top \rangle, \langle \top, \top \rangle, \langle y, y \rangle\}$.

If $u_1 \not\sim_{r'} u$, then there is some context $c \in \mathcal{C}_\Sigma(y)$ such that $c \cdot u_1 \not\sim_{r'} c \cdot u$. More precisely, we get that

$$\text{mgu}(r'(c \cdot u_1), r'(c \cdot u)) = \langle \text{di}(c)(u_1) \cdot \sigma(u), \top \rangle .$$

Otherwise, the trees would be r' -equivalent. Thus, $r'(c \cdot u)$ is inessential. Additionally, for every tree u_2 (including $c \cdot u_1$) with $\text{mgu}(r'(u_2), r'(c \cdot u)) = \langle v_2, \top \rangle$ for some ground term $v_2 \in \mathcal{T}_\Delta$, there is a tree u'_2 with $r'(u_2) = r'(u'_2) = r(u'_2)$ (Property 5.30). Thus,

$$\text{mgu}(r(u'_2), r'(c \cdot u)) = \langle v_2, \top \rangle . \tag{5.31}$$

On the other hand, we get for every context $c' \in \text{dom}(r(c \cdot u))$:

$$\begin{aligned} r(c \cdot u)(c') \cdot \sigma(c \cdot u) &= r(u)(c' \cdot c) \cdot \sigma(u) && (\text{consistent}) \\ &= r'(u)(c' \cdot c) \cdot \sigma'(u) && (\text{equivalence}) \\ &= r'(c \cdot u)(c') \cdot \sigma'(c \cdot u) && (\text{consistent}) \\ &= r'(c \cdot u)(c') && (\text{inessential}) \end{aligned}$$

It follows that $\text{mgu}(r(c \cdot u), r'(c \cdot u)) \in \{\langle \top, \top \rangle, \langle \sigma(c \cdot u), \top \rangle\}$. First, assume that $r(c \cdot u)$ is inessential and equals $r'(c \cdot u)$. Since $u_1 \equiv_r u$ and \equiv_r is a congruence, we get that also $r(c \cdot u_1) = r'(c \cdot u)$. With Equation 5.29 it follows:

$$\text{mgu}(r'(c \cdot u_1), r'(c \cdot u)) = \text{mgu}(r'(c \cdot u_1), r(c \cdot u_1)) = \langle \sigma'(c \cdot u_1), \perp \rangle$$

This contradicts the fact that (σ', r') is unified earliest because $\sigma'(c \cdot u_1)$ is realizable in $r'(c \cdot u_1)$. Thus, we get $\text{mgu}(r(c \cdot u), r'(c \cdot u)) = \langle \sigma(c \cdot u), \top \rangle$ and with Equation 5.31

$$\text{mgu}(r(c \cdot u), r(u'_2)) \geq \langle \sigma(c \cdot u), v_2 \rangle .$$

Since $\sigma(c \cdot u)$ is realizable in $r(c \cdot u)$ and (σ, r) is unified earliest, the most-general unifier must equal $\langle y, y \rangle$ or $\langle \top, \top \rangle$, i.e., the trees are r -equivalent. We get $r(c \cdot u) = r(u'_2) = r'(u'_2) = r'(u_2)$. In particular, also $r(c \cdot u) = r'(c \cdot u_1)$. Summarizing, for every tree u_2 with $\text{mgu}(r'(u_2), r'(c \cdot u)) = \langle v_2, \top \rangle$ for some ground $v_2 \in \mathcal{T}_\Delta$, we get $u_2 \equiv_{r'} c \cdot u_1$. This is a contradiction to $c \cdot u_1 \not\sim_{r'} c \cdot u$. Thus, $u_1 \sim_{r'} u$ and since (σ', r') is minimal, $u_1 \equiv_{r'} u$. However, we assumed $r'(u)$ is inessential, whereas $r'(u_1)$ is essential. — It follows that there are no trees u with $r(u)$ essential and $r'(u)$ inessential.

Finally, assume that $r(u)$ and $r'(u)$ are both essential, but different. Then, we know that for all $c \in \text{dom}(r(u))$ holds $r(u)(c) \cdot \sigma(u) = r'(u)(c) \cdot \sigma'(u)$. Since the residuals are different and essential, the most-general unifier is

$$\text{mgu}(r(u), r'(u)) = \langle \sigma(u), \sigma'(u) \rangle .$$

With Property 5.30 there is a tree $u' \equiv_{r'} u$ with $r(u') = r'(u)$. Thus,

$$\text{mgu}(r(u), r(u')) = \langle \sigma(u), \sigma'(u) \rangle .$$

However, $\sigma(u)$ is realizable in $r(u)$. — Contradiction to (σ, r) is unified earliest.

Altogether, the two residual functions r and r' are the same. Thus, the minimal partition is unique. \square

Summarizing the results from the previous sections, we obtain from the Theorems 5.7, 5.8, 5.9, 5.10, 5.12, 5.13, and 5.14:

Theorem 5.15. *For every tree function τ with path-finite partition of finite index, a minimal, path-finite partition of finite index exists, which is unique.* \square

We call this unique partition *bottom-up partition* if it exists:

Definition 5.10 (Bottom-up Partition). A partition (σ, r) of a function τ is called the *bottom-up partition* of τ if it is minimal, path-finite, and has finite index.

We write (σ_τ, r_τ) for the bottom-up partition of τ . In the following we will see that there exists a bottom-up tree transducer for a tree function if and only if this unique minimal partition (which is path-finite and has finite index) exists for τ .

Chapter 6

Myhill-Nerode Theorem

In Section 5.1, we have seen a correspondence between partitions and DBTTs: Every deterministic bottom-up tree transducer defines a partition, which is path-finite and has finite index. Not only that the properties path-finiteness and finite index are necessary to define the transformation of a DBTT, but also they are sound. More generally, we get the following Myhill-Nerode style theorem. It says that for every tree function, which is definable by a DBTT, the bottom-up partition exists. Furthermore, the DBTT described by the bottom-up partition is the unique minimal bottom-up transducer of this function.

Theorem 6.1. *The following three properties are equivalent for a partial tree function $\tau : \mathcal{T}_\Sigma \dashrightarrow \mathcal{T}_\Delta$:*

1. *There is a DBTT T with $\tau^T = \tau$.*
2. *There is a unique minimal DBTT T_{min} with $\tau^{T_{min}} = \tau$.*
3. *The bottom-up partition (σ_τ, r_τ) exists.*

We will prove this theorem in this chapter. In the first section, we show that a unique minimal transducer can be constructed for every deterministic bottom-up tree transducer. The idea of the construction is inspired by the properties of the bottom-up normal form for a partition. We get a unified earliest normal form and prove that there is a unique minimal transducer. For a given “proper” deterministic bottom-up tree transducer, i.e., one where every state produces either none or infinitely many outputs, the corresponding unique minimal transducer can be constructed in polynomial time.

In a second step, we will show that the partition $(\text{out}^{T_{min}}, r^{T_{min}})$ given by the minimal DBTT T_{min} (cf. Section 5.1) is the bottom-up partition of its transformation $\tau^{T_{min}}$. Finally, we show that (3.) implies (1.): Given a bottom-up partition, we construct a DBTT, which describes the same tree function. Additionally, we will show that this transducer is already the minimal DBTT.

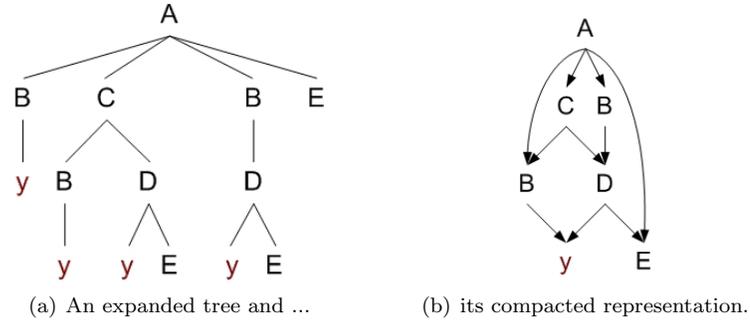


Figure 6.1: The compacted representation of a tree.

6.1 Minimization of Bottom-Up Transducers

In this section, we prove the first part of the Theorem 6.1, i.e., we show that for every deterministic bottom-up tree transducer, there is a unique equivalent DBTT in normal form. Even though the idea is similar to [EMS09], obtaining the normal form is quite different for DBTTs than for top-down transducers. Generating the “earliest” normal form for a given trim DBTT proceeds in four steps. The properties of a DBTT induced by these steps are similar to some of the properties of a bottom-up partition.

1. First, we make the transducer *proper*, i.e., we remove all output from states, which only produce finitely many different outputs. The output for such states is postponed until the root node of the input tree. This is similar to the “proper normal form” of [AU71, EM03c] (which removes states that produce finitely many outputs, using regular look-ahead).
2. We make the transducer *earliest*, i.e., every state produces output as early as possible during translation.
3. We remove pairwise *ground context unifiers* (this is a technical property, achievable in quadratic time on the transducer).
4. Last, we *minimize* in the usual way (by merging isomorphic states).

Steps (2)–(4) can be done in polynomial time, i.e., given a proper transducer, its unique unified earliest transducer is constructed in polynomial time. As a positive side effect, equivalence checking for proper transducers can be done in polynomial time. Constructing a proper transducer (Step 1) takes double-exponential time in the worst case.

Also in Steps 2 and 3, the constructions could result in an exponential blow-up if we consider trees as compositions of irreducible terms. To avoid that, we assume that terms are given in a *compacted representation*. Thereto, isomorphic subterms are shared as depicted in Figure 6.1. Considering trees in this compacted representation, the composition of two trees is only a substitution in one

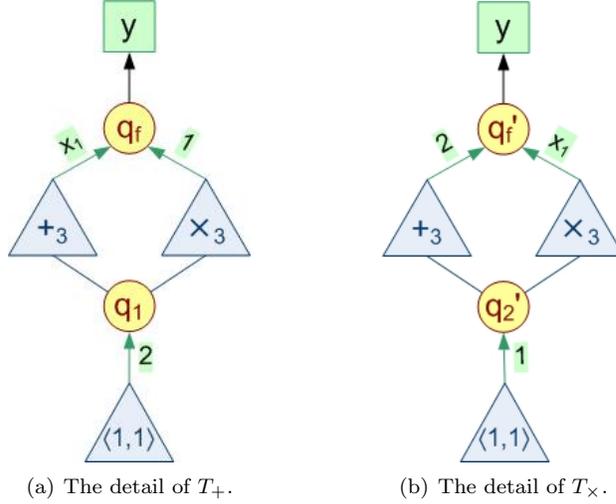


Figure 6.2: Details depending on input label $\langle 1, 1 \rangle$ of the DBTTs of Example 6.1.

occurrence of the variable y . Thus, we have no multiple copying of subterms and the size of the new terms in Steps 2 and 3 is restricted.

6.1.1 Proper Transducers

For a given trim transducer, there is not necessarily a unique minimal equivalent transducer. Recall the Example 5.4 of the quotient ring of integers modulo 3. The two intuitive partitions are not proper. For the transformation of that example, we can also define different transducers, which are equivalent and have the same size.

Example 6.1. Assume that Σ and Δ are defined as in Example 5.4. Consider the transducers $T_+ = (Q_+, \Sigma, \Delta, R_+, F_+)$ and $T_\times = (Q_\times, \Sigma, \Delta, R_\times, F_\times)$ defined below (Figure 6.2). Both describe the tree function τ of that example in Section 5.2.2. The sets of states are $Q_+ = \{q_f, q_0, q_1, q_2\}$ and $Q_\times = \{q'_f, q'_0, q'_1, q'_2\}$. The final functions are only defined for q_f and q'_f , respectively, by $F_+(q_f) = y$ and $F_\times(q'_f) = y$. Finally, the set of rules R_+ for all $a, b \in \{0, 1, 2\}$ is given by

$$\begin{aligned} \langle a, b \rangle &\mapsto q_i(j) && \text{for } i = a \times_3 b, \quad j = a +_3 b \\ +_3(q_a) &\mapsto q_f(x_1) \\ \times_3(q_a) &\mapsto q_f(a), \end{aligned}$$

whereas the set of rules R_\times for all $a, b \in \{0, 1, 2\}$ is given by

$$\begin{aligned} \langle a, b \rangle &\mapsto q_i(j) && \text{for } i = a +_3 b, \quad j = a \times_3 b \\ \times_3(q_a) &\mapsto q_f(x_1) \\ +_3(q_a) &\mapsto q_f(a). \end{aligned}$$

The transducer T_+ produces the output $\tau(+_3(p))$ for every pair p at the node labeled by p . It stores the output of $\times_3(p)$ in the state reached by p . At the root, it suffices either to pass the produced output (if the label of the root is $+_3$) or produces the correct output depending on the state (with root label \times_3). On the other side, T_\times produces the output of $\times_3(p)$ at the leaf and stores the output of $+_3(p)$ in the state. Accordingly, the transitions for the root nodes are swapped, too. In Figure 6.2, the difference of these transducers is illustrated. Both parts of the figure show the transitions, which are necessary to translate a tree containing a leaf with label $\langle 1, 1 \rangle$ for the deterministic bottom-up tree transducers T_+ and T_\times , respectively.

Both transducers are trim and describe the same transformation τ . It is not clear how a unique normal form for the transformation with less than four states could look like. \triangleleft

Considering the partitions of these DBTTs (cf. Section 5.1), we get exactly the partitions of Example 5.4, i.e., (σ_+, r_+) and $(\sigma_\times, r_\times)$, respectively. Likewise, we introduce the property “proper” for DBTTs.

Let T denote a trim transducer with set of states Q . A state $q \in Q$ is called *essential* if the set of results $\{\llbracket t \rrbracket^T \mid t \in \mathcal{L}^T(q)\}$ for input trees reaching q is infinite. Otherwise, q is called *inessential*. Note that all states of the transducers in Example 6.1 are inessential. A proper transducer postpones outputs at inessential states.

Definition 6.1 (Proper). The trim transducer T is called *proper* if every inessential state does not produce any output, i.e., is in Q_* .

For every trim bottom-up tree transducer exists an equivalent proper transducer:

Proposition 6.2. [AU71, EM03c] *For every trim DBTT T , a DBTT T' can be constructed with the following properties:*

1. T' is equivalent to T
2. T' is proper
3. $|T'| \leq \Gamma \cdot |T|$

where Γ is the sum of sizes of all outputs produced for inessential states of T .

In the worst case, an inessential state may have exponentially many outputs – even if the input alphabet has maximal rank 1. In case that both input and output alphabets have symbols of ranks greater than 1, doubly exponentially many outputs of inessential states are possible.

Proof of Proposition 6.2. Let $T = (Q, \Sigma, \Delta, R, F)$ be a trim DBTT. First, we determine the set $\text{Finite} \subseteq Q$ of inessential states of T . For that, we construct the *dependence* graph $G_T = (Q, E_T)$ of T . The set of nodes of G_T is given by the set Q of states of T , and the set E_T of edges is given by:

$$E_T = \{(q_i, q) \mid \mathbf{a}(q_1, \dots, q_k) \rightarrow q(z) \in R \text{ and } x_i \text{ occurs in } z\}$$

An edge $(q_i, q) \in E_T$ is called essential if there is a transition $\mathbf{a}(q_1, \dots, q_k) \rightarrow q(z)$ in R where z contains x_i , but is different from x_i . A strongly connected component of G_T is called essential if it contains an essential edge. Otherwise, it is called inessential. Then a state $q \in Q$ is essential iff it is reachable from some essential strongly connected component. Since the set of all strongly connected components can be computed in linear time [Tar72], also the set of states within all essential strongly connected components can be computed in linear time. Therefore, both the set of all essential states, as well as the set Finite of inessential states, can be computed in linear time.

Consider an inessential state $q \in \text{Finite}$. We partition q into new states $\langle q, z \rangle$ where $q(z)$ is a possible result for some input tree $t \in \mathcal{L}^T(q)$. If q occurs on a left-hand side of a rule as the state for the i -th argument where the state in the right-hand side is essential, a new rule is generated where q is replaced with $\langle q, z \rangle$ and the corresponding variable x_i is replaced with z . Also, the final function F' should be modified accordingly for inessential states.

Formally, let Finite denote the set of all inessential states of T . Let the set of new inessential states $Q_f \subseteq \text{Finite} \times \mathcal{T}_\Delta$ be defined as the least set satisfying the inequation:

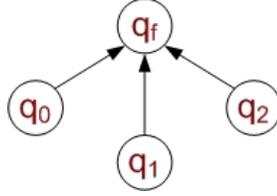
$$Q_f \supseteq \{ \langle q, z[v_1/x_{j_1}, \dots, v_r/x_{j_r}] \rangle \mid \mathbf{a}(q_1, \dots, q_m) \rightarrow q(z) \in R, q \in \text{Finite}, \\ x_{j_1}, \dots, x_{j_r} \text{ occur in } z, \langle q_{j_1}, v_1 \rangle, \dots, \langle q_{j_r}, v_r \rangle \in Q_f \}$$

Then we construct a transducer $T' = (Q', \Sigma, \Delta, R', F')$ by

- $Q' = (Q \setminus \text{Finite}) \cup Q_f$;
- Assume that $F(q) = z$. If $q \notin \text{Finite}$, then $F'(q) = z$. If $\langle q, v \rangle \in Q_f$, then $F'(\langle q, v \rangle) = z \cdot v$.
- Assume that $\mathbf{a}(q_1, \dots, q_m) \rightarrow q(z) \in R$. Then T' has transitions of the form $\mathbf{a}(q'_1, \dots, q'_m) \rightarrow q'(z')$ with the following properties:
 - If $q_i \notin \text{Finite}$, then $q'_i = q_i$. If $q_i \in \text{Finite}$, then $q'_i = \langle q_i, v_i \rangle \in Q_f$. Let z_i be equal to x_i if $q_i \notin \text{Finite}$ and $z_i = v_i$ otherwise.
 - If $q \notin \text{Finite}$, then $q' = q$ and $z' = z[z_1, \dots, z_m]$. If $q \in \text{Finite}$, then $q' = \langle q, z[z_1, \dots, z_m] \rangle$ and $z' = *$.

This deterministic bottom-up tree transducer T' is trim and proper. For the size, we observe that a transition of T is duplicated for every sequence of possible outputs of the inessential states on the left-hand side. Furthermore, x_i -nodes in right-hand sides are replaced by outputs of inessential states. Thus, the size of T' is bounded by $\Gamma \cdot |T|$, where Γ is the sum of sizes of all outputs produced for inessential states of T . \square

With this theorem we get a proper deterministic bottom-up tree transducer for the transformation of Example 5.4. Thereto, we apply the construction of the proof to one of the two transducers T_+ and T_\times of Example 6.1.

Figure 6.3: The dependency graph G_{T_+} for DBTT T_+ .

Example 6.2. Consider again the transducer T_+ of Example 6.1. The dependence graph is

$$G_{T_+} = (\{q_f, q_0, q_1, q_2\}, \{(q_0, q_f), (q_1, q_f), (q_2, q_f)\})$$

with only inessential edges (Figure 6.3). We determine that all states are inessential. The equivalent proper DBTT $T'_+ = (Q'_+, \Sigma, \Delta, R'_+, F'_+)$ has the following set of states:

$$Q'_+ = \{\langle q_f, i \rangle, \langle q_0, i \rangle, \langle q_1, 1 \rangle, \langle q_1, 2 \rangle, \langle q_2, 0 \rangle \mid 0 \leq i \leq 2\}.$$

Since every state of Q'_+ is inessential, the output is postponed to the final function, whereas the right-hand sides of the transitions R'_+ are of the form $\langle q, z \rangle(*)$. More precisely, for all pairs $\langle a, b \rangle \in \Sigma^{(0)}$, the transitions have the form

$$\langle a, b \rangle \rightarrow \langle q_i, j \rangle(*)$$

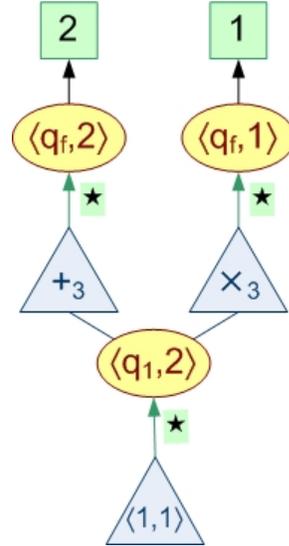
where $i = a \times_3 b$ and $j = a +_3 b$. Additionally, we have for states $\langle q_a, b \rangle \in Q'_+$ with $a \neq f$ the transitions:

$$\begin{aligned} +_3(\langle q_a, b \rangle) &\rightarrow \langle q_f, b \rangle(*) \\ \times_3(\langle q_a, b \rangle) &\rightarrow \langle q_f, a \rangle(*) \end{aligned}$$

The final function is given by

$$F'_+(\langle q_f, a \rangle) = a \quad \text{for all } a \in \{0, 1, 2\}.$$

Figure 6.4 represents the transitions of DBTT T'_+ , which are potential applied to nodes of trees with leaf labeled $\langle 1, 1 \rangle$. It is the same part as we have seen for the two transducers T_+ and T_\times in Figure 6.2. If we construct T'_\times for the second transducer T_\times of Example 6.1, we get an isomorphic transducer to T'_+ . Both deterministic bottom-up tree transducers T'_\times and T'_+ are proper and realize the transformation τ . \triangleleft

Figure 6.4: The details depending on label $\langle 1, 1 \rangle$ of the DBTT of Example 6.2.

6.1.2 Earliest Transducers

Assume that we are given a proper DBTT T . Analog as for partitions, we now want this transducer to produce the output at essential states as *early* as possible. Thereto, we compute the *greatest common suffix* of all non-ground images of contexts for a state q and produce it at q directly.

For an essential state q , let $\hat{\mathcal{Z}}(q)$ denote the set of images $z \in \hat{\mathcal{T}}_\Delta(y)$ produced for contexts of q , i.e, if q is reachable by a tree u then $\hat{\mathcal{Z}}(q) = \hat{\mathcal{Z}}(u)$ w.r.t. the partition of the DBTT (cf. Page 55 for $\hat{\mathcal{Z}}(u)$). Thus, every tree in $\hat{\mathcal{Z}}(q)$ contains an occurrence of the variable y . The *greatest common suffix* of all trees in $\hat{\mathcal{Z}}(q)$ is denoted by $\text{gcs}(q)$, i.e.,

$$\text{gcs}(q) = \bigsqcup \hat{\mathcal{Z}}(q)$$

with respect to the order \sqsubseteq on $\hat{\mathcal{T}}_\Delta(y)_\perp$ from Page 19. It equals the greatest common suffix $\text{gcs}(u)$ for all $u \in \mathcal{L}^T(q)$. In contrast to partitions, for every proper DBTT, the greatest common suffices of the states can effectively be computed:

Lemma 6.3. *For a proper DBTT T , the trees $\text{gcs}(q)$ for all essential states q of T can be computed in polynomial time.*

Proof. Assume that $z \in \mathcal{T}_\Delta(\mathcal{X}_k)$ and x_i occurs in z . Then $\text{suff}_i(z)$ denotes the largest subtree z_i of $z[y/x_i]$ with the following properties:

- y is the only variable occurring in z_i , i.e., $z_i \in \hat{\mathcal{T}}_\Delta(y)$;
- $z[y/x_i] = z' \cdot z_i$ for some z' , i.e., $z' \in \hat{\mathcal{T}}_{\Delta \cup \mathcal{X}_k \setminus \{x_i\}}(y)$.

Then the trees $\text{gcs}(q)$ are the least solution of the inequations

$$\begin{array}{ll} \text{gcs}(q_i) & \sqsupseteq \text{suff}_i(\text{gcs}(q) \cdot z), & \mathbf{a}(q_1, \dots, q_k) \rightarrow q(z) \in R \text{ and } x_i \text{ occurs in } z, \\ \text{gcs}(q) & \sqsupseteq z, & F(q) = z \text{ and } y \text{ occurs in } z. \end{array}$$

Since T is proper, this system contains inequations only for essential states q . The system has a unique least solution because the right-hand sides are monotonic. In addition, since the complete lattice $\hat{\mathcal{T}}_\Delta(y)_\perp$ satisfies the ascending chain condition, this least solution can effectively be computed. Using a standard worklist algorithm, it can be shown that each inequation is evaluated at most $\mathcal{O}(|T|)$ times. If we represent elements from $\hat{\mathcal{T}}_\Delta(y)$ as compositions of irreducible trees in the compact representation, then each right-hand side also can be evaluated in polynomial time. This proves the complexity bound stated in the proposition. \square

The greatest common suffix is different from \perp only for essential states. Since for these states the set $\hat{\mathcal{Z}}(q)$ is not empty, the least upper bound is always in $\hat{\mathcal{T}}_\Delta(y)$.

Definition 6.2. A proper bottom-up tree transducer T is called *earliest* if the greatest common suffix of every essential state q equals y .

With Lemma 6.3 we know that the $\text{gcs}(q)$ are computable in polynomial time. Thus, we can construct an earliest DBTT in polynomial time by moving these common suffices to the output of the trees reaching the state.

Theorem 6.4. *For each proper tree transducer T , a tree transducer T' can be constructed in polynomial time with the following properties:*

- T is equivalent to T'
- T' is earliest.

Proof. Let $T = (Q, \Sigma, \Delta, R, F)$ be a proper tree transducer. According to Lemma 6.3, we can compute the greatest common suffix $\text{gcs}(q)$ for every essential state q of T . The corresponding earliest transducer $T' = (Q, \Sigma, \Delta, R', F')$ has the same set of states as T as well as the same input and output alphabets, but only differs in the transition function and the final function.

Let us first construct the final function F' of the new transducer T' . Let q be a state of T' . Then $F'(q)$ is defined iff $F(q)$ is defined. If q is an inessential state, then $F'(q) = F(q)$. Now assume that q is essential and $F(q) = z$. Then the greatest common suffix $\text{gcs}(q)$ of q is a suffix of z , i.e., $z = u \cdot \text{gcs}(q)$ for some $u \in \hat{\mathcal{T}}_{\Delta}(y)$. Since we assume that $\text{gcs}(q)$ has already been output, we set $F'(q) = u$.

We now construct the transition function R' . Then $R'(\mathbf{a}, q_1 \dots q_m)$ is defined iff $R(\mathbf{a}, q_1 \dots q_m)$ is defined. Assume that $\mathbf{a}(q_1, \dots, q_m) \rightarrow q(z)$ is a transition in R . If q is inessential, then $z = *$ and R' contains the transition $\mathbf{a}(q_1, \dots, q_m) \rightarrow q(*)$ as well. Now assume q is essential. Then we construct the output of the corresponding transition in R' in two steps. First, we add the greatest common suffix corresponding to q to z , i.e., we define $\bar{z} = \text{gcs}(q) \cdot z$. Then we remove from \bar{z} the greatest common suffices of all variables occurring in (z and thus also in) \bar{z} . Let x_{i_1}, \dots, x_{i_r} be an enumeration of the variables occurring in \bar{z} . Then \bar{z} can be uniquely decomposed into:

$$\bar{z} = u[\text{gcs}(q_{i_1})/x_{i_1}, \dots, \text{gcs}(q_{i_r})/x_{i_r}]$$

where $u \in \mathcal{T}_{\Delta}(\{x_{i_1}, \dots, x_{i_r}\})$. Then R' has the transition $\mathbf{a}(q_1, \dots, q_m) \rightarrow q(u)$.

Due to the one-to-one correspondence of the final functions and transition functions, T' is trim. In order to prove the equivalence of T and T' , it suffices to verify by induction on the structure of an input tree $t \in \mathcal{T}_{\Sigma}$ and every essential state q ,

$$\llbracket t \rrbracket^T = q(z) \quad \text{iff} \quad \llbracket t \rrbracket^{T'} = q(\text{gcs}(q) \cdot z).$$

Moreover, for every inessential state of T , we have:

$$\llbracket t \rrbracket^T = q(*) \quad \text{iff} \quad \llbracket t \rrbracket^{T'} = q(*)$$

In particular, this invariant implies that T' is still proper. In order to prove that T' is earliest, it suffices to verify for every context $c \in \hat{\mathcal{T}}_{\Sigma}(y)$ of an essential state q that the following holds

$$\tau_q^T(c) = \tau_q^{T'}(c) \cdot \text{gcs}(q).$$

This invariant can again be proven by induction on the length of the context c (using the first invariant).

Given the trees $\text{gcs}(q)$, the construction can be performed in polynomial time. We need to care, however, not to expand the representation of suffices $\text{gcs}(q)$: Such an expansion could result in an exponential blow-up of the sizes of resulting trees. The factorization of output trees, which is necessary for constructing the tree u in transitions of R' reaching essential states, however, can also be performed with the compacted representation directly. \square

Example 6.3. Assume that $\Sigma = \{\mathbf{a}^{(2)}, \mathbf{b}^{(1)}, \mathbf{c}^{(0)}, \mathbf{d}^{(0)}\}$ and $\Delta = \{\mathbf{D}^{(2)}, \mathbf{E}^{(0)}\}$. Consider the proper DBTT $T = (Q, \Sigma, \Delta, R, F)$, with set of (essential) states $Q = \{q_1, q_2\}$ where the final function is $F = \{q_1 \mapsto \mathbf{D}(\mathbf{D}(y, \mathbf{E}), \mathbf{D}(y, \mathbf{E}))\}$ and the transition function R is given by:

$$\begin{array}{lll} \mathbf{a}(q_1, q_2) & \rightarrow q_1(\mathbf{D}(x_2, \mathbf{D}(x_1, \mathbf{E}))) & \mathbf{d} \rightarrow q_1(\mathbf{E}) \\ \mathbf{b}(q_2) & \rightarrow q_2(\mathbf{D}(x_1, \mathbf{D}(\mathbf{E}, \mathbf{E}, \mathbf{E}))) & \mathbf{c} \rightarrow q_2(\mathbf{E}) \end{array}$$

To compute the greatest common suffices, we consider the following inequations:

$$\begin{array}{llll} \text{gcs}(q_1) & \sqsupseteq & \text{suff}_1(\text{gcs}(q_1) \cdot \mathbf{D}(x_2, \mathbf{D}(x_1, \mathbf{E}))) & = \mathbf{D}(y, \mathbf{E}) \\ \text{gcs}(q_2) & \sqsupseteq & \text{suff}_2(\text{gcs}(q_1) \cdot \mathbf{D}(x_2, \mathbf{D}(x_1, \mathbf{E}))) & = y \\ \text{gcs}(q_2) & \sqsupseteq & \text{suff}_1(\text{gcs}(q_2) \cdot \mathbf{D}(x_1, \mathbf{D}(\mathbf{E}, \mathbf{E}, \mathbf{E}))) & = \text{gcs}(q_2) \cdot \mathbf{D}(y, \mathbf{D}(\mathbf{D}(\mathbf{E}, \mathbf{E}), \mathbf{E})) \\ \text{gcs}(q_1) & \sqsupseteq & F(q_1) & = \mathbf{D}(\mathbf{D}(y, \mathbf{E}), \mathbf{D}(y, \mathbf{E})) \end{array}$$

For q_2 , we obtain $\text{gcs}(q_2) = y$. Moreover, since

$$\mathbf{D}(y, \mathbf{E}) \sqcup \mathbf{D}(\mathbf{D}(y, \mathbf{E}), \mathbf{D}(y, \mathbf{E})) = \mathbf{D}(y, \mathbf{E}) ,$$

we have $\text{gcs}(q_1) = \mathbf{D}(y, \mathbf{E})$. The final function of the earliest DBTT for T' thus is given by $F' = \{q_1 \mapsto \mathbf{D}(y, y)\}$. In order to construct the new transition function, first consider the right-hand side for $\mathbf{a}(q_1, q_2)$ in R' where the right-hand side in T is $R(\mathbf{a}, q_1 q_2) = q_1(\mathbf{D}(x_2, \mathbf{D}(x_1, \mathbf{E})))$ (Figure 6.5). We have already computed the greatest common suffix for the state q_1 of the right-hand side (Step 1 in Figure 6.5). In the second step, we construct

$$\bar{z} = \text{gcs}(q_1) \cdot \mathbf{D}(x_2, \mathbf{D}(x_1, \mathbf{E})) = \mathbf{D}(y, \mathbf{E}) \cdot \mathbf{D}(x_2, \mathbf{D}(x_1, \mathbf{E})) = \mathbf{D}(\mathbf{D}(x_2, \mathbf{D}(x_1, \mathbf{E})), \mathbf{E}) .$$

From this tree, we remove the suffices for q_1 and q_2 at the variables x_1 and x_2 , respectively. This was the third step in the Figure 6.5. It results in the tree $u = \mathbf{D}(\mathbf{D}(x_2, x_1), \mathbf{E})$. Therefore, we obtain in Step 4 the transition

$$\mathbf{a}(q_1, q_2) \rightarrow q_1(\mathbf{D}(\mathbf{D}(x_2, x_1), \mathbf{E})) .$$

Analogously, we obtain the transitions

$$\begin{array}{l} \mathbf{d} \rightarrow q_1(\mathbf{D}(\mathbf{E}, \mathbf{E})) \\ \mathbf{b}(q_2) \rightarrow q_2(\mathbf{D}(x_1, \mathbf{D}(\mathbf{D}(\mathbf{E}, \mathbf{E}), \mathbf{E}))) \\ \mathbf{c} \rightarrow q_2(\mathbf{E}) . \end{array}$$

For this transducer T' , the greatest common suffices for all (essential) states equal y and the transducer is now earliest. \triangleleft

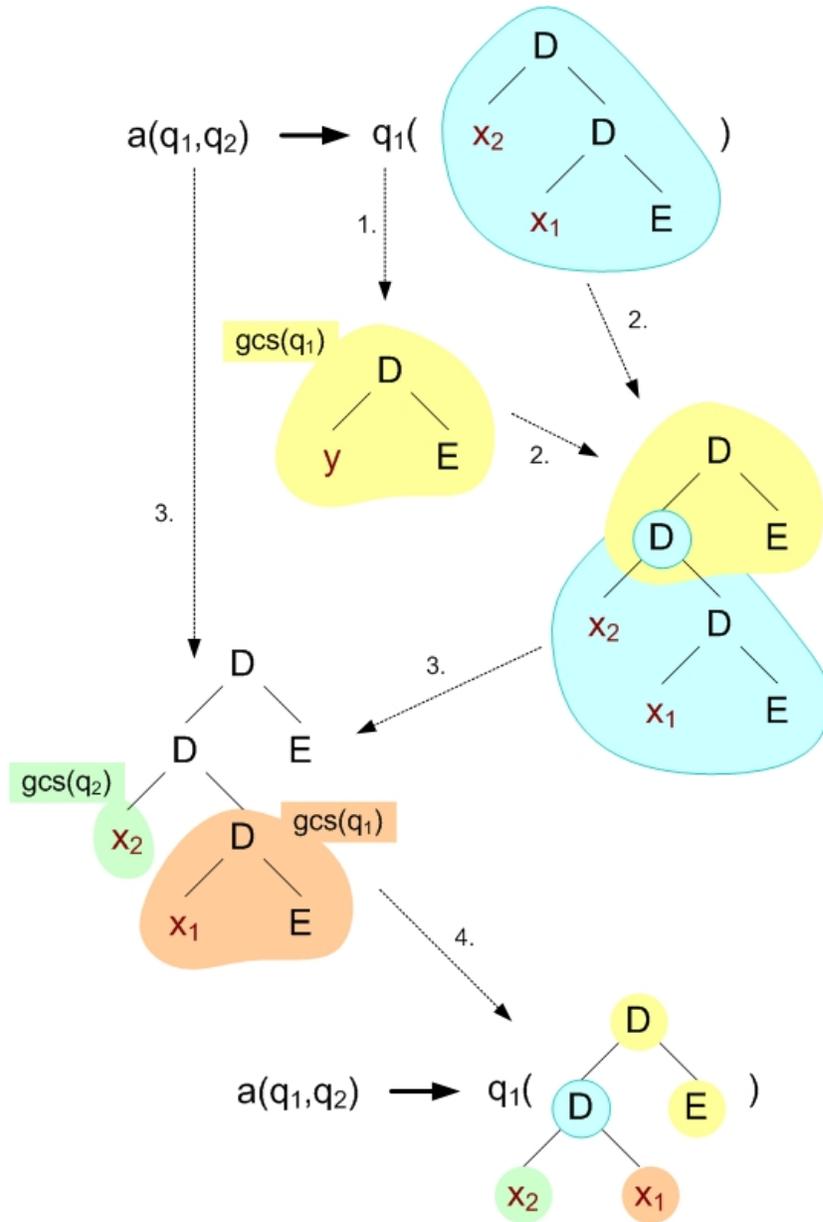


Figure 6.5: Construction of the right-hand side $R'(a, q_1q_2)$ in the earliest DBTT T' of Example 6.3.

6.1.3 Unified Transducers

For an earliest DBTT, contexts of states may disagree except for a pair of output trees.

Example 6.4. Recall the tree function τ_{mgu} of Example 5.8 (Figure 5.4). Similar to the two partitions presented in that example, we can construct different deterministic bottom-up tree transducers for τ_{mgu} with the same number of states. The input and output alphabets are

$$\Sigma = \{\mathbf{a}^{(0)}, \mathbf{b}^{(1)}, \mathbf{d}^{(1)}, \mathbf{e}^{(0)}, \mathbf{f}^{(1)}, \mathbf{g}^{(1)}\} \quad \text{and} \quad \Delta = \{\mathbf{A}^{(2)}, \mathbf{B}^{(1)}, \mathbf{D}^{(0)}, \mathbf{E}^{(0)}, \mathbf{G}^{(2)}\},$$

respectively. First, consider the earliest transducer $T_1 = (Q, \Sigma, \Delta, R_1, F)$ with set of states $Q = \{q_1, q'_1, q_2, q'_2, q_3\}$ and R_1, F given by:

$$\begin{array}{llll} \mathbf{e} & \rightarrow q_1(\mathbf{E}) & \mathbf{b}(q_1) & \rightarrow q_1(\mathbf{B}(x_1)) \\ \mathbf{a} & \rightarrow q'_1(\mathbf{E}) & \mathbf{b}(q'_1) & \rightarrow q'_1(\mathbf{B}(x_1)) & F(q'_1) = y \\ \mathbf{d}(q_1) & \rightarrow q_2(x_1) & \mathbf{g}(q_2) & \rightarrow q_3(\mathbf{G}(\mathbf{D}, y)) & F(q_2) = \mathbf{A}(y, \mathbf{E}) \\ \mathbf{d}'(q_1) & \rightarrow q'_2(x_1) & \mathbf{g}(q'_2) & \rightarrow q_3(\mathbf{G}(\mathbf{D}, \mathbf{B}(y))) & F(q'_2) = \mathbf{A}(\mathbf{B}(y), y) \\ \mathbf{f}(q'_1) & \rightarrow q'_2(\mathbf{E}) & & & F(q_3) = y \end{array}$$

This transducer is illustrated in Figure 6.6. There is an additional dotted arrow from the \mathbf{f} -labeled transition to the state q_2 , which we ignore for T_1 . We observe that the states q_2 and q'_2 are both essential and have the same contexts. The images of the context $\mathbf{g}(y)$,

$$\tau_{q_2}^{T_1}(\mathbf{g}(y)) = \mathbf{G}(\mathbf{D}, y) \quad \text{and} \quad \tau_{q'_2}^{T_1}(\mathbf{g}(y)) = \mathbf{G}(\mathbf{D}, \mathbf{B}(y))$$

differ only in the suffix $\mathbf{B}(y)$. The images of the context y are

$$\tau_{q_2}^{T_1}(y) = \mathbf{A}(y, \mathbf{E}) \quad \text{and} \quad \tau_{q'_2}^{T_1}(y) = \mathbf{A}(\mathbf{B}(y), y).$$

If y is substituted by $\mathbf{B}(\mathbf{E})$ in the image at q_2 and \mathbf{E} at q'_2 , they become equal. Thus, for each context c of q_2 , we get

$$\tau_{q_2}^{T_1}(c) \cdot \mathbf{B}(\mathbf{E}) = \tau_{q'_2}^{T_1}(c) \cdot \mathbf{E}.$$

The transition with left-hand side $\mathbf{f}(q'_1)$ produces \mathbf{E} as output and reaches state q'_2 . Thus, we can also define a transducer $T_2 = (Q, \Sigma, \Delta, R_2, F)$, which equals T_1 but $R_2(\mathbf{f}, q'_1) = q_2(\mathbf{B}(\mathbf{E}))$ (by the dotted arrow in the figure). Both transducers are earliest and have the same number of states. \triangleleft

We see that earliest deterministic bottom-up tree transducers may have unifiable states. Assume that $T = (Q, \Sigma, \Delta, R, F)$ is an earliest DBTT and that $q_1, q_2 \in Q$ are states. Every state of T defines a context function τ_q^T . We define the *most-general unifier* $\text{mgu}(q_1, q_2)$ of the states q_1, q_2 as the most-general unifier of their context functions (cf. Page 61), i.e.,

$$\text{mgu}(q_1, q_2) = \text{mgu}(\tau_{q_1}^T, \tau_{q_2}^T).$$

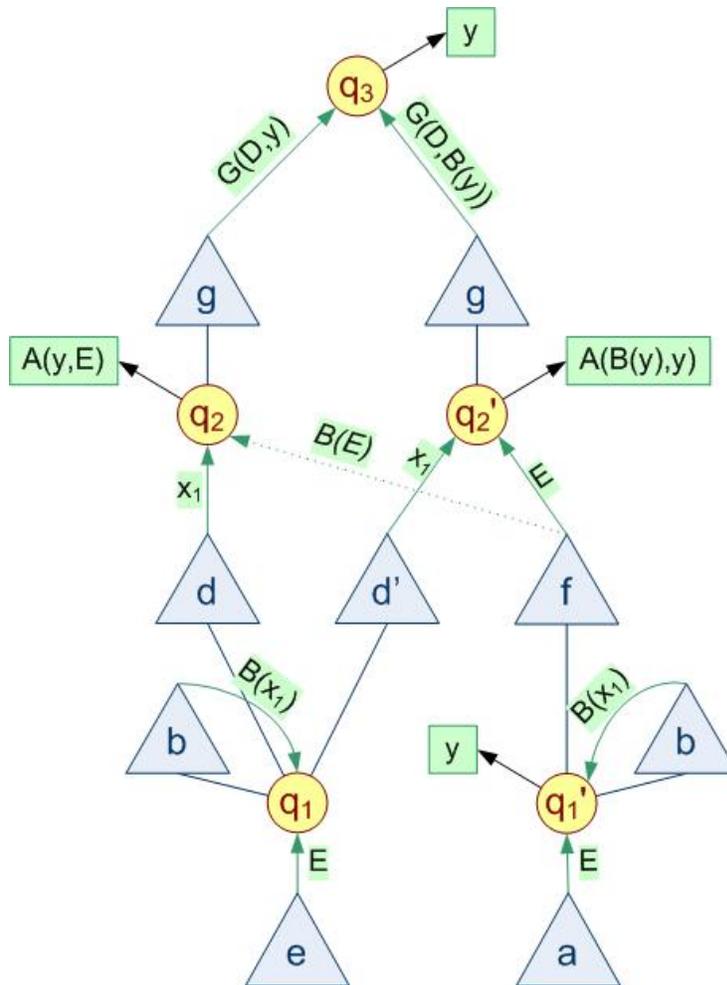


Figure 6.6: The transducers T_1 (without the $B(E)$ -labeled arrow from 'f' to q_2) and T_2 (without the E -labeled arrow from 'f' to q_2') of Example 6.4.

Note that this most-general unifier is a pair if the two states have the same set of contexts, i.e., $\mathcal{C}^T(q_1) = \mathcal{C}^T(q_2)$. Otherwise, it is $\text{mgu}(q_1, q_2) = \perp$. We call states q_1, q_2 *unifiable* if the most-general unifier is not \perp .

For unifiable states q_1 and q_2 , we observe similar properties for the most-general unifier $\text{mgu}(q_1, q_2) = \langle z_1, z_2 \rangle$ as for the most-general unifier of unifiable residuals:

- If q_i is inessential, then for every context c of q_i holds $\tau_{q_i}^T(c) \in \mathcal{T}_\Delta$. Therefore, $z_i = \top$.
- Moreover, z_1 contains y iff z_2 contains y . If both z_1 and z_2 contain y , the most-general unifier must equal $\langle y, y \rangle$, otherwise, T would not be earliest.

A ground term v is called *realizable* in a state q if v is part of the set of outputs of q . Ground terms v occurring in most-general unifiers of states are not necessarily realizable.

Definition 6.3 (Unified Earliest). The earliest DBTT T is called *unified earliest* if no ground term in most-general unifiers of states of T is realizable.

In the following, we show that for every earliest DBTT, a unified earliest DBTT can be constructed in polynomial time. For this construction, the following result is required.

Theorem 6.5. *Assume that T is an earliest deterministic bottom-up tree transducer. Then all most-general unifiers $\text{mgu}(q_1, q_2)$ can be constructed in polynomial time.*

Proof. Let the transducer $T = (Q, \Sigma, \Delta, R, F)$. We determine the greatest mapping $\mu : Q^2 \rightarrow \mathbb{D}_\Delta$ from pairs of states to candidate unifiers, which for all q, q' satisfies the following constraints:

1.

$$\mu(q, q') \leq \begin{cases} \langle \top, \top \rangle & F(q) \text{ and } F(q') \text{ are undefined} \\ \perp & F(q) \text{ defined} \Leftrightarrow F(q') \text{ undefined} \\ \text{mgu}(F(q), F(q')) & \text{otherwise} \end{cases}$$

2. Consider $\mathbf{a} \in \Sigma$ of rank m , $i \in [m]$, $q_1 \dots, q_{i-1}, q_{i+1}, \dots, q_m \in Q$. If it is not the case that the right-hand side $R(\mathbf{a}, q_1 \dots q_{i-1} q q_{i+1} \dots q_m)$ is defined iff $R(\mathbf{a}, q_1 \dots q_{i-1} q' q_{i+1} \dots q_m)$ is defined, then $\mu(q, q') \leq \perp$. Therefore, now assume that both right-hand sides $R(\mathbf{a}, q_1 \dots q_{i-1} q q_{i+1} \dots q_m)$ and $R(\mathbf{a}, q_1 \dots q_{i-1} q' q_{i+1} \dots q_m)$ are defined and equal $q_0(z)$ and $q'_0(z')$, respectively.

- If $\mu(q_0, q'_0) = \perp$, then $\mu(q, q') \leq \perp$.
- Assume $\mu(q_0, q'_0) = \langle s, s' \rangle$ where both s and s' contain y , and one of them equals y . If $s' \cdot z = z_1 \cdot u$ and $s \cdot z' = z_1 \cdot u'$ for some trees $u, u' \in \mathcal{T}_\Delta(x_i)$, then $\mu(q, q') \leq \text{mgu}(u[y/x_i], u'[y/x_i])$. If no such decompositions exist, we have $\mu(q, q') \leq \perp$.

- Assume that $\mu(q_0, q'_0) = \langle v, v' \rangle$ for ground terms $v, v' \in \mathcal{T}_\Delta$. If a variable $x_j \neq x_i$ occurs in z or z' , then we have $\mu(q, q') \leq \perp$. Now assume that $z[y/x_i], z'[y/x_i] \in \mathcal{T}_\Delta(y)$ and, that $v = z[v_i/x_i]$ and $v' = z'[v'_i/x_i]$ for some $v_i, v'_i \in \mathcal{T}_\Delta$. If no such decomposition exists, we have $\mu(q, q') \leq \perp$.
Therefore, assume that such a decomposition exists.
If $z[y/x_i], z'[y/x_i] \in \hat{\mathcal{T}}_\Delta(y)$, then $\mu(q, q') \leq \langle v_i, v'_i \rangle$.
If $z[y/x_i] \in \hat{\mathcal{T}}_\Delta(y)$ and $z' \in \mathcal{T}_\Delta$, then $\mu(q, q') \leq \langle v_i, \top \rangle$.
If $z \in \mathcal{T}_\Delta$ and $z'[y/x_i] \in \hat{\mathcal{T}}_\Delta(y)$, then $\mu(q, q') \leq \langle \top, v'_i \rangle$.
- Now assume that $\mu(q_0, q'_0) = \langle v, \top \rangle$ for a ground term v . If a variable $x_j \neq x_i$ occurs in z , then $\mu(q, q') \leq \perp$. Assume that $z[y/x_i] \in \mathcal{T}_\Delta(y)$ and $v = z[v_i/x_i]$ for some $v_i \in \mathcal{T}_\Delta$.
If no such decomposition exists, we have $\mu(q, q') \leq \perp$.
If such a decomposition exists and $z[y/x_i] \in \hat{\mathcal{T}}_\Delta(y)$, then we have $\mu(q, q') \leq \langle v_i, \top \rangle$.
- The case where $\mu(q_0, q'_0)$ equals $\langle \top, v \rangle$ for a ground term $v \in \mathcal{T}_\Delta$, is analogous.

Each constraint induced by a pair of final outputs $F(q), F(q')$ as well as each constraint induced by a matching pair of transitions

$$\begin{aligned} \mathbf{a}(q_1, \dots, q_{i-1}, q, q_{i+1}, \dots, q_m) &\rightarrow q_0(z) \text{ and} \\ \mathbf{a}(q_1, \dots, q_{i-1}, q', q_{i+1}, \dots, q_m) &\rightarrow q'_0(z') \end{aligned}$$

is monotonic w.r.t. the ordering \leq and distributes over pairwise greatest lower bounds.

By induction on the length of contexts $c \in \hat{\mathcal{T}}_\Delta(y)$, we verify that for each pair of states q, q' it holds:

- If $\tau_q(c)$ and $\tau_{q'}(c)$ are defined, then $\mu(q, q') \leq \text{mgu}(\tau_q(c), \tau_{q'}(c))$.
- If either $\tau_q(c)$ or $\tau_{q'}(c)$ is defined and not the other, then $\mu(q, q') = \perp$.

It follows that $\mu(q, q') \leq \text{mgu}(q, q')$ for all states q, q' . Since mgu is a solution of the constraint system, thus, it is the greatest solution.

The complete lattice (\mathbb{D}, \leq) of candidate unifiers (cf. Page 20) has finite height, i.e., each strictly ascending chain $d_1 \leq d_2 \leq \dots \leq d_k$ has at most length $k = 4$. Using the compacted representations of trees where isomorphic subterms are shared, this construction of the most-general unifiers works in polynomial time. \square

Assume now that we are given all most-general unifiers of an earliest DBTT T . Then we can construct a unified earliest transducer T' , which is equivalent to T . The idea of the proof is similar as for partitions, cf. Theorem 5.10. Realizable ground terms (and parts of it) are stored in new inessential states and with that are postponed. We have:

Theorem 6.6. *For each earliest DBTT T , a DBTT T' can be constructed in polynomial time with the following properties:*

- T' is equivalent to T and
- T' is unified earliest.

Proof. Let $T = (Q, \Sigma, \Delta, R, F)$ be an earliest DBTT. By Theorem 6.5, we may assume that we are given all most-general unifiers $\text{mgu}(q, q')$ for states $q, q' \in Q$. Then we construct the unified earliest transducer $T' = (Q', \Sigma, \Delta, R', F')$ in two steps.

First, we introduce new states and get a transducer T_1 . Whenever an output v of an input u at state q is produced by T , which will contribute to a ground unifier of q , then the computation on u is redirected to a new state $\langle q, v \rangle$, which memorizes v and does produce $*$ only. Instead, the output v is delayed to the images of the contexts. This implies that the new state $\langle q, v \rangle$ is inessential. Furthermore, for states q' used to evaluate subtrees of u whose outputs v' may contribute to v , further states $\langle q', v' \rangle$ should be introduced.

Formally, the DBTT $T_1 = (Q_1, \Sigma, \Delta, R_1, F_1)$ is defined as follows. We denote by $V \subseteq \mathcal{T}_\Delta \cup \{\perp\}$ the set of subterms of terms occurring as ground unifiers of states or \perp . Let Q_1 denote the set of pairs

$$Q_1 = \{\langle q, v \rangle \mid q \in Q, v \in V\}.$$

Let R_1 denote the least set of transitions, which contains for each transition $\mathbf{a}(q_1, \dots, q_m) \rightarrow q(z)$ of R the following transitions. Assume $v_1, \dots, v_m \in V$ and let $v = z[v_1, \dots, v_m]$.

- If $v \in V$, then $\mathbf{a}(\langle q_1, v_1 \rangle, \dots, \langle q_m, v_m \rangle) \rightarrow \langle q, v \rangle(*) \in R_1$.
- If $v \notin V$, then $\mathbf{a}(\langle q_1, v_1 \rangle, \dots, \langle q_m, v_m \rangle) \rightarrow \langle q, \perp \rangle(z[v'_1, \dots, v'_m]) \in R_1$ where $v'_i = x_i$ if $v_i = \perp$, and $v'_i = v_i$ otherwise.

Let $\langle q, v \rangle \in Q_1$. The final function F_1 is defined for $\langle q, v \rangle$ iff F is defined for q where $F_1(\langle q, \perp \rangle) = F(q)$ and $F_1(\langle q, v \rangle) = F(q) \cdot v$ if $v \neq \perp$. Some of the new states $\langle q, v \rangle$ of T_1 for $v \neq \perp$ may be unreachable. The unified earliest transducer $T' = (Q', \Sigma, \Delta, R', F')$, therefore, is defined as the trim DBTT equivalent to T_1 , obtained according to Proposition 4.1.

By induction on the length of contexts c and depth of input trees u , we obtain:

- $\forall \langle q, v \rangle \in Q' : c \in \mathcal{C}^T(q)$ iff $c \in \mathcal{C}^{T'}(\langle q, v \rangle)$;
- $\forall \langle q, \perp \rangle \in Q', c \in \mathcal{C}^T(q) : \tau_q^T(c) = \tau_{\langle q, \perp \rangle}^{T'}(c)$;
- $\forall \langle q, v \rangle \in Q', v \neq \perp, c \in \mathcal{C}^T(q) : \tau_q^T(c) \cdot v = \tau_{\langle q, v \rangle}^{T'}(c)$;
- $\forall u \in \mathcal{T}_\Sigma : \llbracket u \rrbracket^T = q(v)$ iff $\llbracket u \rrbracket^{T'} = \begin{cases} \langle q, v \rangle(*) & v \in V \\ \langle q, \perp \rangle(v) & v \notin V \end{cases}$

The number of outputs produced at a state q of T , which is not produced by $\langle q, \perp \rangle$ in T' , is finite (it is bounded by the number of subtrees of ground unifiers of T). Thus, if q is an essential state of T , $\langle q, \perp \rangle$ is an essential state of T' and each state $\langle q, v \rangle$ with $v \neq \perp$ is inessential. Thus, the transducer T' is *proper*. Since the images of contexts of state q w.r.t. T equal the images of contexts of $\langle q, \perp \rangle$ w.r.t. T' , and the new states $\langle q, v \rangle$ with $v \neq \perp$ are inessential, the transducer T' is still *earliest*. Let $v, v' \in \mathcal{T}_\Delta$. The most-general unifier of T' is given by

$$\begin{aligned} & - \text{mgu}^{T'}(\langle q', \perp \rangle, \langle q, \perp \rangle) = \text{mgu}^T(q', q), \\ & - \text{mgu}^{T'}(\langle q', \perp \rangle, \langle q, v \rangle) = \begin{cases} \langle v', \top \rangle & \text{mgu}^T(q', q) = \langle v', v \rangle \\ \langle v, \top \rangle & \text{mgu}^T(q', q) = \langle y, y \rangle \\ \perp & \text{otherwise} \end{cases} \\ & - \text{mgu}^{T'}(\langle q', v' \rangle, \langle q, v \rangle) = \begin{cases} \langle \top, \top \rangle & \text{mgu}^T(q', q) = \langle v', v \rangle \\ \perp & \text{otherwise} \end{cases} \end{aligned}$$

The new bottom-up tree transducer T' is *unified earliest*: Assume there is a unifier $\text{mgu}(\langle q_1, v_1 \rangle, \langle q_2, v_2 \rangle) = \langle v'_1, v'_2 \rangle$ in T' with ground $v'_1 \in \mathcal{T}_\Delta$. We show that v'_1 is not realizable in $\langle q_1, v_1 \rangle$. Since $v'_1 \neq \top$, the state $\langle q_1, v_1 \rangle$ is essential and with that, $v_1 = \perp$.

- If $v_2 = \perp$, then $\langle v'_1, v'_2 \rangle$ was a most-general unifier in T of q_1, q_2 . Consequently, $\langle q_1, v'_1 \rangle$ is a state in Q' .
- If $v_2 \neq \perp$ and $v'_2 = \top$ and $\text{mgu}^T(q_1, q_2) \neq \langle y, y \rangle$, then $\langle v'_1, v_2 \rangle$ was a most-general unifier in T for q_1, q_2 . Thus, $\langle q_1, v'_1 \rangle \in Q'$.
- If $v_2 \neq \perp$ and $v'_2 = \top$ and $\text{mgu}^T(q_1, q_2) = \langle y, y \rangle$, then $v_2 = v'_1 \in V$ and $\langle q_1, v'_1 \rangle$ is a state in Q' .

Now we show that T and T' are *equivalent*: Consider a tree $u = c \cdot u'$ with $\llbracket u' \rrbracket^T = q(v)$. If $v \neq \perp$ and $\langle q, v \rangle \in Q'$ we have:

$$\tau^T(u) = \tau_q^T(c) \cdot v = \tau_{\langle q, v \rangle}^{T'}(c) = \tau_{\langle q, v \rangle}^{T'}(c) \cdot * = \tau^{T'}(u)$$

Otherwise, we have $\llbracket u' \rrbracket^{T'} = \langle q, \perp \rangle(v)$ and:

$$\tau^T(u) = \tau_q^T(c) \cdot v = \tau_{\langle q, v \rangle}^{T'}(c) \cdot v = \tau^{T'}(u) . \quad \square$$

Example 6.5. Consider again the transducer $T_1 = (Q, \Sigma, \Delta, R_1, F)$ of Example 6.4. The most-general unifiers are

$$\text{mgu}(q_2, q'_2) = \langle \mathbf{B}(\mathbf{E}), \mathbf{E} \rangle$$

and $\text{mgu}(q, q') = \perp$, otherwise. We get the set $V = \{\mathbf{B}(\mathbf{E}), \mathbf{E}, \perp\}$ of subterms of terms occurring as ground unifiers of states or \perp . All states of $Q \times V$ are possible new states. Except from $\langle q_3, \mathbf{B}(\mathbf{E}) \rangle$ and $\langle q_3, \mathbf{E} \rangle$, all are reachable. Starting with

The unified earliest transducer $T'_1 = (Q', \Sigma, \Delta, R'_1, F')$ has the set of states

$$Q' = \{ \langle q_1, \mathbf{E} \rangle, \langle q_1, \mathbf{B}(\mathbf{E}) \rangle, \langle q_1, \perp \rangle, \langle q'_1, \mathbf{E} \rangle, \langle q'_1, \mathbf{B}(\mathbf{E}) \rangle, \langle q'_1, \perp \rangle, \\ \langle q_2, \mathbf{E} \rangle, \langle q_2, \mathbf{B}(\mathbf{E}) \rangle, \langle q_2, \perp \rangle, \langle q'_2, \mathbf{E} \rangle, \langle q'_2, \mathbf{B}(\mathbf{E}) \rangle, \langle q'_2, \perp \rangle, \langle q_3, \perp \rangle \} .$$

The transition function R'_1 is given by

$$\begin{array}{ll} \mathbf{a} \rightarrow \langle q'_1, \mathbf{E} \rangle (*) & \mathbf{b}(\langle q'_1, \mathbf{E} \rangle) \rightarrow \langle q'_1, \mathbf{B}(\mathbf{E}) \rangle (*) \\ \mathbf{e} \rightarrow \langle q_1, \mathbf{E} \rangle (*) & \mathbf{b}(\langle q_1, \mathbf{E} \rangle) \rightarrow \langle q_1, \mathbf{B}(\mathbf{E}) \rangle (*) \\ \mathbf{b}(\langle q'_1, \mathbf{B}(\mathbf{E}) \rangle) \rightarrow \langle q'_1, \perp \rangle (\mathbf{B}(\mathbf{B}(\mathbf{E}))) & \mathbf{b}(\langle q'_1, \perp \rangle) \rightarrow \langle q'_1, \perp \rangle (\mathbf{B}(x_1)) \\ \mathbf{b}(\langle q_1, \mathbf{B}(\mathbf{E}) \rangle) \rightarrow \langle q_1, \perp \rangle (\mathbf{B}(\mathbf{B}(\mathbf{E}))) & \mathbf{b}(\langle q_1, \perp \rangle) \rightarrow \langle q_1, \perp \rangle (\mathbf{B}(x_1)) \\ \mathbf{d}(\langle q_1, \mathbf{E} \rangle) \rightarrow \langle q_2, \mathbf{E} \rangle (*) & \mathbf{g}(\langle q_2, \mathbf{E} \rangle) \rightarrow \langle q_3, \perp \rangle (\mathbf{G}(\mathbf{D}, \mathbf{E})) \\ \mathbf{d}'(\langle q_1, \mathbf{E} \rangle) \rightarrow \langle q'_2, \mathbf{E} \rangle (*) & \mathbf{g}(\langle q'_2, \mathbf{E} \rangle) \rightarrow \langle q_3, \perp \rangle (\mathbf{G}(\mathbf{D}, \mathbf{B}(\mathbf{E}))) \\ \mathbf{d}(\langle q_1, \mathbf{B}(\mathbf{E}) \rangle) \rightarrow \langle q_2, \mathbf{B}(\mathbf{E}) \rangle (*) & \mathbf{g}(\langle q_2, \mathbf{B}(\mathbf{E}) \rangle) \rightarrow \langle q_3, \perp \rangle (\mathbf{G}(\mathbf{D}, \mathbf{B}(\mathbf{E}))) \\ \mathbf{d}'(\langle q_1, \mathbf{B}(\mathbf{E}) \rangle) \rightarrow \langle q'_2, \mathbf{B}(\mathbf{E}) \rangle (*) & \mathbf{g}(\langle q'_2, \mathbf{B}(\mathbf{E}) \rangle) \rightarrow \langle q_3, \perp \rangle (\mathbf{G}(\mathbf{D}, \mathbf{B}(\mathbf{B}(\mathbf{E})))) \\ \mathbf{d}(\langle q_1, \perp \rangle) \rightarrow \langle q_2, \perp \rangle (x_1) & \mathbf{g}(\langle q_2, \perp \rangle) \rightarrow \langle q_3, \perp \rangle (\mathbf{G}(\mathbf{D}, x_1)) \\ \mathbf{d}'(\langle q_1, \perp \rangle) \rightarrow \langle q'_2, \perp \rangle (x_1) & \mathbf{g}(\langle q'_2, \perp \rangle) \rightarrow \langle q_3, \perp \rangle (\mathbf{G}(\mathbf{D}, \mathbf{B}(x_1))) \\ \mathbf{f}(\langle q'_1, \mathbf{E} \rangle) \rightarrow \langle q'_2, \mathbf{E} \rangle (*) & \mathbf{f}(\langle q'_1, \mathbf{B}(\mathbf{E}) \rangle) \rightarrow \langle q'_2, \mathbf{E} \rangle (*) \\ \mathbf{f}(\langle q'_1, \perp \rangle) \rightarrow \langle q_2, \mathbf{E} \rangle (*) & \end{array}$$

and the final function is F' defined as follows:

$$\begin{array}{ll} F'(\langle q'_1, \mathbf{E} \rangle) = \mathbf{E} & F'(\langle q'_1, \mathbf{B}(\mathbf{E}) \rangle) = \mathbf{B}(\mathbf{E}) \\ F'(\langle q_2, \mathbf{E} \rangle) = \mathbf{A}(\mathbf{E}, \mathbf{E}) & F'(\langle q_2, \mathbf{B}(\mathbf{E}) \rangle) = \mathbf{A}(\mathbf{B}(\mathbf{E}), \mathbf{E}) \\ F'(\langle q'_2, \mathbf{E} \rangle) = \mathbf{A}(\mathbf{B}(\mathbf{E}), \mathbf{E}) & F'(\langle q'_2, \mathbf{B}(\mathbf{E}) \rangle) = \mathbf{A}(\mathbf{B}(\mathbf{B}(\mathbf{E})), \mathbf{B}(\mathbf{E})) \\ F'(\langle q_2, \perp \rangle) = \mathbf{A}(y, \mathbf{E}) & F'(\langle q'_2, \perp \rangle) = \mathbf{A}(\mathbf{B}(y), y) \\ F'(\langle q'_1, \perp \rangle) = y & F'(\langle q_3, \perp \rangle) = y \end{array}$$

Figure 6.7: The unified earliest DBTT T'_1 of Example 6.5.

left-hand side \mathbf{a} , we get the new transition $\mathbf{a} \rightarrow \langle q'_1, \mathbf{E} \rangle (*)$ because $\mathbf{E} \in V$. Furthermore, for the transition $\mathbf{b}(q'_1) \rightarrow q'_1(\mathbf{B}(x_1))$, we get the transition

$$\mathbf{b}(\langle q'_1, \mathbf{E} \rangle) \rightarrow \langle q'_1, \mathbf{B}(\mathbf{E}) \rangle (*)$$

because $\mathbf{B}(x_1)[\mathbf{E}/x_1] = \mathbf{B}(\mathbf{E}) \in V$. Now consider the left-hand side $\mathbf{b}(\langle q'_1, \mathbf{B}(\mathbf{E}) \rangle)$. The potential output $\mathbf{B}(x_1)[\mathbf{B}(\mathbf{E})/x_1] = \mathbf{B}(\mathbf{B}(\mathbf{E}))$ is not in V . Thus, the right-hand side should be

$$R'_1(\mathbf{b}, \langle q'_1, \mathbf{B}(\mathbf{E}) \rangle) = \langle q'_1, \perp \rangle (\mathbf{B}(\mathbf{B}(\mathbf{E}))) .$$

In addition, for the left-hand side $\mathbf{b}(\langle q'_1, \perp \rangle)$, we get the transition

$$\mathbf{b}(\langle q'_1, \perp \rangle) \rightarrow \langle q'_1, \perp \rangle (\mathbf{B}(x_1)) .$$

Now for all states $\langle q'_1, v \rangle$ for $v \in V$, we get the transition

$$\mathbf{f}(\langle q'_1, v \rangle) \rightarrow \langle q'_2, \mathbf{E} \rangle (*) .$$

The whole unified earliest transducer T'_1 is given in Figure 6.7.

On the other hand, if we start with DBTT T_2 , we get a unified earliest transducer T'_2 with the transition

$$\mathbf{f}(\langle q'_1, v \rangle) \rightarrow \langle q_2, \mathbf{B}(\mathbf{E}) \rangle (*)$$

for all $v \in \{\mathbf{B}(\mathbf{E}), \mathbf{E}, \perp\}$. However, both $\langle q'_2, \mathbf{E} \rangle$ in T'_1 and $\langle q_2, \mathbf{B}(\mathbf{E}) \rangle$ in T'_2 are inessential. Since they define the same context function, they are equivalent. Altogether, the resulting DBTTs T'_1 and T'_2 are isomorphic up to renaming. \triangleleft

6.1.4 Minimal Transducers

In the last example, we get two equivalent tree transducers, which are unified earliest. But the construction of these transducers yields in a number of redundant states. Consider for example the states $\langle q'_1, \mathbf{E} \rangle$, $\langle q'_1, \mathbf{B}(\mathbf{E}) \rangle$, and $\langle q'_1, \perp \rangle$ of transducer T'_1 in Figure 6.7. If we merge them, the transducer remains unified earliest. Moreover, the resulting DBTT is smaller.

Thus, as a last step of the normalization of deterministic bottom-up tree transducers, we merge equivalent states by preserving the properties of a unified earliest DBTT. Therefore, we consider a congruence relation on states. It is similarly defined as the congruence relation \sim_r on subtrees in Section 5.2.6. Let $T = (Q, \Sigma, \Delta, R, F)$ be a DBTT and let \sim'_T denote the smallest equivalence relation with the following properties:

- If $\text{mgu}(q, q') = \langle y, y \rangle$ or $\text{mgu}(q, q') = \langle \top, \top \rangle$ then $q \sim'_T q'$;
- Assume that $\text{mgu}(q, q_1) = \langle \top, s_1 \rangle$ for some ground term $s_1 \in \mathcal{T}_\Delta$. If for all states q_2 with $\text{mgu}(q, q_2) = \langle \top, s_2 \rangle$ for some $s_2 \neq \top$, $\text{mgu}(q_1, q_2) = \langle y, y \rangle$ holds then $q \sim'_T q_1$.

Then, the relation \sim_T is the greatest equivalence relation, which is a refinement of \sim'_T such that, $q_1 \sim_T q_2$, whenever for every symbol $\mathbf{a} \in \Sigma$ of rank m , every $i \in [m]$, and all states $p_1, \dots, p_{i-1}, p_{i+1}, \dots, p_m \in Q$, the following holds. There is a transition $\mathbf{a}(p_1, \dots, p_{i-1}, \mathbf{q}_1, p_{i+1}, \dots, p_m) \rightarrow q'_1(z_1)$ in R iff there is a transition $\mathbf{a}(p_1, \dots, p_{i-1}, \mathbf{q}_2, p_{i+1}, \dots, p_m) \rightarrow q'_2(z_2)$ in R . If such two transitions exist then $q'_1 \sim_T q'_2$. Two states q_1, q_2 are called *similar* if $q_1 \sim_T q_2$.

Definition 6.4 (Minimal). A unified earliest transducer $T = (Q, \Sigma, \Delta, R, F)$ is said to be *minimal* iff all distinct states $q_1, q_2 \in Q$ are not similar, i.e., $q_1 \not\sim_T q_2$.

Theorem 6.7. *For each unified earliest DBTT T , a unified earliest DBTT T' can be constructed in polynomial time with the following properties:*

- T' is equivalent to T
- T' is minimal
- $|T'| \leq |T|$.

Proof. Let $T = (Q, \Sigma, \Delta, R, F)$ be a unified earliest DBTT. By fixpoint iteration, we compute the equivalence relation \sim_T on Q . Now we build a transducer T' with the equivalence classes of \sim_T as states. Let $[q] = \{q' \mid q \sim_T q'\}$ the equivalence class of q . We call $[q]$ inessential if all states in $[q]$ are inessential. Otherwise, it is called essential. For each class $[q]$, we mark a representative state $p_q \in Q$, which is essential iff $[q]$ is essential.

Formally, we get $T' = (Q', \Sigma, \Delta, R', F')$ with $Q' = \{[q] \mid q \in Q\}$. The function F' is given by $F'([q]) = F(p_q)$. For R' , assume that $q_1, \dots, q_m \in Q$ are representatives of their classes and that $\mathbf{a}(q_1, \dots, q_m) \rightarrow q(z) \in R$.

- If q is essential, then $\mathbf{a}([q_1], \dots, [q_m]) \rightarrow [q](z) \in R'$.
- If p_q is inessential, then $\mathbf{a}([q_1], \dots, [q_m]) \rightarrow [q](z) \in R'$.
- Otherwise, if $\text{mgu}(q, p_q) = \langle \top, s \rangle$, then $\mathbf{a}([q_1], \dots, [q_m]) \rightarrow [q](s) \in R'$.

By induction on the depth of input trees u , we obtain for all $u \in \mathcal{T}_\Sigma$:

$$\llbracket u \rrbracket^{T'} = [q](z) \quad \text{iff} \quad \exists q' \in [q] \text{ with } \llbracket u \rrbracket^T = \begin{cases} q'(z) & \text{if } \text{mgu}(q', p_q) = \langle y, y \rangle \\ q'(z) & \text{if } \text{mgu}(q', p_q) = \langle \top, \top \rangle \\ q'(*) & \text{if } \text{mgu}(q', p_q) = \langle \top, z \rangle \end{cases}$$

In addition, by induction on the length of context c , we get for all $c \in \hat{\mathcal{T}}_\Sigma(y)$:

$$\tau_{[q]}^{T'}(c) = z \quad \text{iff} \quad \tau_{p_q}^T(c) = z$$

It follows that $\tau^T = \tau^{T'}$ and that T' is trim, proper, and earliest.

To prove that T' is unified earliest, assume that $\text{mgu}(p_1, p_2) = \langle v_1, v_2 \rangle$ for states $p_i \in Q'$ and $v_1 \in \mathcal{T}_\Delta$ is realizable in p_1 . If v_2 is ground, it follows that there are states $q_1 \in p_1$ and $q_2 \in p_2$ of T with $\text{mgu}(q_1, q_2) = \langle v_1, v_2 \rangle$ and v_1 is realizable in q_1 . If $v_2 = \top$, it follows that $q_1 \sim_T q_2$. \square

In the following, we show that the minimal transducer is unique. Before that, let us consider the construction of the minimal DBTT in an example.

Example 6.6. Consider again the unified earliest DBTT of Example 6.5, i.e., $T'_1 = (Q', \Sigma, \Delta, R'_1, F')$. The complete unified earliest transducer is presented on Page 99. Let \sim'_T and \sim_T denote the equivalence relations $\sim'_{T'_1}$ and $\sim_{T'_1}$, respectively. We compute the following most-general unifiers different to \perp :

$$\text{mgu}(\langle q_2, \perp \rangle, \langle q'_2, \perp \rangle) = \langle \mathbf{B}(\mathbf{E}), \mathbf{E} \rangle \quad (6.1)$$

$$\text{mgu}(\langle q_2, \mathbf{B}(\mathbf{E}) \rangle, \langle q'_2, \mathbf{E} \rangle) = \langle \top, \top \rangle \quad (6.2)$$

$$\text{mgu}(\langle q_2, \mathbf{B}(\mathbf{E}) \rangle, \langle q'_2, \perp \rangle) = \langle \top, \mathbf{E} \rangle \quad (6.3)$$

$$\text{mgu}(\langle q_2, \perp \rangle, \langle q'_2, \mathbf{E} \rangle) = \langle \mathbf{B}(\mathbf{E}), \top \rangle \quad (6.4)$$

Additionally, for every pair $(\langle q, \perp \rangle, \langle q, v \rangle)$ of states with $q \in \{q_1, q'_1, q_2, q'_2\}$ and $v \in \{\mathbf{E}, \mathbf{B}(\mathbf{E})\}$, the most-general unifier is

$$\text{mgu}(\langle q, \perp \rangle, \langle q, v \rangle) = \langle v, \top \rangle. \quad (6.5)$$

Every state q of the form $\langle q_2, v \rangle$ or $\langle q'_2, v \rangle$ for $v \in \{\mathbf{E}, \mathbf{B}(\mathbf{E}), \perp\}$ only occurs on the left-hand side $g(q)$ with $\langle q_3, \perp \rangle$ on the corresponding right-hand side. Thus, two states q, q' of this form are similar if $q \sim'_T q'$. With Equation 6.2, this holds for the pair $\langle q_2, \mathbf{B}(\mathbf{E}) \rangle, \langle q'_2, \mathbf{E} \rangle$, i.e., $\langle q_2, \mathbf{B}(\mathbf{E}) \rangle \sim_T \langle q'_2, \mathbf{E} \rangle$. Also, we get with Equation 6.5:

$$\langle q_2, \mathbf{E} \rangle \sim_T \langle q_2, \perp \rangle \quad \text{and} \quad \langle q'_2, \mathbf{B}(\mathbf{E}) \rangle \sim_T \langle q'_2, \perp \rangle$$

Now consider $\langle q_2, \mathbf{B}(\mathbf{E}) \rangle$. There are two essential states, which have a most-general unifier of the form $\langle \top, v \rangle$ with this inessential state: $\langle q'_2, \perp \rangle$ and $\langle q_2, \perp \rangle$ by Equations 6.3 and 6.5, respectively. Thus, $\langle q_2, \mathbf{B}(\mathbf{E}) \rangle$ is not similar to any essential state. The same follows for $\langle q'_2, \mathbf{E} \rangle$.

On the other hand, consider the states with q_1 as first component. The left-hand side $d(\langle q_1, \top \rangle)$ has $\langle q_2, \top \rangle$ on the right-hand side, whereas for $\langle q_1, \mathbf{B}(\mathbf{E}) \rangle$ it is $R(d, \langle q_1, \mathbf{B}(\mathbf{E}) \rangle) = \langle q_2, \mathbf{B}(\mathbf{E}) \rangle(*)$. Since $\langle q_2, \top \rangle$ and $\langle q_2, \mathbf{B}(\mathbf{E}) \rangle$ are not similar, also

$$\langle q_1, \top \rangle \not\sim_T \langle q_1, \mathbf{B}(\mathbf{E}) \rangle$$

follows. Analogously, we get with left-hand side $d'(\langle q_1, \top \rangle)$ and $d'(\langle q_1, \mathbf{E} \rangle)$ that these states are not similar. Additionally, for both labels d and d' follows

$$\langle q_1, \mathbf{E} \rangle \not\sim_T \langle q_1, \mathbf{B}(\mathbf{E}) \rangle.$$

Last, we have the states, with first component q'_1 . There, we get with Equation 6.5 that $\langle q'_1, \mathbf{E} \rangle \sim_T \langle q'_1, \perp \rangle$ and $\langle q'_1, \perp \rangle \sim_T \langle q'_1, \mathbf{B}(\mathbf{E}) \rangle$. Since \sim'_T is the smallest equivalence relation such that this holds, this implies

$$\langle q'_1, \mathbf{E} \rangle \sim'_T \langle q'_1, \mathbf{B}(\mathbf{E}) \rangle.$$

The minimal transducer $T_{min} = (Q_{min}, \Sigma, \Delta, R_{min}, F_{min})$ has the set of states

$$Q_{min} = \{[\langle q_1, \mathbf{E} \rangle], [\langle q_1, \mathbf{B}(\mathbf{E}) \rangle], [\langle q_1, \perp \rangle], [\langle q'_1, \perp \rangle], \\ [\langle q_2, \perp \rangle], [\langle q'_2, \mathbf{E} \rangle], [\langle q'_2, \perp \rangle], [\langle q_3, \perp \rangle]\}.$$

The transition function R_{min} is given by

$$\begin{array}{ll} \mathbf{a} \rightarrow [\langle q'_1, \perp \rangle](\mathbf{E}) & \mathbf{b}([\langle q'_1, \perp \rangle]) \rightarrow [\langle q'_1, \perp \rangle](\mathbf{B}(x_1)) \\ \mathbf{e} \rightarrow [\langle q_1, \mathbf{E} \rangle](*) & \mathbf{b}([\langle q_1, \mathbf{E} \rangle]) \rightarrow [\langle q_1, \mathbf{B}(\mathbf{E}) \rangle](*) \\ \mathbf{b}([\langle q_1, \mathbf{B}(\mathbf{E}) \rangle]) \rightarrow [\langle q_1, \perp \rangle](\mathbf{B}(\mathbf{B}(\mathbf{E}))) & \mathbf{b}([\langle q_1, \perp \rangle]) \rightarrow [\langle q_1, \perp \rangle](\mathbf{B}(x_1)) \\ \mathbf{d}([\langle q_1, \mathbf{E} \rangle]) \rightarrow [\langle q_2, \perp \rangle](\mathbf{E}) & \\ \mathbf{d}'([\langle q_1, \mathbf{E} \rangle]) \rightarrow [\langle q'_2, \mathbf{E} \rangle](*) & \mathbf{g}([\langle q'_2, \mathbf{E} \rangle]) \rightarrow [\langle q_3, \perp \rangle](\mathbf{G}(\mathbf{D}, \mathbf{B}(\mathbf{E}))) \\ \mathbf{d}'([\langle q_1, \mathbf{B}(\mathbf{E}) \rangle]) \rightarrow [\langle q'_2, \perp \rangle](\mathbf{B}(\mathbf{E})) & \\ \mathbf{d}([\langle q_1, \mathbf{B}(\mathbf{E}) \rangle]) \rightarrow [\langle q'_2, \mathbf{E} \rangle](*) & \\ \mathbf{d}([\langle q_1, \perp \rangle]) \rightarrow [\langle q_2, \perp \rangle](x_1) & \mathbf{g}([\langle q_2, \perp \rangle]) \rightarrow [\langle q_3, \perp \rangle](\mathbf{G}(\mathbf{D}, x_1)) \\ \mathbf{d}'([\langle q_1, \perp \rangle]) \rightarrow [\langle q'_2, \perp \rangle](x_1) & \mathbf{g}([\langle q'_2, \perp \rangle]) \rightarrow [\langle q_3, \perp \rangle](\mathbf{G}(\mathbf{D}, \mathbf{B}(x_1))) \\ \mathbf{f}([\langle q'_1, \perp \rangle]) \rightarrow [\langle q'_2, \mathbf{E} \rangle](*) & \end{array}$$

and the final function is F_{min} defined as follows:

$$\begin{array}{ll} F_{min}([\langle q'_2, \mathbf{E} \rangle]) = \mathbf{A}(\mathbf{B}(\mathbf{E}), \mathbf{E}) & \\ F_{min}([\langle q_2, \perp \rangle]) = \mathbf{A}(y, \mathbf{E}) & F_{min}([\langle q'_2, \perp \rangle]) = \mathbf{A}(\mathbf{B}(y), y) \\ F_{min}([\langle q'_1, \perp \rangle]) = y & F_{min}([\langle q_3, \perp \rangle]) = y \end{array}$$

Figure 6.8: The minimal DBTT T_{min} of Example 6.6.

For each of these states we have the right-hand sides

$$\begin{aligned} R'_1(\mathbf{f}, \langle q'_1, v \rangle) &= \langle q'_2, \mathbf{E} \rangle (*) \\ R'_1(\mathbf{b}, \langle q'_1, v \rangle) &= \langle q'_1, v' \rangle (z) \end{aligned}$$

for $v \in \{\mathbf{E}, \mathbf{B}(\mathbf{E}), \perp\}$ and suitable $v' \in \{\mathbf{B}(\mathbf{E}), \perp\}$ and $z \in \{*, \mathbf{B}(x_1)\}$. Thus, we get that $\langle q'_1, \mathbf{E} \rangle$, $\langle q'_1, \perp \rangle$, and $\langle q'_1, \mathbf{B}(\mathbf{E}) \rangle$ are equivalent under \sim_T . Summarizing, we get the following equivalence classes of \sim_T with more than one element:

$$\begin{aligned} [\langle q_2, \perp \rangle] &= \{\langle q_2, \perp \rangle, \langle q_2, \mathbf{E} \rangle\} \\ [\langle q'_2, \mathbf{E} \rangle] &= \{\langle q_2, \mathbf{B}(\mathbf{E}) \rangle, \langle q'_2, \mathbf{E} \rangle\} \\ [\langle q'_2, \perp \rangle] &= \{\langle q'_2, \perp \rangle, \langle q'_2, \mathbf{B}(\mathbf{E}) \rangle\} \\ [\langle q'_1, \perp \rangle] &= \{\langle q'_1, \perp \rangle, \langle q'_1, \mathbf{B}(\mathbf{E}) \rangle, \langle q'_1, \mathbf{E} \rangle\} \end{aligned}$$

To get the correct transitions, let us consider, e.g., the transition $\mathbf{a} \rightarrow \langle q'_1, \mathbf{E} \rangle (*)$ of T'_1 . The state $\langle q'_1, \mathbf{E} \rangle$ is inessential, whereas its equivalence class $[\langle q'_1, \perp \rangle]$ is essential. The most-general unifier of the two states is (cf. Equation 6.5)

$$\text{mgu}(\langle q'_1, \mathbf{E} \rangle, \langle q'_1, \perp \rangle) = \langle \top, \mathbf{E} \rangle .$$

Thus, we get the new transition $\mathbf{a} \rightarrow [\langle q'_1, \mathbf{E} \rangle](\mathbf{E})$ in T_{min} .

For sake of completeness, the whole minimal deterministic bottom-up tree transducer $T_{min} = (Q_{min}, \Sigma, \Delta, R_{min}, F_{min})$ for the tree function τ_{mgu} is presented in Figure 6.8. ◁

In the following, we show that the equivalent minimal DBTT for a given earliest unified DBTT is unique. Let the deterministic bottom-up tree transducers $T_1 = (Q_1, \Sigma, \Delta, R_1, F_1)$ and $T_2 = (Q_2, \Sigma, \Delta, R_2, F_2)$ be two equivalent minimal DBTTs, i.e., $\tau^{T_1} = \tau^{T_2}$.

For each state $q \in Q_1$, a state $p_q \in Q_2$ is said to be *related* to q if both are reached by at least one same input tree, i.e., $\exists u \in \mathcal{L}^{T_1}(q) \cap \mathcal{L}^{T_2}(p_q)$. Since q is reachable, there exists $u \in \mathcal{L}^{T_1}(q)$ and since q is meaningful, there should also exist a state $p_q \in Q_2$ with $u \in \mathcal{L}^{T_2}(p_q)$. Thus, for each state $q \in Q_1$ exists at least one related state p_q in T_2 . We will show that there exists exactly one related state in T_2 for each state $q \in Q_1$. That will give us a mapping from T_1 to T_2 .

Lemma 6.8. *Assume T_1 and T_2 are two minimal DBTTs, which are equivalent. Then for each state q of T_1 , there exists exactly one related state p_q in T_2 and the following holds:*

- Every context c of q is a context of p_q and $\tau_q^{T_1}(c) = \tau_{p_q}^{T_2}(c)$.
- $\mathcal{L}^{T_1}(q) = \mathcal{L}^{T_2}(p_q)$ and for each input tree u holds

$$\llbracket u \rrbracket^{T_1} = q(s) \quad \text{iff} \quad \llbracket u \rrbracket^{T_2} = p_q(s) .$$

Proof. If the states q and p_q are related then their sets of contexts are equal, i.e., $\mathcal{C}^{T_1}(q) = \mathcal{C}^{T_2}(p_q)$, because the transducers define the same transformation.

First assume that q is inessential. Then for every context $c \in \mathcal{C}^{T_1}(q)$, the image is ground, i.e., $\tau_q^{T_1}(c) \in \mathcal{T}_\Delta$. Assume two different states $p_1 \neq p_2$ of T_2 are related to q in T_1 . For both $i \in \{1, 2\}$, let tree $u_i \in \mathcal{L}^{T_1}(q) \cap \mathcal{L}^{T_2}(p_i)$ with $\llbracket u_i \rrbracket^{T_2} = p_i(s_i)$. Since the transducers T_1 and T_2 are equivalent, for every context $c \in \mathcal{C}^{T_1}(q)$ it holds

$$\tau_q^{T_1}(c) = \tau^{T_1}(c \cdot u_i) = \tau^{T_2}(c \cdot u_i) = \tau_{p_i}^{T_2}(c) \cdot s_i .$$

Thus, for all context $c \in \mathcal{C}^{T_1}(q)$ it holds $\tau_{p_1}^{T_2}(c) \cdot s_1 = \tau_{p_2}^{T_2}(c) \cdot s_2$. Consequently, we get for the most-general unifier:

$$\text{mgu}(p_1, p_2) \in \{ \langle s_1, s_2 \rangle, \langle s_1, \top \rangle, \langle \top, s_2 \rangle, \langle \top, \top \rangle, \langle y, y \rangle \}$$

Since T_2 is unified earliest and s_i is an output of p_i , it is not part of the most-general unifier, i.e., $\text{mgu}(p_1, p_2) \notin \{ \langle s_1, s_2 \rangle, \langle s_1, \top \rangle, \langle \top, s_2 \rangle \}$. Moreover, since T_2 is minimal, $\text{mgu}(p_1, p_2) \notin \{ \langle \top, \top \rangle, \langle y, y \rangle \}$. – Contradiction.

Now assume that q is essential. Then, there are infinitely many outputs of q . Since T_2 is finite, there should be a related state p in T_2 with infinitely many common input trees for which q produces infinitely many different outputs, i.e., the range $\text{out}^{T_1}(\mathcal{L}^{T_1}(q) \cap \mathcal{L}^{T_2}(p))$ is infinite. Consider a context $c \in \mathcal{C}^{T_1}(q)$ with $\tau_q^{T_1}(c) \in \hat{\mathcal{T}}_\Delta(y)$. We get for every $u \in \mathcal{L}^{T_1}(q) \cap \mathcal{L}^{T_2}(p)$:

$$\tau_q^{T_1}(c) \cdot \text{out}^{T_1}(u) = \tau_p^{T_2}(c) \cdot \text{out}^{T_2}(u) .$$

With Proposition 2.1, we know that the image $\tau_q^{T_1}(c)$ in T_1 is a prefix of the image $\tau_p^{T_2}(c)$ in T_2 or vice versa. W.l.o.g., assume that $\tau_q^{T_1}(c) \cdot s = \tau_p^{T_2}(c)$ for some $s \in \hat{\mathcal{T}}_\Delta(y)$. Then for infinitely many input trees u , $\text{out}^{T_1}(u) = s \cdot \text{out}^{T_2}(u)$. Therefore, also for any other context c' of q ,

$$\tau_q^{T_1}(c') \cdot s = \tau_p^{T_2}(c') .$$

Since T_2 is earliest, the context s must equal y , and therefore, $\tau_q^{T_1}(c) = \tau_p^{T_2}(c)$ for every context c of q . Furthermore, it follows $\text{out}^{T_1}(u) = \text{out}^{T_2}(u)$ for each $u \in \mathcal{L}^{T_1}(q) \cap \mathcal{L}^{T_2}(p)$. Thus, if essential states q and p are related, they produce the same output for common input trees and induce the same image for their contexts.

Now assume there exists an input tree u_1 of q with $\llbracket u_1 \rrbracket^{T_2} = p_1(s')$ and $p_1 \neq p$. Let $\llbracket u_1 \rrbracket^{T_1} = q(s)$. If p_1 is essential, there is a related state q_1 of p_1 with $\tau_{q_1}^{T_1}(c) = \tau_{p_1}^{T_2}(c)$ for every context c of p_1 . Then we have for every context c of p_1 (and thus also of q and q_1):

$$\tau_{q_1}^{T_1}(c) \cdot s' = \tau_{p_1}^{T_2}(c) \cdot s' = \tau^{T_2}(c \cdot u_1) = \tau^{T_1}(c \cdot u_1) = \tau_q^{T_1}(c) \cdot s$$

If $q_1 \neq q$, then $\langle s', s \rangle$ is a unifier of q_1 and q . The most-general unifier of q_1 and q cannot equal $\langle y, y \rangle$ or $\langle \top, \top \rangle$, since T_1 is minimal. Also, $\langle s', s \rangle$ or $\langle \top, s \rangle$ cannot

equal the most-general unifier, since s is realizable at q . Finally, $\langle s', \top \rangle$ is also no possible most-general unifier of q_1 and q because q is essential. Consequently, $q_1 = q$, and, therefore, also $p = p_1$.

Now assume that p_1 is inessential. In the following, we prove that $p_1 \sim_{T_2} p$ and therefore, $p_1 = p$ since T_2 is minimal. First, we note that for every context $c \in \mathcal{C}^{T_2}(p_1)$ holds

$$\tau_{p_1}^{T_2}(c) = \tau^{T_2}(c \cdot u_1) = \tau^{T_1}(c \cdot u_1) = \tau_q^{T_1}(c) \cdot s = \tau_p^{T_2}(c) \cdot s.$$

For a contradiction assume that $p_1 \neq p$. Then the most-general unifier of the states p_1 and p is $\langle \top, s \rangle$.

Assume that the reason for $p_1 \not\sim_{T_2} p$ is another essential state p_2 of T_2 such that $\text{mgu}(p_1, p_2) = \langle \top, s_2 \rangle$. Then the mgu of the two essential states p and p_2 is given by $\langle s, s_2 \rangle$. Let q_2 be the essential state of T_1 , which is related to p_2 . Then q_2 and p_2 also have the same sets of contexts where for all $c \in \mathcal{C}^{T_1}(q_2)$, $\tau_{q_2}^{T_1}(c) = \tau_{p_2}^{T_2}(c)$. It follows that $\text{mgu}(q, q_2) = \text{mgu}(p, p_2) = \langle s, s_2 \rangle$. Since s is realizable in q , this is a contradiction.

If $p_1 \not\sim_{T_2} p$ and there is no other essential state p_2 of T_2 with most-general unifier $\text{mgu}(p_1, p_2) = \langle \top, s_2 \rangle$ for some s_2 , then there exists a context $c' \in \mathcal{C}_\Sigma$ (not necessarily in $\mathcal{C}^{T_1}(q)$, $\mathcal{C}^{T_2}(p)$, or $\mathcal{C}^{T_2}(p_1)$) together with states q' of T_1 and distinct states p', p'_1 , and p'_2 of T_2 together with output trees z and z_1 with the properties:

- $\llbracket c' \rrbracket_q^{T_1} = q'(z)$, $\llbracket c' \rrbracket_p^{T_2} = p'(z)$, and $\llbracket c' \rrbracket_{p_1}^{T_2} = p'_1(z_1)$;
- p'_1 inessential as well, and $\text{mgu}(p', p'_1) = \langle z \cdot s, \top \rangle$;
- there exists another essential state p'_2 of T_2 such that $\text{mgu}(p'_1, p'_2) = \langle \top, s_2 \rangle$.

With the same argument as above, this implies that there exists an essential state q'_2 of T_1 such that $\text{mgu}(q', q'_2) = \langle z \cdot s, s_2 \rangle$ where $z \cdot s$ is realizable at q' — which is a contradiction.

We conclude that the related state of an essential state is also unique (and essential). It remains to prove for inessential states that the related state produces the same output (i.e., is also inessential) and induces the same image for a context. For a contradiction, assume that the related state p_q of an inessential state q is essential. Then p_q again is related to a unique essential state q' that must be different from q — which is not possible. Consequently, the related state p_q of an inessential state q must be inessential as well. Thus, the output for each input tree $u \in \mathcal{L}^{T_1}(q)$ is $*$ in T_1 and T_2 , i.e., $\llbracket u \rrbracket^{T_1} = q(*)$ and $\llbracket u \rrbracket^{T_2} = p_q(*)$. In addition, for each context $c \in \mathcal{C}^{T_1}(q)$, we have for every input tree u of q :

$$\tau_q^{T_1}(c) = \tau^{T_1}(c \cdot u) = \tau^{T_2}(c \cdot u) = \tau_{p_q}^{T_2}(c) \quad \square$$

Subsequently, if we map each state of the transducer T_1 to its related state of the second minimal transducer T_2 , we get an isomorphism. The previous Lemma 6.8 yields that minimal transducers are unique.

Theorem 6.9. *The minimal transducer T for a transformation τ is unique.*

Proof. Assume T_1 and T_2 are minimal transducers with $\tau^{T_1} = \tau^{T_2}$. We define a mapping $\varphi : Q_1 \rightarrow Q_2$ by $\varphi(q) = p_q$ where p_q is the related state of q . By the previous lemma, this mapping is well-defined and a bijective. In particular, q is the related state of p_q , too. It remains to show that φ is an isomorphism w.r.t. the transition and final functions, i.e.,

1. $F_1(q)$ is defined iff $F_2(\varphi(q))$ is defined, and if they are defined, then they are equal, i.e., $F_1(q) = F_2(\varphi(q))$, and
2. $\mathbf{a}(q_1, \dots, q_m) \rightarrow q_0(z_0) \in R_1 \iff \mathbf{a}(\varphi(q_1), \dots, \varphi(q_m)) \rightarrow \varphi(q_0)(z_0) \in R_2$.

Both follow from Lemma 6.8:

1. Because $\varphi(q)$ is the related state of q and vice versa, every context c is a context of q iff it is a context of p_q , and their images are the same. In particular for $c = y$,

$$F_1(q) = \tau_q^{T_1}(y) = \tau_{\varphi(q)}^{T_2}(y) = F_2(\varphi(q)) .$$

2. For each $i \in [m]$, consider an input tree u_i of q_i . If $\mathbf{a}(q_1, \dots, q_m) \rightarrow q_0(z)$ in R_1 , then the tree $u_0 = \mathbf{a}(u_1, \dots, u_m)$ is an input tree of q_0 , and also of the related state $\varphi(q_0)$. Therefore, there are states p_1, \dots, p_m of T_2 such that u_i are input trees of p_i , and there is a transition of the form $\mathbf{a}(p_1, \dots, p_m) \rightarrow \varphi(q_0)(z')$ in R_2 . Let $p_0 = \varphi(q_0)$. Since u_i are input trees of p_i (for all $0 \leq i \leq m$), the states q_i and p_i are related and hence, by Lemma 6.8, we get $p_i = \varphi(q_i)$. It remains to show that $z = z'$. By Lemma 6.8, $\text{out}^{T_1}(u_i) = \text{out}^{T_2}(u_i)$ holds for all $i = 0, \dots, m$. In particular, this means that $z = *$ iff $z' = *$.

Now assume that z is ground but different from $*$. Then there exists a context c of q_0 (and p_0) such that $s = \tau_{q_0}^{T_1}(c) = \tau_{p_0}^{T_2}(c)$ contains y . If z' contains an occurrence of a variable, then there are two distinct output trees z_1, z_2 at p_0 such that $s \cdot z = s \cdot z_1 = s \cdot z_2$ — which is impossible. Hence, z' must be ground as well and equal to z . It remains to consider the case where z contains occurrences of variables x_{j_1}, \dots, x_{j_r} and where all states q_{j_i}, p_{j_i} are essential. Then $z = z'$ since for all i , the images for all contexts of q_{j_i} and p_{j_i} must agree. \square

Summarizing, we obtain from Propositions 4.1, and 6.2 and Theorems 6.4, 6.6, 6.7, and 6.9:

Theorem 6.10. *For each DBTT T an equivalent minimal transducer can be constructed, which is unique up to renaming of states. If the DBTT T is already proper, the construction can be performed in polynomial time.* \square

This proves the first part of the Myhill-Nerode theorem for DBTTs (cf. Page 83), i.e., (1.) implies (2.).

6.2 From Minimal Transducers to Partitions

In Section 5.1, we defined the partition (out^T, r^T) of a DBTT T and we observe that the partition of any trim DBTT T

- is a partition of the tree function τ^T of the transducer (Corollary 5.3),
- is path-finite (Theorem 5.6), and
- its equivalence relation $\equiv_{r,T}$ has finite index (Lemma 5.4).

Now let $T = (Q, \Sigma, \Delta, R, F)$ be a minimal DBTT, which describes the tree function τ . We show that the partition of T is the bottom-up partition of its tree function τ , i.e., $(\text{out}^T, r^T) = (\sigma_\tau, r_\tau)$. This leads to the next part of Theorem 6.1. It proves that (2.) implies (3.).

First, we show that for all trees $u, u' \in \text{dom}(\text{out}^T)$ it holds $u \equiv_{r,T} u'$ iff u and u' reach the same state in T , i.e., $\exists q \in Q$ with $u, u' \in \mathcal{L}^T(q)$. Thereto, we define the equivalence relation \equiv_T on the set of all input trees of states of T , i.e., $\bigcup\{\mathcal{L}^T(q) \mid q \in Q\}$, as follows:

$$u \equiv_T u' \quad \text{if } \exists q \in Q : u, u' \in \mathcal{L}^T(q)$$

It is the same equivalence relation as $\equiv_{r,T}$:

Lemma 6.11. *Assume $T = (Q, \Sigma, \Delta, R, F)$ be a minimal deterministic bottom-up tree transducer. Then $\equiv_{r,T}$ equals \equiv_T .*

Proof. If $u \equiv_T u'$, there is a state q such that $u, u' \in \mathcal{L}^T(q)$. For their residuals, we know $r^T(u) = \tau_q^T = r^T(u')$. Thus, $u \equiv_{r,T} u'$.

Now assume $u \equiv_{r,T} u'$, but $u \not\equiv_T u'$. There are two states $q \neq q'$ with $u \in \mathcal{L}^T(q)$ and $u' \in \mathcal{L}^T(q')$. Since $u \equiv_{r,T} u'$, we get

$$\tau_q^T = r^T(u) = r^T(u') = \tau_{q'}^T.$$

The most-general unifier of q and q' is $\text{mgu}(q, q') \in \{\langle y, y \rangle, \langle \top, \top \rangle\}$. If q and q' are not similar (cf. Page 101), i.e., $q \not\sim q'$, there must be a context c such that $\llbracket c \rrbracket_q^T = q_1(z_1)$ and $\llbracket c \rrbracket_{q'}^T = q'_1(z'_1)$ with $\text{mgu}(q_1, q'_1) \in \{\langle v_1, \top \rangle, \langle \top, v_2 \rangle\}$ for some $v_1, v_2 \in \mathcal{T}_\Delta$. W.l.o.g. assume $\text{mgu}(q_1, q'_1) = \langle v_1, \top \rangle$. For each context $c_1 \in \mathcal{C}^T(q_1)$ of q_1 we know that c_1 is also a context of q'_1 . Let $\tau_{q_1}^T(c_1) = z_2$ and $\tau_{q'_1}^T(c_1) = z'_2$. Since $\text{mgu}(q, q') \in \{\langle y, y \rangle, \langle \top, \top \rangle\}$, we get

$$z_2 \cdot z_1 = \tau_q^T(c_1 \cdot c) = \tau_{q'}^T(c_1 \cdot c) = z'_2 \cdot z'_1.$$

On the other hand, since $\text{mgu}(q_1, q'_1) = \langle v_1, \top \rangle$,

$$z_2 \cdot v_1 = z'_2.$$

If $v_1 = z_1$, then v_1 is realizable in q_1 , which is a contradiction to the assumption that T is unified earliest. Otherwise, Proposition 2.1 provides $v_1 \in \hat{\mathcal{T}}_\Delta(y)$ or $z_2 = z'_2 \in \mathcal{T}_\Delta$. Both contradict the definition of the most-general unifier. It follows: $u \equiv_{r,T} u' \implies u \equiv_T u'$. \square

The equivalence relation of the partition is the same as the equivalence relation defined by the states of the DBTT. Now we show that the partition of the minimal DBTT fulfills similar properties as the transducer, i.e., the partition is also trim, proper, earliest, unified earliest, and minimal. Furthermore, the equivalence relation is already a congruence. Consequently, the partition is a bottom-up partition. Since the bottom-up partition of a transformation is unique, it is the unique bottom-up partition of the transformation defined by the transducer.

Theorem 6.12. *Assume T is a minimal DBTT. Then its partition (out^T, r^T) is the bottom-up partition of τ^T .*

Proof. Let $T = (Q, \Sigma, \Delta, R, F)$. By Corollary 5.3, Lemma 5.4, and Theorem 5.6, we know that (out^T, r^T) is a path-finite partition of τ^T with finite index. We will show that the properties trim, proper, earliest, and unified earliest are carried over from T to the partition. We will furthermore prove that \equiv_T is a congruence, which implies that \equiv_{r^T} is a congruence (Lemma 6.11) and the partition is minimal. For each $u \in \text{dom}(\text{out}^T)$ we denote q_u as the state reached by u in T , i.e., $u \in \mathcal{L}^T(q_u)$. Let $u \in \text{dom}(\text{out}^T)$. With Lemma 6.11 we get

$$[u]_{r^T} = \mathcal{L}^T(q_u) . \quad (6.6)$$

With that, we prove that (out^T, r^T) is the bottom-up partition of τ^T :

- It is trim: Since T is trim, we get the following equivalence:

$$\begin{aligned} r^T(u)(\mathcal{C}_\Sigma(y)) \subseteq \mathcal{T}_\Delta &\Leftrightarrow \tau_{q_u}^T(\mathcal{C}_\Sigma(y)) \subseteq \mathcal{T}_\Delta \\ &\Leftrightarrow \llbracket u \rrbracket^T = q_u(*) && T \text{ is trim} \\ &\Leftrightarrow \text{out}^T(u) = * \end{aligned}$$

- The partition is proper: With Equation 6.6, we get

$$\text{out}^T([u]_{r^T}) = \text{out}^T(\mathcal{L}^T(q_u)) .$$

Thus, q_u is essential iff $r^T(u)$ is essential. Since T is proper, it follows that $\text{out}^T([u]_{r^T})$ is either infinite or only contains $*$.

- The partition (out^T, r^T) is earliest: Assume $r^T(u)$ is essential. Then we know that $\text{gcs}(q_u) = \text{gcs}(u)$. Since T is earliest, the greatest common suffix equals y . On the other hand, if $r^T(u)$ is inessential, the greatest common suffix is defined as \perp .
- It is unified earliest: Assume there is a tree $u' \in \text{dom}(\text{out}^T)$ such that the most-general unifier of $r^T(u)$ and $r^T(u')$ equals $\langle v, z \rangle$ with v is a realizable ground term. Assume $\text{out}^T(u) = v$. Then $\langle v, z \rangle$ is also the most-general unifier of q_u and $q_{u'}$ (by definition of the most-general unifier of states). Furthermore, $\text{out}^T(u) = v$ and $u \in \mathcal{L}^T(q_u)$ imply, v is realizable in q_u . – Contradiction.

- The relation $\equiv_{r,T}$ is a congruence: It suffices (with Lemma 6.11) to prove that \equiv_T is a congruence relation. If $u \equiv_T u'$ then both trees u and u' reach the same state q_u . Thus, for every context $c \in \mathcal{C}_\Sigma$, the tree $c \cdot u$ in $\text{dom}(\text{out}^T)$ reaches the state $q_{c \cdot u}$ iff the tree $c \cdot u'$ reaches the same state. It follows that $c \cdot u \equiv_T c \cdot u'$.
- Finally, the partition is minimal: Let $u, u' \in \text{dom}(\text{out}^T)$. By definition of the most-general unifier of states, it holds:

$$\text{mgu}(u, u') = \text{mgu}(q_u, q_{u'})$$

It follows that

$$u \sim'_{r,T} u' \quad \text{iff} \quad q_u \sim'_T q_{u'} . \quad (6.7)$$

Furthermore, we get for all trees $u_1, u_2 \in \text{dom}(\text{out}^T)$ the following:

$$\begin{aligned} u_1 \equiv_{r,T} u_2 &\Leftrightarrow u_1 \equiv_T u_2 && \text{Lemma 6.11} \\ &\Leftrightarrow q_{u_1} = q_{u_2} && \text{Definition of } \equiv_T \\ &\Leftrightarrow q_{u_1} \sim_T q_{u_2} && \text{Definition of } \sim_T \quad (6.8) \\ &\Rightarrow q_{u_1} \sim'_{r,T} q_{u_2} && \text{Definition of } \sim'_{r,T} \quad (6.9) \end{aligned}$$

Now we know that $\equiv_{r,T}$ is a congruence and that it is a refinement of $\sim'_{r,T}$ (Equations 6.7 and 6.9). Assume that $\equiv_{r,T}$ is not the greatest congruence, which is a refinement of $\sim'_{r,T}$. Then there are trees $u, u' \in \text{dom}(\text{out}^T)$ with $u \sim_{r,T} u'$ but $u \not\equiv_{r,T} u'$. Since $\sim_{r,T}$ is a congruence, it holds for every context c with $c \cdot u \in \text{dom}(\text{out}^T)$ that $c \cdot u' \in \text{dom}(\text{out}^T)$ and $c \cdot u \sim_{r,T} c \cdot u'$. Thus, we get for all such contexts $c \cdot u \sim'_{r,T} c \cdot u'$ by Equation 6.7. Consequently, $u \sim_T u'$ and with Equation 6.8 also $u \equiv_{r,T} u'$. – Contradiction. We get that $\equiv_{r,T}$ is the greatest congruence $\sim_{r,T}$, which is a refinement of $\sim'_{r,T}$. \square

Finally, we get the result that for every transformation, which is definable by a deterministic bottom-up tree transducer, the bottom-up partition exists and is given by the unique minimal DBTT, which can be constructed (given an arbitrary DBTT of the transformation). The following example illustrates this on the tree function τ_{mgu} of Example 5.8.

Example 6.7. Both in Section 5.2.6 and Section 6.1.4 we considered in the examples the tree function τ_{mgu} . In Example 5.12 we defined the unique minimal partition (σ', r') of τ_{mgu} , i.e. the bottom-up partition of τ_{mgu} (Figure 5.7). In addition, the unique minimal DBTT T_{min} of τ_{mgu} (cf. Figure 6.8) is constructed in Example 6.6.

If we deduce the partition $(\text{out}^{T_{\text{min}}}, r^{T_{\text{min}}})$ of the DBTT, we get exactly the bottom-up partition (σ', r') . For instance, the state $[\langle q_2, \perp \rangle]$ is reached by the tree $\mathbf{d}(\mathbf{e})$. The transitions

$$\mathbf{e} \rightarrow [\langle q_1, \mathbf{E} \rangle](*) \quad \text{and} \quad \mathbf{d}([\langle q_1, \mathbf{E} \rangle]) \rightarrow [\langle q_2, \perp \rangle](\mathbf{E})$$

give us $\text{out}^{T_{min}}(\mathbf{d}(\mathbf{e})) = \mathbf{E}$. Furthermore, since $F_{min}([\langle q_2, \perp \rangle]) = \mathbf{A}(y, \mathbf{E})$, we get that $r^{T_{min}}(\mathbf{d}(\mathbf{e}))(y) = \mathbf{A}(y, \mathbf{E})$. Additionally, there is the transition:

$$\mathbf{g}([\langle q_2, \perp \rangle]) \rightarrow [\langle q_3, \perp \rangle](\mathbf{G}(\mathbf{D}, x_1))$$

Since $F_{min}([\langle q_3, \perp \rangle]) = y$, it follows

$$r^{T_{min}}(\mathbf{d}(\mathbf{e}))(\mathbf{g}(y)) = y \cdot \mathbf{G}(\mathbf{D}, y) = \mathbf{G}(\mathbf{D}, y) .$$

This corresponds to the residual $r'(\mathbf{d}(\mathbf{e}))$. ◁

6.3 From Partitions to Minimal Transducers

In this section, we prove the last part of the Theorem 6.1. We have already shown that for every DBTT there is a minimal DBTT and that the partition of a minimal DBTT is the bottom-up partition of its tree function, i.e., we know (1.) \Rightarrow (2.) \Rightarrow (3.). Since every minimal DBTT is a deterministic bottom-up tree transducer, we also have (2.) \Rightarrow (1.). It remains to prove that for every bottom-up partition a DBTT exists, which describes the same tree function. Here, we construct such a transducer and, more specific, we will show that it is already the minimal DBTT.

In the following, let (σ_τ, r_τ) be the bottom-up partition of a tree function $\tau : \mathcal{T}_\Sigma \dashrightarrow \mathcal{T}_\Delta$. Our goal is to define a DBTT T_τ , which describes τ . Thereto, we have to build transitions of the form

$$\mathbf{a}(q_1, \dots, q_k) \rightarrow q(z)$$

where z is a tree with variables x_1, \dots, x_k . Assume $\mathbf{a}(u_1, \dots, u_k)$ is a subtree and this transition would be applied to the root, i.e., for $i = 1, \dots, k$ the subtree u_i reaches the state q_i . We build the right-hand side by replacing dedicated subtrees of the form $\sigma_\tau(u_i)$ by variables x_i , such that

$$z[\sigma_\tau(u_1), \dots, \sigma_\tau(u_k)] = \sigma_\tau(\mathbf{a}(u_1, \dots, u_k))$$

holds. The correct trees $z_i \in \mathcal{T}_\Delta(y)$ such that $z_i[\sigma_\tau(u_i)] = \sigma_\tau(\mathbf{a}(u_1, \dots, u_k))$ holds, are given by the dedicated images (cf. Definition 5.8):

$$z_i = \text{di}(u_i)(\mathbf{a}(u_1, \dots, u_{i-1}, y, u_{i+1}, \dots, u_k))$$

Subsequently, we define z as the greatest common prefix of z_1, \dots, z_k .

Definition 6.5 (Greatest Common Prefix). Let $z_1, \dots, z_k \in \mathcal{T}_\Sigma(y)$. The *greatest common prefix* of z_1, \dots, z_k is

$$\text{gcp}(z_1, \dots, z_k) = \begin{cases} \mathbf{a}(z'_1, \dots, z'_n) & \forall j \leq k : z_j = \mathbf{a}(z_{j,1}, \dots, z_{j,n}) \wedge \forall i \leq n : \\ & z'_i = \text{gcp}(z_{1,i}, \dots, z_{k,i}) \text{ are defined} \\ x_i & z_i = y \wedge \forall j, j' \in [k] \setminus \{i\} : \\ & z_j = z_{j'} \in \mathcal{T}_\Sigma(y) \setminus \{y\} \\ \text{undefined} & \text{otherwise} \end{cases}$$

The greatest common prefix is a tree containing variables x_1, \dots, x_k . Let p be a y -path of z_i , i.e., $p \in \text{Paths}_y(z_i)$. The trees $z_j[p]$, to which p refers in the other trees z_j (for $j \neq i$), are the same if the greatest common prefix is defined. Furthermore, p refers to the variable x_i in the greatest common prefix. Vice versa, if a path p in the greatest common prefix refers to a variable x_i then p is a y -path in the corresponding tree z_i . With the definition of the greatest common prefix, we now define the bottom-up tree transducer of τ using its bottom-up partition (σ_τ, r_τ) :

Definition 6.6. The bottom-up tree transducer T_τ is defined by

$$T_\tau = (\text{dom}(\tau) / \equiv_{r_\tau}, \Sigma, \Delta, R_\tau, F_\tau)$$

where the right-hand sides of R_τ are defined as follows:

- For all $\mathbf{b} \in \Sigma^{(0)}$ with $\mathbf{b} \in \text{dom}(\sigma_\tau)$: $R_\tau(\mathbf{b}) = [\mathbf{b}]_{r_\tau}(\sigma_\tau(\mathbf{b}))$.
- For all $k > 0$, $\mathbf{a} \in \Sigma^{(k)}$, $u_1, \dots, u_k \in \text{dom}(r_\tau)$: If $\mathbf{a}(u_1, \dots, u_k) \in \text{dom}(\sigma_\tau)$ then

$$R_\tau(\mathbf{a}, [u_1]_{r_\tau} \dots [u_k]_{r_\tau}) = [\mathbf{a}(u_1, \dots, u_k)]_{r_\tau}(\text{gcp}(z_1, \dots, z_k))$$

where for all $1 \leq i \leq k$

$$z_i = \text{di}(u_i)(\mathbf{a}(u_1, \dots, u_{i-1}, y, u_{i+1}, \dots, u_k)) .$$

The function F_τ is defined for all $u \in \text{dom}(\sigma_\tau)$ with $y \in \text{dom}(r_\tau(u))$ by

$$F_\tau([u]_{r_\tau}) = r_\tau(u)(y) .$$

In the following, we prove that this DBTT is well-defined and correct (Theorems 6.19 and 6.21). First, we show that the greatest common prefix of the dedicated images exists. Thereto, consider a tree $\mathbf{a}(u_1, \dots, u_k) \in \text{dom}(\sigma_\tau)$ and the corresponding dedicated images z_1, \dots, z_k as defined in the previous definition. The following two lemmas show that y -paths of different trees z_i and z_j are no prefixes of each other. This result is summarized in Corollary 6.15. With this property and the fact that for each $i \in [k]$ holds

$$z_i \cdot \sigma_\tau(u_i) = \sigma_\tau(\mathbf{a}(u_1, \dots, u_k)) ,$$

we get that the greatest common prefix of z_1, \dots, z_k is defined (Lemma 6.16).

After that, we prove that the choice of the right-hand side is independent of the representatives of the states on the left-hand side (Lemma 6.18). Roughly, this follows because trees are in the same equivalent class if they have the same residual and, subsequently, they have the same dedicated images. Finally, we show that the DBTT describes τ (Theorem 6.21). Additionally, we prove that the transducer T_τ is already the minimal DBTT for this tree function (Theorem 6.23).

First, we show that y -paths in different images of related contexts w.r.t. subtrees of the same node in an input tree of τ , are no prefixes of each other. Note that we use the enhanced definition (cf. Page 20) of the partial order \sqsupseteq on trees in $\mathcal{T}_\Sigma(y)$. For two ground trees $t, t' \in \mathcal{T}_\Sigma$, $t \sqsupseteq t'$ means that t occurs as a subtree in t' .

Lemma 6.13. *Assume $c \cdot \mathbf{a}(u_1, \dots, u_k) \in \text{dom}(\tau)$. For all $i \in [k]$ let*

$$r_\tau(u_i)(c \cdot \mathbf{a}(u_1, \dots, u_{i-1}, y, u_{i+1}, \dots, u_k)) = z_i .$$

Then, it holds for all y -paths $p_{i_1} \in \text{Paths}_y(z_{i_1})$, $p_{i_2} \in \text{Paths}_y(z_{i_2})$, for some $i_1, i_2 \in [k]$ that p_{i_2} is no proper prefix of p_{i_1} .

Proof. Let $i_1, i_2 \in [k]$ and $p_{i_1} \in \text{Paths}_y(z_{i_1})$, $p_{i_2} \in \text{Paths}_y(z_{i_2})$. Assume $i_1 = i_2$. Since y -nodes have rank 0, p_{i_2} is no proper prefix of p_{i_1} . Thus, assume $i_1 \neq i_2$. Let $c_{\mathbf{a}}$ be the tree $c \cdot \mathbf{a}(u_1, \dots, u_k)$ where u_{i_1} and u_{i_2} are replaced by x_1 and x_2 , respectively, i.e.:

$$c_{\mathbf{a}} = \begin{cases} c \cdot \mathbf{a}(u_1, \dots, u_{i_1-1}, x_1, u_{i_1+1}, \dots, u_{i_2-1}, x_2, u_{i_2+1}, \dots, u_k) & i_1 < i_2 \\ c \cdot \mathbf{a}(u_1, \dots, u_{i_2-1}, x_2, u_{i_2+1}, \dots, u_{i_1-1}, x_1, u_{i_1+1}, \dots, u_k) & i_1 > i_2 \end{cases}$$

Under the assumption that p_{i_2} is a proper prefix of p_{i_1} , we show that σ_τ is not path-finite. Thereto, we prove that the residual $r_\tau(u_{i_2})$ maps an infinite chain of input contexts $c_2^{(1)}, c_2^{(2)}, \dots$ to images $z_2^{(1)}, z_2^{(2)}, \dots$, respectively, such that the input contexts have all the same y -path p . But there are y -paths $p^{(i)}$ of $z_2^{(i)}$ such that $p^{(1)}$ is a proper prefix of $p^{(2)}$, which is a proper prefix of $p^{(3)}$ and so on. First, we show that there is a chain of pairs

$$(u_1^{(1)}, u_2^{(1)}), (u_1^{(2)}, u_2^{(2)}), \dots$$

such that $\sigma_\tau(u_1^{(1)})$ is a proper subtree of $\sigma_\tau(u_2^{(1)})$, which is a proper subtree of $\sigma_\tau(u_1^{(2)})$ and so forth. Furthermore, every second element $u_2^{(i)}$ in this chain is r_τ -equivalent to u_{i_2} and the other elements $u_1^{(i)}$ are r_τ -equivalent to u_{i_1} .

We inductively construct such a chain of pairs $(u_1^{(i)}, u_2^{(i)})$. Thereto, let $P(u_1^{(i)}, u_2^{(i)})$ define the property of $u_1^{(i)}$ and $u_2^{(i)}$ as informally described above. Formally, we define for trees $u_1, u_2 \in \text{dom}(\sigma_\tau)$:

$$P(u_1, u_2) = \forall j \in [2] : u_j \equiv_{r_\tau} u_{i_j} \tag{P_1}$$

$$\wedge c_{\mathbf{a}}[u_1, u_2] \in \text{dom}(\tau) \tag{P_2}$$

$$\wedge r_\tau(u_1)(c_{\mathbf{a}}[y, u_2]) = z_1 \in \hat{\mathcal{T}}_\Delta(y) \tag{P_3}$$

$$\wedge r_\tau(u_2)(c_{\mathbf{a}}[u_1, y]) = z_2 \in \hat{\mathcal{T}}_\Delta(y) \tag{P_4}$$

$$\wedge \exists v \in \hat{\mathcal{T}}_\Delta(y) \setminus \{y\} : \sigma_\tau(u_2) = v \cdot \sigma_\tau(u_1) \tag{P_5}$$

$$\wedge \exists p \in \text{Paths}_y(z_2) : z_1[p] = v \tag{P_6}$$

Assume there is a set $\{(u_1^{(i)}, u_2^{(i)}) \mid i \in \mathbb{N}\}$ such that for all $i \in \mathbb{N}$ holds $P(u_1^{(i)}, u_2^{(i)})$. We denote by $p^{(i)}$ the path $p \in \text{Paths}_y(r_\tau(u_2^{(i)})(c_{\mathbf{a}}[u_1^{(i)}, y]))$, which

exists according to P_6 for $(u_1^{(i)}, u_2^{(i)})$. In addition, assume that for all $i \in \mathbb{N}$ holds: $p^{(i)}$ is a proper prefix of $p^{(i+1)}$. Then, we get an infinite set of equivalent trees $u_2^{(i)} \equiv_{r_\tau} u_{i_2}$. All these trees have the same residual, hence, for all $i \in \mathbb{N}$ holds $c_a[u_1^{(i)}, y] \in \text{dom}(r_\tau(u_{i_2}))$ (it follows by P_4 for $(u_1^{(i)}, u_2^{(i)})$). The path p_{in} referring to the node labeled x_2 in c_a is a y -path in each input context $c_a[u_1^{(i)}, y]$. However, the set of y -paths in the output (as considered in the definition of path-finiteness) contains all $p^{(i)}$:

$$\{p^{(i)} \mid i \in \mathbb{N}\} \subseteq \bigcup \{\text{Paths}_y(r_\tau(u_2)(c)) \mid c \in \text{dom}(r_\tau(u_2)) \wedge p \in \text{Paths}_y(c)\}$$

Thus, this set is infinite and r_τ is not path-finite. — Contradiction. It follows that p_{i_2} is no proper prefix of p_{i_1} .

Now we prove by induction that such a set $\{(u_1^{(i)}, u_2^{(i)}) \mid i \in \mathbb{N}\}$ exists if p_{i_2} is a proper prefix of p_{i_1} . More precisely, we prove the following:

If p_{i_2} is a proper prefix of p_{i_1} , then there is a set $\{(u_1^{(i)}, u_2^{(i)}) \mid i \in \mathbb{N}\}$ such that for all $i \in \mathbb{N}$ holds $P(u_1^{(i)}, u_2^{(i)})$ and $p^{(i)}$ is a proper prefix of $p^{(i+1)}$.

Again, we refer by $p^{(i)}$ to the path $p \in \text{Paths}_y(r_\tau(u_2^{(i)})(c_a[u_1^{(i)}, y]))$, which fulfills Property P_6 for the i^{th} pair $(u_1^{(i)}, u_2^{(i)})$. Assume the path p_{i_2} is a proper prefix of path p_{i_1} .

Basis case: First, we show that u_{i_1} and u_{i_2} fulfill $P(u_{i_1}, u_{i_2})$. With $z_j = z_{i_j}$ for both $j \in [2]$ follow immediately P_1 and P_2 . The terms z_{i_1}, z_{i_2} are in $\hat{\mathcal{T}}_\Delta(y)$ because p_{i_1} and p_{i_2} are y -paths of them, respectively. P_3 and P_4 follow. If p_{i_2} is a proper prefix of p_{i_1} , then $z_{i_2}[p_{i_2}] = y$, whereas there is a tree $v \in \hat{\mathcal{T}}_\Delta(y)$ different to y with $z_{i_1}[p_{i_2}] = v$. Besides P_6 , we also get P_5 :

$$\begin{aligned} \sigma_\tau(u_{i_2}) &= (z_{i_2} \cdot \sigma_\tau(u_{i_2}))[p_{i_2}] \\ &= \tau(c_a[u_{i_1}, u_{i_2}])[p_{i_2}] \\ &= (z_{i_1} \cdot \sigma_\tau(u_{i_1}))[p_{i_2}] \\ &= v \cdot \sigma_\tau(u_{i_1}) \end{aligned}$$

Property P_5 holds with $p = p_{i_2}$. Let u_{i_1} and u_{i_2} be $u_1^{(1)}$ and $u_2^{(1)}$, respectively. In addition, we have $p^{(1)} = p_{i_2}$.

Inductive step: Now assume we have trees $u_1^{(i)}$ and $u_2^{(i)}$, which fulfill the property $P(u_1^{(i)}, u_2^{(i)})$. Let the trees z_j in P_3 and P_4 for $j \in [2]$, v in P_5 , and the path p in P_6 of $P(u_1^{(i)}, u_2^{(i)})$ be denoted by $z_j^{(i)}$, $v^{(i)}$, and $p^{(i)}$, respectively. Furthermore, we denote $\tau(c_a[u_1^{(i)}, u_2^{(i)}])$ by $z^{(i)}$. It follows by the definition and the consistency of (σ_τ, r_τ) :

$$z_1^{(i)} \cdot \sigma_\tau(u_1^{(i)}) = z^{(i)} = z_2^{(i)} \cdot \sigma_\tau(u_2^{(i)}). \quad (6.10)$$

With Properties P_3 and P_4 it follows that the residuals of $u_1^{(i)}$ and $u_2^{(i)}$ are essential. Since the bottom-up partition is proper, there are infinitely many

trees $\bar{u}_j \equiv_{r_\tau} u_j^{(i)}$ for both $j \in [2]$. Because of this, we get the following equations for each of these pairs (\bar{u}_1, \bar{u}_2) :

- $r_\tau(\bar{u}_1)(c_a[y, u_2^{(i)}]) = z_1^{(i)}$ and
- $r_\tau(\bar{u}_2)(c_a[u_1^{(i)}, y]) = z_2^{(i)}$.

Consequently: $\tau(c_a[\bar{u}_1, u_2^{(i)}]) = z_1^{(i)} \cdot \sigma_\tau(\bar{u}_1)$ and $\tau(c_a[u_1^{(i)}, \bar{u}_2]) = z_2^{(i)} \cdot \sigma_\tau(\bar{u}_2)$. Thus, there exist \bar{z}_1 and \bar{z}_2 with

- $r_\tau(u_1^{(i)})(c_a[y, \bar{u}_2]) = \bar{z}_1$ and
- $r_\tau(u_2^{(i)})(c_a[\bar{u}_1, y]) = \bar{z}_2$.

Note that \bar{z}_1 depends on \bar{u}_2 and \bar{z}_2 on \bar{u}_1 . Then, we have

$$\bar{z}_1 \cdot \sigma_\tau(u_1^{(i)}) = \tau(c_a[u_1^{(i)}, \bar{u}_2]) = z_2^{(i)} \cdot \sigma_\tau(\bar{u}_2) \quad (6.11)$$

$$z_1^{(i)} \cdot \sigma_\tau(\bar{u}_1) = \tau(c_a[\bar{u}_1, u_2^{(i)}]) = \bar{z}_2 \cdot \sigma_\tau(u_2^{(i)}) . \quad (6.12)$$

Because of the equivalence of $u_j^{(i)}$ and \bar{u}_j (for both $j \in [2]$), also

- $r_\tau(\bar{u}_1)(c_a[y, \bar{u}_2]) = \bar{z}_1$ and
- $r_\tau(\bar{u}_2)(c_a[\bar{u}_1, y]) = \bar{z}_2$.

Consequently:

$$\bar{z}_1 \cdot \sigma_\tau(\bar{u}_1) = \tau(c[\bar{u}_1, \bar{u}_2]) = \bar{z}_2 \cdot \sigma_\tau(\bar{u}_2) . \quad (6.13)$$

For each pair (\bar{u}_1, \bar{u}_2) hold the Properties P_1 and P_2 .

Since the partition is proper, the set of output trees $\{\sigma_\tau(\bar{u}_j) \mid \bar{u}_j \equiv_{r_\tau} u_j^{(i)}\}$ is infinite (for both $j \in [2]$), whereas the number of subtrees of $\sigma_\tau(u_1^{(i)})$ and $\sigma_\tau(u_2^{(i)})$ is finite. Furthermore, also for every tree \bar{v}_1 , also $v^{(i)} \cdot \bar{v}_1$ and \bar{v}_1 have only finitely many subtrees. Thus, there are infinitely many different pairs (\bar{v}_1, \bar{v}_2) satisfying the conditions:

1. $\bar{v}_j = \sigma_\tau(\bar{u}_j)$ for some $\bar{u}_j \equiv_{r_\tau} u_j^{(i)}$ (for both $j \in [2]$)
2. $\bar{v}_2 \neq v^{(i)} \cdot \bar{v}_1$
3. $\bar{v}_j \not\sqsubseteq \sigma_\tau(u_{j'}^{(i)})$ for $j, j' \in [2]$ (and with it $\sigma_\tau(u_{j'}^{(i)}) \neq \bar{v}_j$)
4. $\bar{v}_2 \not\sqsubseteq \bar{v}_1$

Additionally, we restrict the set of pairs (\bar{v}_1, \bar{v}_2) , which fulfill these conditions such that the following two conditions hold, too:

5. $r_\tau(u_1^{(i)})(c_a[y, \bar{u}_2]) \in \hat{\mathcal{T}}_\Delta(y)$
6. $r_\tau(u_2^{(i)})(c_a[\bar{u}_1, y])[p^{(i)}] \in \hat{\mathcal{T}}_\Delta(y)$

To show that such a pair exists, assume there is no tree $\bar{v}_2 = \sigma_\tau(\bar{u}_2)$, which fulfills Conditions 1.-5., i.e., there are at least two different trees $\bar{v}_2 = \sigma_\tau(\bar{u}_2)$ and $\bar{v}'_2 = \sigma_\tau(\bar{u}'_2)$ with

- $\bar{z}_1 = r(u_1^{(i)})(c_a[y, \bar{u}_2]) \in \mathcal{T}_\Delta$ and
- $\bar{z}'_1 = r(u_1^{(i)})(c_a[y, \bar{u}'_2]) \in \mathcal{T}_\Delta$.

With Equations 6.11 and 6.13 for both \bar{u}_2 and \bar{u}'_2 and an arbitrary \bar{u}_1 , we get:

$$\begin{aligned} z_2^{(i)} \cdot \bar{v}_2 &= \bar{z}_1 \cdot \sigma_\tau(u_1^{(i)}) = \bar{z}_1 = \bar{z}_1 \cdot \bar{v}_1 = \bar{z}_2 \cdot \bar{v}_2 \\ z_2^{(i)} \cdot \bar{v}'_2 &= \bar{z}'_1 \cdot \sigma_\tau(u_1^{(i)}) = \bar{z}'_1 = \bar{z}'_1 \cdot \bar{v}_1 = \bar{z}_2 \cdot \bar{v}'_2 \end{aligned}$$

By Proposition 2.1 it follows that it exists $s \in \hat{\mathcal{T}}_\Delta(y)$ with $z_2^{(i)} = \bar{z}_2 \cdot s$ or $z_2^{(i)} \cdot s = \bar{z}_2$ (because $z_2^{(i)} \notin \mathcal{T}_\Delta$ according to Property P_4). It follows (in both cases) that $s \cdot \bar{v}_2 = \bar{v}_2$ and in order that, s equals y and $z_2^{(i)} = \bar{z}_2$. With Equations 6.10 and 6.12, we get:

$$\begin{aligned} z_1^{(i)} \cdot \sigma_\tau(u_1^{(i)}) &= z_2^{(i)} \cdot \sigma_\tau(u_2^{(i)}) && \text{(Equation 6.10)} \\ &= \bar{z}_2 \cdot \sigma_\tau(u_2^{(i)}) \\ &= z_1^{(i)} \cdot \bar{v}_1 && \text{(Equation 6.12)} \end{aligned}$$

Since $\sigma_\tau(u_1^{(i)}) \neq \bar{v}_1$ (Condition 3), $z_1^{(i)}$ must be a ground tree in \mathcal{T}_Δ , which is a contradiction to Property P_3 for $(u_1^{(i)}, u_2^{(i)})$. It follows that there are infinitely many pairs fulfilling Conditions 1.-5.

Now assume there is no tree $\bar{v}_1 = \sigma_\tau(\bar{u}_1)$, which fulfills Conditions 1.-6., i.e., there are at least two different trees $\bar{v}_1 = \sigma_\tau(\bar{u}_1)$ and $\bar{v}'_1 = \sigma_\tau(\bar{u}'_1)$ with

- $\bar{z}_2 = r(u_2^{(i)})(c_a[\bar{u}_1, y])$ and $\bar{z}_2[p^{(i)}] \in \mathcal{T}_\Delta$ and
- $\bar{z}'_2 = r(u_2^{(i)})(c_a[\bar{u}'_1, y])$ and $\bar{z}'_2[p^{(i)}] \in \mathcal{T}_\Delta$.

With Equations 6.12 for both \bar{u}_1 and \bar{u}'_1 and an arbitrary \bar{u}_2 , we get:

$$\begin{aligned} \bar{z}_2[p^{(i)}] &= (\bar{z}_2 \cdot \sigma_\tau(u_2^{(i)}))[p^{(i)}] = (z_1^{(i)} \cdot \bar{v}_1)[p^{(i)}] = v^{(i)} \cdot \bar{v}_1 \\ \bar{z}'_2[p^{(i)}] &= (\bar{z}'_2 \cdot \sigma_\tau(u_2^{(i)}))[p^{(i)}] = (z_1^{(i)} \cdot \bar{v}'_1)[p^{(i)}] = v^{(i)} \cdot \bar{v}'_1 \end{aligned}$$

Since $\bar{v}_1 \neq \bar{v}'_1$, it follows that $z_1^{(i)}[p^{(i)}] = v^{(i)}$. But then we get

$$\begin{aligned} \bar{v}_2 &= (z_2^{(i)} \cdot \bar{v}_2)[p^{(i)}] && P_6 \text{ for } (u_1^{(i)}, u_2^{(i)}) \\ &= (\bar{z}_1 \cdot \sigma_\tau(u_1^{(i)}))[p^{(i)}] && \text{Equation 6.11} \\ &= v^{(i)} \cdot \sigma_\tau(u_1^{(i)}) \\ &= \sigma_\tau(u_2^{(i)}) && P_5 \text{ for } (u_1^{(i)}, u_2^{(i)}) . \end{aligned}$$

This contradicts the third condition, i.e., $\sigma_\tau(u_2^{(i)}) \neq \bar{v}_2$ (Page 115). Thus, there must be a pair $(u_1^{(i+1)}, u_2^{(i+1)})$ such that the Conditions 1.-6. hold for trees $v_j^{(i+1)} = \sigma_\tau(u_j^{(i+1)})$, for $j \in [2]$, with

- $z_1^{(i+1)} = r_\tau(u_1^{(i)})(c_a[y, u_2^{(i+1)}]) \in \hat{\mathcal{T}}_\Delta(y)$ and
- $z_2^{(i+1)} = r_\tau(u_2^{(i)})(c_a[u_1^{(i+1)}, y])$ with $z_2^{(i+1)}[p^{(i)}] \in \hat{\mathcal{T}}_\Delta(y)$.

Consequently, the pair $(u_1^{(i+1)}, u_2^{(i+1)})$ fulfills Properties $P_1 - P_4$. Now assume that Property P_5 is not true for this pair, i.e., $v_1^{(i+1)} \not\sqsubseteq v_2^{(i+1)}$. Equation 6.12 and Property P_6 for pair $(u_1^{(i)}, u_2^{(i)})$ state

$$v^{(i)} \cdot v_1^{(i+1)} = (z_1^{(i)} \cdot v_1^{(i+1)})[p^{(i)}] = (z_2^{(i+1)} \cdot \sigma_\tau(u_2^{(i)}))[p^{(i)}].$$

Since $v_1^{(i+1)} \not\sqsubseteq \sigma_\tau(u_2^{(i)})$ and $\sigma_\tau(u_2^{(i)}) = v^{(i)} \cdot \sigma_\tau(u_1^{(i)})$, there are the following possibilities:

1. $v_1^{(i+1)} = v \cdot \sigma_\tau(u_2^{(i)})$ and $z_2^{(i+1)}[p^{(i)}] = v^{(i)} \cdot v \in \hat{\mathcal{T}}_\Delta(y)$ or
2. $z_2^{(i+1)}[p^{(i)}] = v^{(i)} \cdot v_1^{(i+1)}$.

Assume the first possibility: Equation 6.13 gives us

$$(z_1^{(i+1)} \cdot v_1^{(i+1)})[p^{(i)}] = (z_2^{(i+1)} \cdot v_2^{(i+1)})[p^{(i)}] = v^{(i)} \cdot v \cdot v_2^{(i+1)}.$$

We know that $v_1^{(i+1)}$ is no subtree of $v_2^{(i+1)}$ and vice versa. Furthermore, it holds $v_1^{(i+1)} = v \cdot v^{(i)} \cdot \sigma_\tau(u_1^{(i)})$ and, thus, $v_1^{(i+1)}$ is no subtree of v and $v^{(i)}$. It follows that

$$z_1^{(i+1)}[p^{(i)}] = v^{(i)} \cdot v \cdot v_2^{(i+1)}.$$

With Equation 6.11 and P_6 for $(u_1^{(i)}, u_2^{(i)})$, it follows

$$v^{(i)} \cdot v \cdot v_2^{(i+1)} = z_1^{(i+1)} \cdot \sigma_\tau(u_1^{(i)})[p^{(i)}] = (z_2^{(i)} \cdot v_2^{(i+1)})[p^{(i)}] = v_2^{(i+1)}.$$

Thus, $v^{(i)} \cdot v = y$, which is a contradiction to Property P_5 for $(u_1^{(i)}, u_2^{(i)})$.

Now assume the second possibility, i.e., $z_2^{(i+1)}[p^{(i)}] = v^{(i)} \cdot v_1^{(i+1)}$: With Equation 6.13, we get

$$(z_1^{(i+1)} \cdot v_1^{(i+1)})[p^{(i)}] = (z_2^{(i+1)} \cdot v_2^{(i+1)})[p^{(i)}] = v^{(i)} \cdot v_1^{(i+1)}$$

and with Equation 6.11 and Property P_6 for $(u_1^{(i)}, u_2^{(i)})$:

$$(z_1^{(i+1)} \cdot \sigma_\tau(u_1^{(i)}))[p^{(i)}] = (z_2^{(i)} \cdot v_2^{(i+1)})[p^{(i)}] = v_2^{(i+1)}.$$

Let us consider $z_1^{(i+1)}[p^{(i)}]$:

- If $z_1^{(i+1)}[p^{(i)}] = v^{(i)} \cdot v_1^{(i+1)}$, it equals $v_2^{(i+1)}$. However, this contradicts the assumption that $v_1^{(i+1)}$ is no subtree of $v_2^{(i+1)}$.

- If $z_1^{(i+1)}[p^{(i)}] = v^{(i)}$, we get $v_2^{(i+1)} = v^{(i)} \cdot \sigma_\tau(u_1^{(i)}) = \sigma_\tau(u_2^{(i)})$. However, this contradicts the choice of $v_2^{(i+1)}$.
- If $z_1^{(i+1)}[p^{(i)}] = v'$ for some $v' \neq v^{(i)}$, we get $v' \cdot \sigma_\tau(u_1^{(i)}) = v_2^{(i+1)}$. Then either $v_1^{(i+1)}$ is a subtree of v' and with that $v_1^{(i+1)} \supseteq v_2^{(i+1)}$ (which is a contradiction to the assumption), or it is a subtree of $v^{(i)}$ and with that $v_1^{(i+1)} \supseteq v^{(i)} \cdot v_1^{(i)} = v_2^{(i)}$ (which contradicts Condition 3).

Thus, we get Property P_5 for $(u_1^{(i+1)}, u_2^{(i+1)})$:

$$\exists v^{(i+1)} \in \hat{\mathcal{T}}_\Delta(y) \setminus \{y\} : v_2^{(i+1)} = v^{(i+1)} \cdot v_1^{(i+1)}.$$

We define $v^{(i+1)}$ such that $v_1^{(i+1)}$ is no subtree of it.

Now we will find the path $p^{(i+1)}$: Thereto, consider the path $p^{(i)}$. We know by Property P_6 for $(u_1^{(i)}, u_2^{(i)})$, that $z_1^{(i)}[p^{(i)}] = v^{(i)}$ and $z_2^{(i)}[p^{(i)}] = y$. Equations 6.11 and 6.12 provide:

$$\begin{aligned} (z_1^{(i+1)} \cdot \sigma_\tau(u_1^{(i)}))[p^{(i)}] &= v_2^{(i+1)} \\ (z_2^{(i+1)} \cdot \sigma_\tau(u_2^{(i)}))[p^{(i)}] &= v^{(i)} \cdot v_1^{(i+1)} \end{aligned}$$

With Condition 6, we know that $p^{(i)}$ is a prefix of a y -path of $z_2^{(i+1)}$. More precisely,

$$z_2^{(i+1)}[p^{(i)}] = v^{(i)} \cdot v' \tag{6.14}$$

with $v' \cdot \sigma_\tau(u_2^{(i)}) = v_1^{(i+1)}$ and $v' \in \hat{\mathcal{T}}_\Delta(y)$. It must be $v^{(i)} \cdot v'$ instead of only a part of $v^{(i)}$. Assume the opposite, i.e., $z_2^{(i+1)}[p^{(i)}] = v'$ with $v^{(i)}$ is no prefix of v' . Then we get with Equation 6.12 that

$$v' \cdot \sigma_\tau(u_2^{(i)}) = v^{(i)} \cdot v_1^{(i+1)}.$$

Because $v_1^{(i+1)}$ is no subtree of $\sigma_\tau(u_2^{(i)})$ (Condition 3), there must be a path p such that $v^{(i)}[p] = \sigma_\tau(u_2^{(i)})$. Then Property P_5 for the pair $(u_1^{(i)}, u_2^{(i)})$ implies that $\sigma_\tau(u_2^{(i)})$ is a subtree of itself — Contradiction. Thus, Equation 6.14 is true.

With this property and Equation 6.13,

$$(z_1^{(i+1)} \cdot v_1^{(i+1)})[p^{(i)}] = v^{(i)} \cdot v' \cdot v_2^{(i+1)} = v^{(i)} \cdot v' \cdot v^{(i+1)} \cdot v_1^{(i+1)}.$$

There are two possibilities: $z_1^{(i+1)}[p^{(i)}] \in \mathcal{T}_\Delta$ or $z_1^{(i+1)}[p^{(i)}] = v^{(i)} \cdot v' \cdot v^{(i+1)}$. The other cases are not possible: $v_1^{(i+1)}$ is no subtree of $v^{(i)}$ or v' because $v_1^{(i+1)}$ is no subtree of $\sigma_\tau(u_2^{(i)})$ and no proper subtree of itself. In addition, by definition of $v^{(i+1)}$, the tree $v_1^{(i+1)}$ is no subtree of $v^{(i+1)}$.

Assume the first possibility, i.e., $z_1^{(i+1)}[p^{(i)}] = v^{(i)} \cdot v' \cdot v_2^{(i+1)}$. Then we get with Equation 6.11

$$v^{(i)} \cdot v' \cdot v_2^{(i+1)} = (z_1^{(i+1)} \cdot \sigma_\tau(u_1^{(i)}))[p^{(i)}] = v_2^{(i+1)}.$$

Since $v^{(i)} \neq y$, this is a contradiction. Thus, the last possibility, must be true. We have:

$$z_1^{(i+1)}[p^{(i)}] = v^{(i)} \cdot v' \cdot v^{(i+1)} \quad \text{and} \quad z_2^{(i+1)}[p^{(i)}] = v^{(i)} \cdot v'$$

Let p be a y -path of $v^{(i)} \cdot v'$. Then we define

$$p^{(i+1)} = p^{(i)} \cdot p.$$

This path fulfills Property P_6 for $(u_1^{(i+1)}, u_2^{(i+1)})$ and $p^{(i)}$ is a proper prefix of this path $p^{(i+1)}$. \square

The y -paths are not only no proper prefixes of each other. They are also different:

Lemma 6.14. *Assume $c \cdot \mathbf{a}(u_1, \dots, u_k) \in \text{dom}(\tau)$. For all $i \in [k]$ let*

$$r_\tau(u_i)(c \cdot \mathbf{a}(u_1, \dots, u_{i-1}, y, u_{i+1}, \dots, u_k)) = z_i.$$

Then, it holds for all y -paths $p_{i_1} \in \text{Paths}_y(z_{i_1})$, $p_{i_2} \in \text{Paths}_y(z_{i_2})$, for some $i_1, i_2 \in [k]$ that $i_1 \neq i_2$ implies $p_{i_2} \neq p_{i_1}$.

Proof. Let $1 \leq i_1 < i_2 \leq k$, $p_{i_1} \in \text{Paths}_y(z_{i_1})$, and $p_{i_2} \in \text{Paths}_y(z_{i_2})$. Under the assumption that p_{i_2} equals p_{i_1} , we show that there is an image \bar{z}_2 of a context of u_{i_2} such that there is a y -path p in \bar{z}_2 and p_{i_1} is a proper prefix of p . Then we get by Lemma 6.13 a contradiction.

Assume $p_{i_2} = p_{i_1}$. Let $c_{\mathbf{a}}$ be the tree $c \cdot \mathbf{a}(u_1, \dots, u_k)$ where u_{i_1} and u_{i_2} are replaced by x_1 and x_2 respectively:

$$c_{\mathbf{a}} = c \cdot \mathbf{a}(u_1, \dots, u_{i_1-1}, \mathbf{x}_1, u_{i_1+1}, \dots, u_{i_2-1}, \mathbf{x}_2, u_{i_2+1}, \dots, u_k).$$

Analog to the inductive step of Lemma 6.13, there are infinitely many trees $\bar{u}_j \equiv_{r_\tau} u_{i_j}$ for both $j \in [2]$, such that $\sigma_\tau(\bar{u}_j) \not\sqsubseteq \sigma_\tau(u_{i_{j'}})$ for $j, j' \in \{1, 2\}$, and $\sigma_\tau(\bar{u}_1) \neq \sigma_\tau(\bar{u}_2)$. Furthermore, there exist \bar{z}_1, \bar{z}_2 with the following properties:

$$\begin{aligned} r_\tau(\bar{u}_1)(c_{\mathbf{a}}[y, u_{i_2}]) &= z_{i_1} = r_\tau(u_{i_1})(c_{\mathbf{a}}[y, u_{i_2}]) \\ r_\tau(\bar{u}_2)(c_{\mathbf{a}}[u_{i_1}, y]) &= z_{i_2} = r_\tau(u_{i_2})(c_{\mathbf{a}}[u_{i_1}, y]) \\ r_\tau(\bar{u}_1)(c_{\mathbf{a}}[y, \bar{u}_2]) &= \bar{z}_1 = r_\tau(u_{i_1})(c_{\mathbf{a}}[y, \bar{u}_2]) \\ r_\tau(\bar{u}_2)(c_{\mathbf{a}}[\bar{u}_1, y]) &= \bar{z}_2 = r_\tau(u_{i_2})(c_{\mathbf{a}}[\bar{u}_1, y]) \end{aligned}$$

$$z_{i_1} \cdot \sigma_\tau(u_{i_1}) = \tau(c_{\mathbf{a}}[u_{i_1}, u_{i_2}]) = z_{i_2} \cdot \sigma_\tau(u_{i_2}) \quad (6.15)$$

$$\bar{z}_1 \cdot \sigma_\tau(u_{i_1}) = \tau(c_{\mathbf{a}}[u_{i_1}, \bar{u}_2]) = z_{i_2} \cdot \sigma_\tau(\bar{u}_2) \quad (6.16)$$

$$z_{i_1} \cdot \sigma_\tau(\bar{u}_1) = \tau(c_{\mathbf{a}}[\bar{u}_1, u_{i_2}]) = \bar{z}_2 \cdot \sigma_\tau(u_{i_2}) \quad (6.17)$$

$$\bar{z}_1 \cdot \sigma_\tau(\bar{u}_1) = \tau(c[\bar{u}_1, \bar{u}_2]) = \bar{z}_2 \cdot \sigma_\tau(\bar{u}_2). \quad (6.18)$$

Now consider the path p_{i_1} . We know that $z_{i_1}[p_{i_1}] = y = z_{i_2}[p_{i_1}]$. Equation 6.15 provides:

$$\sigma_\tau(u_{i_1}) = (z_{i_1} \cdot \sigma_\tau(u_{i_1}))[p_{i_1}] = (z_{i_2} \cdot \sigma_\tau(u_{i_2}))[p_{i_1}] = \sigma_\tau(u_{i_2})$$

By Equations 6.16 and 6.17, we get

$$\begin{aligned} (\bar{z}_1 \cdot \sigma_\tau(u_{i_1}))[p_{i_1}] &= \sigma_\tau(\bar{u}_2) \\ (\bar{z}_2 \cdot \sigma_\tau(u_{i_2}))[p_{i_1}] &= \sigma_\tau(\bar{u}_1) \end{aligned}$$

Assume that p_{i_1} is no prefix of a y -path of \bar{z}_2 . Then $\bar{z}_2[p_{i_1}] = \sigma_\tau(\bar{u}_1)$ and (with Equation 6.18)

$$(\bar{z}_1 \cdot \sigma_\tau(\bar{u}_1))[p_{i_1}] = \sigma_\tau(\bar{u}_1) .$$

There are two possibilities: $\bar{z}_1[p_{i_1}] \in \mathcal{T}_\Delta$ or $\bar{z}_1[p_{i_1}] = y$. We get

$$\begin{aligned} \sigma_\tau(\bar{u}_1) &= (\bar{z}_1 \cdot \sigma_\tau(u_{i_1}))[p_{i_1}] = \sigma_\tau(\bar{u}_2) \\ \text{or } \sigma_\tau(u_{i_1}) &= (\bar{z}_1 \cdot \sigma_\tau(u_{i_1}))[p_{i_1}] = \sigma_\tau(\bar{u}_2) , \end{aligned}$$

respectively. Both contradict the choice of $\sigma_\tau(\bar{u}_1) \neq \sigma_\tau(\bar{u}_2) \neq \sigma_\tau(u_{i_1})$. Thus, path p_{i_1} must be a prefix of a y -path of \bar{z}_2 . More precisely,

$$\bar{z}_2[p_{i_1}] = v'$$

with $v' \cdot \sigma_\tau(u_{i_2}) = \sigma_\tau(\bar{u}_1)$ and $v' \in \hat{\mathcal{T}}_\Delta(y) \setminus \{y\}$ (because $\sigma_\tau(\bar{u}_1) \neq \sigma_\tau(u_{i_2})$). Let p' be a y -path of v' . Then $p_{i_1} \cdot p' = p$ is a y -path of \bar{z}_2 and we have:

- $c \cdot \mathbf{a}(u'_1, \dots, u'_k) \in \text{dom}(\tau)$ with $u'_i = u_i$ for $i \neq i_1$ and $u'_{i_1} = \bar{u}_1$.
- for all $i \in [k]$: $r_\tau(u'_i)(c \cdot \mathbf{a}(u'_1, \dots, u'_{i-1}, y, u'_{i+1}, \dots, u'_k)) = z'_i$ with $z'_{i_2} = \bar{z}_2$ and $z'_{i_1} = z_{i_1}$.
- paths $p_{i_1} \in \text{Paths}_y(z'_{i_1})$, $p \in \text{Paths}_y(z'_{i_2})$ holds p_{i_1} is a proper prefix of p .

This contradicts Lemma 6.13. \square

By these two lemmas, a corresponding property for paths of dedicated images follows immediately:

Corollary 6.15. *Assume $\mathbf{a}(u_1, \dots, u_k) \in \text{dom}(\sigma_\tau)$. For all $i \in [k]$ let*

$$\text{di}(u_i)(\mathbf{a}(u_1, \dots, u_{i-1}, y, u_{i+1}, \dots, u_k)) = z_i .$$

Then, it holds for all y -paths $p_{i_1} \in \text{Paths}_y(z_{i_1})$, $p_{i_2} \in \text{Paths}_y(z_{i_2})$, for some $i_1, i_2 \in [k]$ that $i_1 \neq i_2$ implies p_{i_2} is no prefix of p_{i_1} .

By the following lemma, we get that the greatest common prefixes, as used in Definition 6.6, are defined. In the definition, the trees $\mathbf{gcp}(z_1, \dots, z_k)$ with dedicated images z_i are part of the right-hand sides. The conditions of the following lemma are fulfilled for these images z_i (with $v_i = \sigma_\tau(u_i)$).

Lemma 6.16. *Assume $z_1, \dots, z_k \in \mathcal{T}_\Sigma(y)$ and there are $v_1, \dots, v_k \in \mathcal{T}_\Sigma$, such that for all $i, j \in [k]$ holds*

$$z_i \cdot v_i = z_j \cdot v_j .$$

If for all $i, j \in [k]$ with $i \neq j$ and $p_i \in \text{Paths}_y(z_i)$, $p_j \in \text{Paths}_y(z_j)$ holds that p_i is no prefix of p_j , then $\mathbf{gcp}(z_1, \dots, z_k)$ is defined.

Proof. First, assume $z_1, \dots, z_k \in \mathcal{T}_\Sigma$. Then,

$$z_i = z_i \cdot v_i = z_j \cdot v_j = z_j$$

for all $i, j \in \{1, \dots, k\}$. Consequently, $\text{gcp}(z_1, \dots, z_k) = z_1$ is defined.

Now assume, w.l.o.g., that $z_1 \in \hat{\mathcal{T}}_\Sigma(y)$ and p is a longest y -path of z_1 , i.e., for all $i \in [k]$ and all $p' \in \text{Paths}_y(z_i)$ holds $|p| \geq |p'|$. We prove that the greatest common prefix is defined, by induction over the length of p . For $p = \epsilon$, it follows that $z_1 = y$ and $z_2, \dots, z_k \in \mathcal{T}_\Sigma$. Otherwise, p would be a prefix of another y -path. Thus, $z_i = z_i \cdot v_i = z_j \cdot v_j = z_j$ for all $i, j \in \{2, \dots, k\}$. We get $\text{gcp}(z_1, \dots, z_k) = x_1$ is defined.

If $p = (\mathbf{a}, d).p'$ for a path p' , $\mathbf{a} \in \Sigma^{(n)}$, and $d \leq n$, then $z_1 = \mathbf{a}(z_{1,1}, \dots, z_{1,n})$ and $p' \in \text{Paths}_y(z_{1,d})$. If $z_i = y$ for some $i \neq 1$, then ϵ is a y -path and a prefix of p , which is a contradiction. Thus, z_2, \dots, z_k do not equal y . Furthermore, for $2 \leq i \leq k$, the equation $z_i \cdot v_i = z_1 \cdot \sigma(u_1)$ implies that $z_i = \mathbf{a}(z_{i,1}, \dots, z_{i,n})$ with $z_{i,j} \cdot v_i = z_{1,j} \cdot \sigma(u_1)$ for $1 \leq j \leq n$. By induction, we get for all $1 \leq j \leq n$ that $z'_j = \text{gcp}(z_{1,j}, \dots, z_{k,j})$ is defined. Consequently, $\text{gcp}(z_1, \dots, z_k)$ is defined and equals $\mathbf{a}(z'_1, \dots, z'_n)$. \square

The definition of a right-hand side in Definition 6.6 depends on the representatives of the states in the left-hand side. However, also if we choose different representatives, we get the same right-hand side. To prove that, we first need the following lemma:

Lemma 6.17. *Assume $z_1, \dots, z_k \in \mathcal{T}_\Sigma(y)$ and there are $v_1, \dots, v_k \in \mathcal{T}_\Sigma$, such that for all $i, j \in [k]$ holds*

$$z_i \cdot v_i = z_j \cdot v_j.$$

If $z = \text{gcp}(z_1, \dots, z_k)$ is defined, then for all $1 \leq i \leq k$

$$z[v_1, \dots, v_{i-1}, y, v_{i+1}, \dots, v_k] = z_i.$$

Proof. W.l.o.g., we prove the equation for $i = 1$, i.e., $z[y, v_2, \dots, v_k] = z_1$, by induction over z : First, assume $z = x_1$. By definition of gcp it follows

$$z_1 = y \quad \text{and} \quad \forall j \in \{2, \dots, k\} : z_j = z_2 \in \mathcal{T}_\Sigma(y) \setminus \{y\}$$

Thus, $z[y, v_2, \dots, v_k] = y = z_1$. Now assume $z = x_i$ for some $1 \neq i \leq k$. By definition of gcp it follows

$$z_i = y \quad \text{and} \quad \forall j \in \{1, \dots, k\} \setminus \{i\} : z_j = z_1 \in \mathcal{T}_\Sigma$$

Otherwise, if $z_1 \in \hat{\mathcal{T}}_\Sigma(y)$, $\epsilon \in \text{Paths}_y(z_i)$ is a prefix of each y -path in z_1 . Thus,

$$z[y, v_2, \dots, v_k] = v_i = z_i \cdot v_i = z_1 \cdot v_1 = z_1.$$

Last, assume $z = \mathbf{a}(z'_1, \dots, z'_n)$ for some $\mathbf{a} \in \Sigma^{(n)}$. By definition of gcp it follows

$$\forall j \leq k : z_j = \mathbf{a}(z_{j,1}, \dots, z_{j,n}) \quad \text{and} \quad \forall i \leq n : z'_i = \text{gcp}(z_{1,i}, \dots, z_{k,i})$$

For all $j \in [k]$ holds $v_a = z_j \cdot v_j$. Since $z_j = \mathbf{a}(z_{j,1}, \dots, z_{j,n})$, we get $v_a = \mathbf{a}(v_{a,1}, \dots, v_{a,n})$. In particular, $z_{j,i} \cdot v_i = v_{a,i}$ for all $i \in [n]$. By induction, we get $z'_i[y, v_2, \dots, v_k] = z_{1,i}$ and this provides:

$$z_1 = \mathbf{a}(z_{1,1}, \dots, z_{1,n}) = \mathbf{a}(z'_1, \dots, z'_n)[y, v_2, \dots, v_k] = z[y, v_2, \dots, v_k] \quad \square$$

In particular, the previous lemmas hold for the dedicated images z_i in the Definition 6.6. Thus, we get that the greatest common prefix of trees with different, but r_τ -equivalent trees are the same:

Lemma 6.18. *Assume $\mathbf{a}(u_1, \dots, u_k), \mathbf{a}(u'_1, \dots, u'_k) \in \text{dom}(\sigma_\tau)$ with $u_i \equiv_{r_\tau} u'_i$ for all $i \in [k]$ and for all $i \in [k]$ let*

$$\begin{aligned} \text{di}(u_i)(\mathbf{a}(u_1, \dots, u_{i-1}, y, u_{i+1}, \dots, u_k)) &= z_i \\ \text{di}(u'_i)(\mathbf{a}(u'_1, \dots, u'_{i-1}, y, u'_{i+1}, \dots, u'_k)) &= z_i. \end{aligned}$$

Then $\text{gcp}(z_1, \dots, z_k) = \text{gcp}(z'_1, \dots, z'_k)$.

Proof. Consider the greatest common prefix $\text{gcp}(z_1, \dots, z_k)$ with respect to the tree $\mathbf{a}(u_1, \dots, u_k)$. If we swap only for one index $j \in [k]$ from u_j to u'_j , then all z_i possibly change (without z_j). Thus, for each sequence $\bar{u}_1, \dots, \bar{u}_k$ where $\bar{u}_i \in \{u_i, u'_i\}$, we get different images \bar{z}_i . For a set $I \subseteq [k]$ of indices, we define gcp^I as the greatest common prefix of the sequence z_1^I, \dots, z_k^I corresponding to the sequence u_1^I, \dots, u_k^I where $u_i^I = u'_i$ if $i \in I$ and $u_i^I = u_i$ otherwise. We prove that

$$\text{gcp}^I[v_1^I, \dots, v_{j-1}^I, x_j, \dots, x_k] = \text{gcp}^{I \cup \{i\}}[v_1^I, \dots, v_{j-1}^I, x_j, \dots, x_k]$$

implies

$$\text{gcp}^I[v_1^I, \dots, v_j^I, x_{j+1}, \dots, x_k] = \text{gcp}^{I \cup \{i\}}[v_1^I, \dots, v_j^I, x_{j+1}, \dots, x_k]$$

for every $i, j \in [k]$.

Let $u_a = \mathbf{a}(u_1, \dots, u_k)$, $v_a = \sigma_\tau(u_a)$, $\sigma_\tau(u_i) = v_i$, and $\sigma_\tau(u'_i) = v'_i$, for all $i \in [k]$. For a set $I \subseteq [k]$, we define:

$$u_i^I = \begin{cases} u_i & i \notin I \\ u'_i & i \in I \end{cases} \quad v_i^I = \begin{cases} v_i & i \notin I \\ v'_i & i \in I \end{cases}$$

Furthermore, we get z_i^I with:

$$\begin{aligned} z_i^I &= \text{di}(u_i)(\mathbf{a}(u_1^I, \dots, u_{i-1}^I, y, u_{i+1}^I, \dots, u_k^I)) \\ &= \text{di}(u'_i)(\mathbf{a}(u_1^I, \dots, u_{i-1}^I, y, u_{i+1}^I, \dots, u_k^I)) \end{aligned}$$

The equality of the two dedicated images follows since $u_i \equiv_{r_\tau} u'_i$. It follows that for a set I and index $i \notin I$ holds:

$$z_i^I = z_i^{I \cup \{i\}} \tag{6.19}$$

We abbreviate greatest common prefixes of z_1^I, \dots, z_k^I by $\text{gcp}^I = \text{gcp}(z_1^I, \dots, z_k^I)$. By Lemma 6.16, we know that for all I the greatest common prefix gcp^I is defined. For readability, we write for a set $J \subseteq [k]$: $\text{gcp}_J^I = \text{gcp}^I[\bar{v}_1, \dots, \bar{v}_k]$ where $\bar{v}_i = x_i$ if $i \in J$, $\bar{v}_i = v_i^I$ otherwise. In particular, $\text{gcp}_{[k]}^I = \text{gcp}^I$.

Thus, we want to prove that

$$\text{gcp}_{[k]}^\emptyset = \text{gcp}_{[k]}^{[k]}. \quad (6.20)$$

Let $I \subseteq [k]$ and $i \notin I$. By Lemma 6.17, we get $\text{gcp}_{\{i\}}^I = z_i^I$ and $\text{gcp}_{\{i\}}^{I \cup \{i\}} = z_i^{I \cup \{i\}}$. With Equation 6.19 it follows:

$$\text{gcp}_{\{i\}}^I = \text{gcp}_{\{i\}}^{I \cup \{i\}} \quad (6.21)$$

In the end of this proof, we prove the following two implications. Thereto, let $I, J \subseteq [k]$ be arbitrary sets and indexes $i \notin I$ and $j \notin J$ with $i \neq j$. It holds

$$v_i = v'_i \implies \text{gcp}_{[k]}^I = \text{gcp}_{[k]}^{I \cup \{i\}} \quad (6.22)$$

$$\text{gcp}_J^I = \text{gcp}_J^{I \cup \{i\}} \implies \text{gcp}_{J \cup \{j\}}^I = \text{gcp}_{J \cup \{j\}}^{I \cup \{i\}}. \quad (6.23)$$

By Implication 6.22 follows, for arbitrary set $I \subseteq [k]$ and index $i \notin I$:

$$v_i = v'_i \implies \text{gcp}_J^I = \text{gcp}_J^{I \cup \{i\}} \text{ for all } J \subseteq [k]$$

By the second implication, it follows immediately, for $I, J \subseteq [k]$ and $i \in J \setminus I$:

$$\text{gcp}_J^I = \text{gcp}_J^{I \cup \{i\}} \implies \text{gcp}_{[k]}^I = \text{gcp}_{[k]}^{I \cup \{i\}} \quad (6.24)$$

Then, we get by Equation 6.21 for every $0 \leq i \leq k-2$ the following:

$$\text{gcp}_{\{i+1\}}^{[i]} = \text{gcp}_{\{i+1\}}^{[i+1]} \quad \text{and} \quad \text{gcp}_{\{i+2\}}^{[i+1]} = \text{gcp}_{\{i+2\}}^{[i+2]}$$

In addition, by Implication 6.23 and its conclusive Equation 6.24, we get

$$\text{gcp}_{[k]}^{[i]} = \text{gcp}_{[k]}^{[i+1]} = \text{gcp}_{[k]}^{[i+2]}.$$

If $k = 1$, Equation 6.21 already provides: $\text{gcp}_{[1]}^\emptyset = \text{gcp}_{[1]}^{[1]}$. This results in Equation 6.20: $\text{gcp}_{[k]}^\emptyset = \text{gcp}_{[k]}^{[k]}$.

It remains to prove the two implications. First, we show Implication 6.22, i.e., for $1 \leq i \leq k$ and $I \subseteq [k]$ with $i \notin I$, it holds

$$v_i = v'_i \implies \text{gcp}_{[k]}^I = \text{gcp}_{[k]}^{I \cup \{i\}}.$$

Assume $v_i = v'_i$. If for all $j \in [k]$ holds $v_j = v'_j$, then follows $\text{gcp}_{[k]}^I = \text{gcp}_{[k]}^{I \cup \{i\}}$ immediately. Because of the equivalence of u_j and u'_j , $z_j^I = z_j^{I \cup \{i\}}$ holds for

all $j \in [k]$. Thus, consider $j \in [k]$ with $v_j \neq v'_j$. Assume $j \notin I$. We get the following equations:

$$\begin{aligned} \text{gcp}_{\{i\}}^I[v_i/x_i] &= \text{gcp}_{\{j\}}^I[v_j/x_j] \\ \text{gcp}_{\{i\}}^{I \cup \{i\}}[v'_i/x_i] &= \text{gcp}_{\{j\}}^{I \cup \{i\}}[v_j/x_j] \\ \text{gcp}_{\{i\}}^{I \cup \{j\}}[v_i/x_i] &= \text{gcp}_{\{j\}}^{I \cup \{j\}}[v'_j/x_j] \\ \text{gcp}_{\{i\}}^{I \cup \{i,j\}}[v'_i/x_i] &= \text{gcp}_{\{j\}}^{I \cup \{i,j\}}[v'_j/x_j] \end{aligned}$$

By Equation 6.21, we could replace $\text{gcp}_{\{i\}}^{I \cup \{i,j\}}$ by $\text{gcp}_{\{i\}}^{I \cup \{j\}}$ and so on. Furthermore, we have the assumption that $v_i = v'_i$. We get

$$\begin{aligned} \text{gcp}_{\{j\}}^I[v_j/x_j] &= \text{gcp}_{\{i\}}^I[v_i/x_i] = \text{gcp}_{\{j\}}^{I \cup \{i\}}[v_j/x_j] \\ \text{gcp}_{\{j\}}^I[v'_j/x_j] &= \text{gcp}_{\{i\}}^{I \cup \{j\}}[v_i/x_i] = \text{gcp}_{\{j\}}^{I \cup \{i\}}[v'_j/x_j] \end{aligned}$$

It follows by Proposition 2.1 that $z_{\{j\}}^I = \text{gcp}_{\{j\}}^I = \text{gcp}_{\{j\}}^{I \cup \{i\}} = z_{\{j\}}^{I \cup \{i\}}$. Now assume $j \in I$. By the same argument for $I' = I \setminus \{j\}$ and Equation 6.21, we get

$$z_{\{j\}}^I = \text{gcp}_{\{j\}}^I = \text{gcp}_{\{j\}}^{I'} = \text{gcp}_{\{j\}}^{I' \cup \{i\}} = \text{gcp}_{\{j\}}^{I \cup \{i\}} = z_{\{j\}}^{I \cup \{i\}}.$$

Since this holds for every $j \in [k]$, we get $\text{gcp}_{[k]}^I = \text{gcp}_{[k]}^{I \cup \{i\}}$ (which are the gcp of the same trees).

Last, we prove Implication 6.23, i.e., for arbitrary sets $I, J \subseteq [k]$ and indices $i \notin I$ and $j \notin J$ with $i \neq j$. It holds

$$\text{gcp}_J^I = \text{gcp}_J^{I \cup \{i\}} \implies \text{gcp}_{J \cup \{j\}}^I = \text{gcp}_{J \cup \{j\}}^{I \cup \{i\}}.$$

First, assume $v_i = v'_i$. By Implication 6.22, it follows:

$$\text{gcp}_J^I = \text{gcp}_J^{I \cup \{i\}} \quad \text{and} \quad \text{gcp}_{J \cup \{j\}}^I = \text{gcp}_{J \cup \{j\}}^{I \cup \{i\}}$$

Thus, assume $v_i \neq v'_i$. Let $I, J \subseteq [k]$ be arbitrary sets with indices $i \notin I$, $j \notin J$. Assume

$$\text{gcp}_J^I = \text{gcp}_J^{I \cup \{i\}} \quad \text{and} \quad \text{gcp}_{J \cup \{j\}}^I \neq \text{gcp}_{J \cup \{j\}}^{I \cup \{i\}}.$$

It follows

$$\text{gcp}_{J \cup \{j\}}^I[v_j^I/x_j] = \text{gcp}_J^I = \text{gcp}_J^{I \cup \{i\}} = \text{gcp}_{J \cup \{j\}}^{I \cup \{i\}}[v_j^{I \cup \{i\}}/x_j].$$

There is a path p_j such that (w.l.o.g.)

$$\text{gcp}_{J \cup \{j\}}^I[p_j] = x_j \quad \text{and} \quad \text{gcp}_{J \cup \{j\}}^{I \cup \{i\}}[p_j] = v \neq x_j.$$

This implies that $z_j^I[p_j] = y$ and $z_j^{I \cup \{i\}}[p_j] = v' \neq y$. W.l.o.g., assume $j \notin I$. We have:

$$z_j^I \cdot v_j = z_i^I \cdot v_i \quad (6.25)$$

$$z_j^{I \cup \{i\}} \cdot v_j = z_i^{I \cup \{i\}} \cdot v'_i \quad (6.26)$$

$$z_j^{I \cup \{j\}} \cdot v'_j = z_i^{I \cup \{j\}} \cdot v_i \quad (6.27)$$

$$z_j^{I \cup \{i,j\}} \cdot v'_j = z_i^{I \cup \{i,j\}} \cdot v'_i \quad (6.28)$$

If $z_j^I[p_j] = y$, then also $z_j^{I \cup \{j\}}[p_j] = y$ (Equation 6.19). With it, we get $\text{gcp}^{I \cup \{j\}}[p_j] = x_j$. By Corollary 6.15, we know that there is no y -path in $z_i^{I \cup \{j\}}$ which is a prefix of p_j or vice versa. Thus, $z_i^{I \cup \{j\}}[p_j] \in \mathcal{T}_\Delta$. It follows:

$$\begin{aligned} v'_j &= (z_j^{I \cup \{j\}} \cdot v'_j)[p_j] = (z_i^{I \cup \{j\}} \cdot v_i)[p_j] && \text{Equation 6.27} \\ &= z_i^{I \cup \{j\}}[p_j] && \text{Lemmas 6.13 and 6.14} \\ &= z_i^{I \cup \{i,j\}}[p_j] && \text{Equation 6.19} \\ &= (z_i^{I \cup \{i,j\}} \cdot v'_i)[p_j] && \text{Lemmas 6.13 and 6.14} \\ &= (z_j^{I \cup \{i,j\}} \cdot v'_j)[p_j] && \text{Equation 6.28} \\ &= (z_j^{I \cup \{i\}} \cdot v'_j)[p_j] && \text{Equation 6.19} \\ &= v' \cdot v'_j \end{aligned}$$

This is a contradiction if $v' \in \hat{\mathcal{T}}_\Delta(y)$. Otherwise, $v' = v'_j$ and we get:

$$\begin{aligned} v'_j &= v' = z_j^{I \cup \{i\}}[p_j] \\ &= (z_j^{I \cup \{i\}} \cdot v_j)[p_j] && v' \in \mathcal{T}_\Delta \\ &= (z_i^{I \cup \{i\}} \cdot v'_i)[p_j] && \text{Equation 6.26} \\ &= z_i^{I \cup \{i\}}[p_j] && \text{Lemmas 6.13 and 6.14: } z_i^{I \cup \{i\}}[p_j] \in \mathcal{T}_\Delta \\ &= z_i^I[p_j] && \text{Equation 6.19} \\ &= (z_i^I \cdot v_i)[p_j] && \text{Lemmas 6.13 and 6.14} \\ &= (z_j^I \cdot v_j)[p_j] && \text{Equation 6.25} \\ &= v_j \end{aligned}$$

This contradicts the assumption. Thus, Implication 6.23 holds and with that Equation 6.20: $\text{gcp}^\emptyset_{[k]} = \text{gcp}^{[k]}_{[k]}$. \square

Summarizing, we get with Corollary 6.15, Lemmas 6.16 and 6.18 that, given a bottom-up partition for a transformation τ , the transducer T_τ of Definition 6.6 is well-defined and correct:

Theorem 6.19. *The transducer T_τ is a well-defined deterministic bottom-up tree transducer.*

Proof. Since \equiv_{r_τ} has finite index, T_τ has finitely many states. We have to show that the definitions of the transition function and the final function are independent of the choice of representatives of involved states.

The partition (σ_τ, r_τ) is the bottom-up partition of τ . Subsequently, the equivalence relation \equiv_{r_τ} is a congruence relation. Let $k \geq 0$, $\mathbf{a} \in \Sigma^{(k)}$ and for all $i \in [k]$ let $u_i, u'_i \in \text{dom}(\sigma_\tau)$ with $u_i \equiv_{r_\tau} u'_i$. Thus, we get

$$\begin{aligned} \mathbf{a}(u_1, \dots, u_k) &\equiv_{r_\tau} \mathbf{a}(u'_1, u_2, \dots, u_k) \\ &\dots \\ &\equiv_{r_\tau} \mathbf{a}(u'_1, \dots, u'_i, u_{i+1}, \dots, u_k) \\ &\dots \\ &\equiv_{r_\tau} \mathbf{a}(u'_1, \dots, u'_k) \end{aligned}$$

It follows that for a transition, the state on the right-hand side is well-defined and independent of the choice of representatives for the states on the left-hand side. By Lemma 6.18, the same holds for the output $\text{gcp}(z_1, \dots, z_k)$ on the right-hand side.

Similarly, for the final function we get: Let $u, u' \in \text{dom}(\sigma_\tau)$ with $u \equiv_{r_\tau} u'$. Since $r_\tau(u) = r_\tau(u')$, the context y is in $\text{dom}(r_\tau(u))$ iff it is in $\text{dom}(r_\tau(u'))$. Furthermore,

$$F_\tau([u]_{r_\tau}) = r_\tau(u)(y) = r_\tau(u')(y) = F_\tau([u']_{r_\tau}) . \quad \square$$

To prove the correctness of T_τ , i.e., that $\tau^{T_\tau} = \tau$, we first show the following equivalence:

Lemma 6.20. *For every tree $u \in \mathcal{T}_\Sigma$ holds*

$$\sigma_\tau(u) = v \wedge q = [u]_{r_\tau} \iff \llbracket u \rrbracket^{T_\tau} = q(v) .$$

Proof. Assume $u \in \text{dom}(\sigma_\tau)$ and $c \in \text{dom}(r_\tau(u))$. We prove the equivalence by induction. If $u = \mathbf{b} \in \Sigma^{(0)}$, then there is a transition $\mathbf{b} \rightarrow [\mathbf{b}](\sigma_\tau(u))$. This implies $\llbracket \mathbf{b} \rrbracket^{T_\tau} = [\mathbf{b}]_{r_\tau}(\sigma_\tau(\mathbf{b}))$. If $u = \mathbf{a}(u_1, \dots, u_k)$ for $k > 0$. For all $i \in [k]$, $\sigma_\tau(u_i)$ is defined because σ_τ is a subtree function. Then, there is a transition

$$\mathbf{a}([u_1]_{r_\tau}, \dots, [u_k]_{r_\tau}) \rightarrow [u]_{r_\tau}(\text{gcp}(z_1, \dots, z_k)) \in R_\tau .$$

Let $\sigma_\tau(u_i) = v_i$ for $i \in [k]$. We know by induction that for all $i \in [k]$ holds: $\llbracket u_i \rrbracket^{T_\tau} = [u_i]_{r_\tau}(v_i)$. By definition, $\llbracket u \rrbracket^{T_\tau} = [u]_{r_\tau}(\text{gcp}(z_1, \dots, z_k)[v_1, \dots, v_k])$. In addition, Lemma 6.17 directs that

$$\text{gcp}(z_1, \dots, z_k)[v_1, \dots, v_k] = v .$$

On the other hand, assume $\llbracket u \rrbracket^{T_\tau} = q(v)$. If $u = \mathbf{b} \in \Sigma^{(0)}$, then there is a transition $\mathbf{b} \rightarrow q(v)$. By construction, this implies that $v = \sigma_\tau(\mathbf{b})$ and $q = [\mathbf{b}]_{r_\tau}$.

If the tree $u = \mathbf{a}(u_1, \dots, u_k)$ for $k > 0$, then there are states q_1, \dots, q_k with $\llbracket u_i \rrbracket^{T_\tau} = q_i(v_i)$ for all $i \in [k]$ and a transition

$$\mathbf{a}(q_1, \dots, q_k) \rightarrow q(z) \in R_\tau .$$

We get the result $\llbracket u \rrbracket^{T_\tau} = q(z[v_1, \dots, v_k])$. By induction, we know the output $\sigma_\tau(u_i) = v_i$ and the state $q_i = [u_i]_{r_\tau}$ for all $i \in [k]$. It follows by the construction of T_τ and determinism that $q = [u]_{r_\tau}$. Since the greatest common prefix is unique, $\mathbf{gcp}(z_1, \dots, z_k)[v_1, \dots, v_k] = v$. \square

It follows that T_τ defines the desired transformation τ :

Theorem 6.21. *The bottom-up tree transducer T_τ describes τ , i.e., $\tau^{T_\tau} = \tau$.*

Proof. With Lemma 6.20, we get for $u \in \mathcal{T}_\Delta$ the following: Assume $\tau(u) = v'$. Then $\sigma_\tau(u) = v$ and $r_\tau(u)(y) = z$ are defined for some v and some z with $v' = z \cdot v$. We get for the transducer that $\llbracket u \rrbracket^{T_\tau} = [u]_{r_\tau}(v)$ (Lemma 6.20) and $F_\tau([u]_{r_\tau}) = r_\tau(u)(y) = z$. Thus, $\tau^{T_\tau}(u) = z \cdot v = \tau(u)$.

Now assume $\tau^{T_\tau}(u) = v'$. Then there is a state $q \in Q$, such that $\llbracket u \rrbracket^{T_\tau} = q(v)$ and $F_\tau(q) = z$ and $v' = z \cdot v$. Thus, $\sigma_\tau(u) = v$ and $q = [u]_{r_\tau}$ (Lemma 6.20) and $r_\tau(u)(y) = F_\tau([u]_{r_\tau}) = z$, which implies $\tau(u) = z \cdot v = \tau^{T_\tau}(u)$. \square

Finally, we get that the DBTT of a bottom-up partition is already the unique minimal transducer of its tree function. Thereto, we need the following invariant:

Lemma 6.22. *For every tree $u \in \mathcal{T}_\Sigma$ and every context $c \in \mathcal{C}_\Sigma(y)$ holds*

$$c \cdot u \in \text{dom}(\sigma_\tau) \implies \llbracket c \rrbracket_{[u]_{r_\tau}}^{T_\tau} = [c \cdot u]_{r_\tau}(\text{di}(u)(c)) .$$

Proof. Assume $c \cdot u \in \text{dom}(\sigma_\tau)$. We prove the implication by induction: First, assume $c = y$, then by definition, we get

$$\llbracket y \rrbracket_{[u]_{r_\tau}}^{T_\tau} = [u]_{r_\tau}(y) = [y \cdot u]_{r_\tau}(\text{di}(u)(y)) .$$

Now assume that $c = \mathbf{a}(c_1, \dots, c_k)$ for some label $\mathbf{a} \in \Sigma^{(k)}$. Since c is a context, there is only one occurrence of y in c . Thus, w.l.o.g. only $c_1 \in \mathcal{C}_\Sigma(y)$ and $c_2, \dots, c_k \in \mathcal{T}_\Sigma$. The result of c starting at $[u]_{r_\tau}$ is defined by

$$\begin{aligned} \llbracket c \rrbracket_{[u]_{r_\tau}}^{T_\tau} &= q'(z[z_1, \dots, z_k]) \\ &\text{with } \forall i \llbracket c_i \rrbracket_{[u]_{r_\tau}}^{T_\tau} = q_i(z_i) \text{ and } R_\tau(\mathbf{a}, q_1 \dots q_k) = q'(z) \end{aligned}$$

In addition, we know:

1. For $i > 1$ holds $\llbracket c_i \rrbracket_{[u]_{r_\tau}}^{T_\tau} = \llbracket c_i \rrbracket^{T_\tau} = [c_i]_{r_\tau}(\sigma_\tau(c_i))$.
2. By induction follows $\llbracket c_1 \rrbracket_{[u]_{r_\tau}}^{T_\tau} = [c_1 \cdot u]_{r_\tau}(\text{di}(u)(c_1))$.
3. $R_\tau(\mathbf{a}, [c_1 \cdot u]_{r_\tau} \cdot [c_2]_{r_\tau} \dots [c_k]_{r_\tau}) = [\mathbf{a}(c_1 \cdot u, c_2, \dots, c_k)]_{r_\tau}(z) = [c \cdot u]_{r_\tau}(z)$

The first equation follows by Lemma 6.20 and the third is given by the construction of T_τ . The output tree z is the greatest common prefix as given in the definition of T_τ . It holds

$$z[y, \sigma_\tau(c_2), \dots, \sigma_\tau(c_k)] = \text{di}(c_1 \cdot u)(\mathbf{a}(y, c_2, \dots, c_k)) \cdot$$

This follows by Lemma 6.17 and the fact that

$$\begin{aligned} & \text{di}(c_1 \cdot u)(\mathbf{a}(y, c_2, \dots, c_k)) \cdot \sigma_\tau(c_1 \cdot u) \\ &= \text{di}(c_i)(\mathbf{a}(c_1 \cdot u, c_2, \dots, c_{i-1}, y, c_{i+1}, \dots, c_k)) \cdot \sigma_\tau(c_i) \end{aligned}$$

holds for every $i > 2$. For the dedicated image of $c_1 \cdot u$, we get for every context $c' \in \text{dom}(r_\tau(c \cdot u))$:

$$\begin{aligned} & r_\tau(c \cdot u)(c') \cdot \text{di}(c_1 \cdot u)(\mathbf{a}(y, c_2, \dots, c_k)) \cdot \text{di}(u)(c_1) \\ &= r_\tau(c_1 \cdot u)(c' \cdot \mathbf{a}(y, c_2, \dots, c_k)) \cdot \text{di}(u)(c_1) \\ &= r_\tau(u)(c' \cdot c) \end{aligned}$$

Since $\text{di}(u)(c)$ is the unique tree for which this equation holds (Definition 5.8), we get

$$\begin{aligned} \text{di}(u)(c) &= \text{di}(c_1 \cdot u)(\mathbf{a}(y, c_2, \dots, c_k)) \cdot \text{di}(u)(c_1) \\ &= z[\text{di}(u)(c_1), \sigma_\tau(c_2), \dots, \sigma_\tau(c_k)], \end{aligned}$$

which is the output of c starting at state $[u]_{r_\tau}$:

$$\llbracket c \rrbracket_{[u]_{r_\tau}}^{T_\tau} = [c \cdot u]_{r_\tau}(\text{di}(u)(c)) \quad \square$$

To prove that the unique minimal bottom-up partition induces the unique minimal DBTT, we show that the transducer T_τ is minimal, i.e., it is trim, proper, earliest, unified earliest, and minimal. These properties are induced by the corresponding properties of the bottom-up partition.

Theorem 6.23. *The transducer T_τ is minimal.*

Proof. The transducer is trim, proper, earliest, unified earliest, and minimal because the partition is trim, proper, earliest, unified earliest, minimal, and its equivalence relation is a congruence:

- Every state is represented by a subtree $u \in \text{dom}(\sigma_\tau)$. Thus, the state is reached by this tree. Both follow by Lemma 6.20. Since the partition is trim, there is a context $c \in \text{dom}(r_\tau(u))$ and with Lemma 6.22, this is also a context of $[u]_{r_\tau}$ such that this state is meaningful. If the output of a state q is useless and there is a transition $\mathbf{a}(q_1, \dots, q_k) \rightarrow q(z)$, then there are trees $u_i \in \mathcal{L}^{T_\tau}(q_i)$ (for all $i \in [k]$) and the residual of $\mathbf{a}(u_1, \dots, u_k)$ is inessential. Then the dedicated images

$$\text{di}(u_i)(\mathbf{a}(u_1, \dots, u_{i-1}, y, u_{i+1}, \dots, u_k)) = *$$

and thus, the greatest common prefix of these trees equals $*$, too. It follows that $z = *$ and with that, the transducer is trim.

- Let $q = [u]_{r_\tau}$ be an inessential state. Then $\{\llbracket u' \rrbracket^{T_\tau} \mid u' \in \mathcal{L}^{T_\tau}(q)\}$ is finite. Thus, $\sigma_\tau([u]_{r_\tau})$ is finite, too. It follows that $r_\tau(u)$ is inessential and since the partition (σ_τ, r_τ) is proper, $\sigma_\tau([u]_{r_\tau}) = \{*\}$. This implies that the inessential state q only produces $*$. — T_τ is proper.
- The greatest common suffix of a state q equals the greatest common suffix of a tree $u \in \mathcal{L}^{T_\tau}(q)$. Because the partition is earliest, the greatest common suffices equal y or \perp . For essential states, it contains y , thus, it must equal y . Thus, the transducer is earliest.
- The most-general unifier of two essential states q, q' is defined by the most-general unifier of its context functions $\tau_q^{T_\tau}$ and $\tau_{q'}^{T_\tau}$, respectively. Let $c \in \text{dom}(\tau_q^{T_\tau})$ be an arbitrary context with the result $\llbracket c \rrbracket_q^{T_\tau} = \bar{q}(z)$. Furthermore, let $F_\tau(\bar{q}) = z'$. By Lemma 6.22, we know that $\bar{q} = [c \cdot u]_{r_\tau}$ for some tree $u \in \mathcal{L}^{T_\tau}(q)$. Then, we get by Lemma 6.22 and the definition of the dedicated image:

$$\tau_q^{T_\tau}(c) = z' \cdot z = r_\tau(c \cdot u)(y) \cdot \text{di}(u)(c) = r_\tau(u)(c)$$

Analogously, we get for a tree $u' \in \mathcal{L}^{T_\tau}(q')$ that $\tau_{q'}^{T_\tau}$ equals the residual of u' , i.e., $\tau_{q'}^{T_\tau} = r_\tau(u')$. We get:

$$\text{mgu}(q, q') = \text{mgu}(r_\tau(u), r_\tau(u'))$$

Assume there is a ground term v in such a most-general unifier, which is realizable in q . That means there is an input tree $\bar{u} \in \mathcal{L}^{T_\tau}(q)$ such that $\text{out}(\bar{u}) = v$. Then $\sigma_\tau(\bar{u}) = v$ and $\bar{u} \equiv_{r_\tau} u$ (both follow by Lemma 6.22). Hence v is realizable in $r_\tau(u)$, which contradicts that the partition is unified earliest. Subsequently, T_τ is unified earliest.

- Finally, we show that the transducer is minimal. Thereto, assume T_τ is not minimal. Then there are distinct states q_1, q_2 with $q_1 \sim_{T_\tau} q_2$. Let $u_1, u_2 \in \text{dom}(\sigma_\tau)$ with $[u_i]_{r_\tau} = q_i$ for $i \in [2]$. For every context $c \in \mathcal{C}_\Sigma(y)$ follows that $[c \cdot u_1]_{r_\tau} \sim_{T_\tau} [c \cdot u_2]_{r_\tau}$ if the states exist. We know that the most-general unifier of two states equals the most-general unifier of the residuals of any of their representative subtrees, i.e.,

$$\text{mgu}(q, q') = \text{mgu}(r_\tau(u), r_\tau(u')) \text{ if } u \in \mathcal{L}^{T_\tau}(q), u' \in \mathcal{L}^{T_\tau}(q') .$$

Thus, it follows that for all contexts $c \in \mathcal{C}_\Sigma(y)$ holds

$$c \cdot u_1 \sim'_{r_\tau} c \cdot u_2 .$$

Since \sim_{r_τ} is the greatest congruence, which is a refinement of \sim'_{r_τ} , we get $u_1 \sim_{r_\tau} u_2$. Because (σ_τ, r_τ) is minimal, u_1 and u_2 are r_τ -equivalent. It follows that $q_1 = [u_1]_{r_\tau} = [u_2]_{r_\tau} = q_2$. — Contradiction.

Thus, the DBTT T_τ of the bottom-up partition (σ_τ, r_τ) is the unique minimal transducer of τ . \square

Summarizing, we proved the last part of the Myhill-Nerode Theorem 6.1. More precise, Theorems 6.19 and 6.21 give us that for every bottom-up partition a deterministic bottom-up tree transducer exists, which describes the same tree function. In addition, Theorem 6.23 states that this transducer is already the unique minimal DBTT.

Chapter 7

Conclusion

In this part, we have shown two main results for deterministic bottom-up finite-state tree-to-tree transducers:

1. We provided a Myhill-Nerode style theorem, i.e., we gave a necessary and sufficient condition for a tree function to be definable by a DBTT.
2. We gave a construction to get a unique DBTT, which can be performed in polynomial time for a large class of transducers.

For the Myhill-Nerode style theorem, tree functions were semantically characterized by partitions, which, in addition, provide an equivalence relation on subtrees of the domain of the tree function. It was shown that a tree function is definable by a DBTT if both the equivalence relation has finite index and the partition is path-finite. The latter guarantees that every path occurring in input trees yields only to finitely many different paths in corresponding output trees. We introduced a normal form for path-finite partitions with equivalence relations of finite index. This bottom-up partition is unique. Step by step, the original partition is improved to guarantee different properties. In particular, the equivalence relation of the bottom-up partition is a congruence. Additionally, the corresponding tree transducer with equivalence classes as states is defined. It is shown that the bottom-up partition leads to the unique DBTT of the same tree function.

The second contribution is the construction of this unique transducer, given an arbitrary deterministic bottom-up tree transducer. In case that the DBTT is already proper, i.e., does only produce output at essential states, the construction can be performed in polynomial time — given that we represent right-hand sides compactly. Though similar in spirit as the corresponding construction for top-down deterministic transducers, the given construction for bottom-up transducers is amazingly involved and relies on a long sequence of transformations of the original DBTT to rule out anomalies in the behavior of the transducer. In a last step, the DBTT was minimized.

It is well known that equivalence for deterministic bottom-up transducers is decidable [Zac80]. Now our results provide an polynomial algorithm for the

equivalence problem for proper DBTTs. For both transducers the unique form has to be constructed and checked for equivalence.

7.1 Future Work

We plan to evaluate in how far our novel normal-form of DBTTs and the Myhill-Nerode theorem can be applied in the context of learning tree-to-tree transformations. In a Gold-style learning algorithm [Gol78], the minimal DBTT would be learned by a characteristic sample set of input and output trees. It would be interesting, how such an algorithm can be defined using the subtree-equivalence of the bottom-up partition. Presumably, deducing the bottom-up partition and mainly identifying trees with essential residuals will be an intrinsic part of such an algorithm. As another example, the normal form can be used to decide certain (semantic) subclasses of DBTTs; e.g., we can decide whether a given DBTT is equivalent to a relabeling, using the normal form. This provides an alternative proof of [Gaz06], for the deterministic case.

Also, it remains to future work to evaluate how these ideas can be enhanced to other classes of tree transformations, e.g., macro tree transducers, which are introduced by Engelfriet in [Eng80]. In [LMN10], Lemay et al. present a Myhill-Nerode Theorem and a Gold-style learning algorithm for deterministic top-down tree transducers. A common superclass of deterministic bottom-up and top-down tree transducers are deterministic top-down tree transducers with regular look-ahead [Eng77]. Such a transducer first assigns informations to the nodes of the input tree by processing a bottom-up relabeling. After that, it uses these informations executing a top-down tree transducer. Is it possible to combine the two minimization algorithms for bottom-up and top-down transducers to get a minimal top-down transducer with regular look-ahead?

There are a lot of different applications of tree transducers, e.g., natural language processing and XSLT (see introduction). An ambitious goal is to enhance the transformation models (as tree transducer) and the corresponding algorithms (e.g., to minimize or learn the models) to approach more complex transformations, which are considered in natural language processing or which occur in practical XSLT examples. Such approaches are often restricted by complexity bounds.

Part II

Type Checking Tree Walking Transducers

Chapter 8

Introduction

The extensible markup language XML is the current standard format for exchanging structured data. Its widespread use has initiated lots of work to support processing of XML on many different levels: Specialized query languages for XML, such as XQuery, transformation languages like XSLT, and programming language support either in the form of special purpose languages like XDuce, or of binding facilities for mainstream programming languages like JAXB. A central problem in XML processing is the *(static) type checking problem*: Given an input and output type and a transformation f , can we statically check whether all outputs generated by f on valid inputs conform to the output type? Since XML types are intrinsically more complex than the types found in conventional programming languages, the type checking problem for XML poses new challenges on the design of type checking algorithms. The excellent survey of [MS05] gives an overview of the different approaches to XML type checking.

In Chapter 1 we have already said that the type checking problem for XML transformations in its most general setting is undecidable. Instead of the approximative approaches, here, we will analyze *exact XML type checking*, i.e., restricting the types and transformations in such a way that type checking becomes decidable. For the exact setting, types can be considered as regular or *recognizable* tree languages — thus, capturing the expressive strength of virtually all known type formalisms for XML [MLM00]. Exact type checking is possible for a large class of transformations (cf. Chapter 1), but even for top-down transformations it is exponential time complete and for more complex transformations, the problem is non-elementary. In this part, we present algorithms to type check expressive classes of transformations defined by tree walking transducers (2TTs), macro tree walking transducers (2MTTs), and macro forest walking transducers (2MFTs). We show that these algorithms work in polynomial time for large subclasses of transformations.

Our approach is *forward type inference*: Given an input type I and a transformation τ , we statically infer all outputs generated by τ on valid inputs. Then we determine whether the set $\tau(I)$ of generated outputs conforms to a given output type. Thereto, we check the intersection of $\tau(I)$ with the complement of

the output type for emptiness. The input and output types are given by tree automata and the transformation is given by a 2TT, 2MTT, or 2MFT. To infer the set $\tau(I)$ of produced output trees, we restrict the transformation by enforcing an additional run of the transducer on the input tree. This additional run simulates the tree automata on the input type and checks if the input tree is valid. Next, for this transducer and a tree automaton describing the output type, a transducer is constructed, which produces desired output trees only, i.e., trees of the complement of the output type. Thereto, this transducer realizes the same transformation as before, but simultaneously checks whether the produced output is accepted by the output automaton. If this transducer does not produce any output, we get an affirmative answer to the type checking problem. To check the transducer for emptiness, we construct a bottom-up tree automaton, which describes the output produced by the transducer. With regard to the different processing of the input trees (walking vs. bottom-up), this construction is complex. Roughly, we subsume all visits of the walking transducer at one node to get the state of the automaton at this node. Thus, if the transducer visits each node only a bounded number of times, the bottom-up automaton can be constructed in polynomial time.

This part is organized as follows: In **Chapter 9**, we define tree walking transducers and related notions. Additionally, we present a normal form for 2TTs, which offers better complexity bounds for the subsequent algorithms. The rules of a 2TT in normal form have restricted right-hand sides.

The type checking problem is formulated in **Chapter 10**. Our approach of using forward type inference, is presented. The types are given by finite tree automata. The basic definitions and properties for bottom-up and deterministic top-down tree automata are summarized in this chapter.

In **Chapter 11**, we give the algorithm to type check tree walking transducers. First, we intersect the transducer with the output type given by a tree automaton. The resulting transducer produces only output trees not in the output type and apart from that, works as the given transducer. If the transformation of the new transducer is empty, the origin 2TT only produces valid output with respect to the given output type. Thus, it remains to decide if the new transducer realizes the empty transformation. Thereto, we deduce an alternating tree walking automaton (ATWA), which accepts the domain of the 2TT. For this automaton, an equivalent bottom-up tree automaton is constructed. Since emptiness for a bottom-up tree automaton is decidable in linear time, we get in Section 11.2 an exponential time algorithm to decide emptiness for both 2TTs and alternating tree automata. In the next section, we show that this algorithm is a polynomial time algorithm if the transducer is *b-bounded*, i.e., every subtree is only processed in at most b different ways. With that, we show that type checking a *strictly* b -bounded 2TT, i.e., every node in an input tree is only visited b times, is polynomial in the size of the transducer and the sizes of deterministic bottom-up automata describing the input and output types.

A similar approach is applied to type check macro tree walking transducers in **Chapter 12**. First, we give the definition and discuss their semantics. Here, we additionally have to deal with parameters for which we consider call-by-

value evaluation. For complexity reasons, we present two different algorithms to compute the intersection of a 2MTT with a deterministic bottom-up tree automaton and with the complement of a deterministic top-down tree automaton, respectively. Then, we construct again an ATWA defining the domain of the resulting macro tree walking transducer. With the algorithm to decide emptiness for an ATWA of the previous chapter, type checking for 2MTTs is decidable in exponential time. For input-linear macro tree walking transducers, i.e., strictly 1-bounded 2MTTs, it takes polynomial time only (Section 12.3).

Finally, we consider in **Chapter 13** macro forest walking transducers. Here, the output type is given by a finite forest monoid (FFM). We present an algorithm to compute the intersection of a 2MFT with an FFM representing the complement of the output type. Since the construction of the complement FFM, in general, is exponential, we also give an algorithm intersecting output-linear 2MFTs with bottom-up forest automata. A bottom-up forest automaton is a bottom-up tree automaton accepting the binary tree representations of forests. Then, an ATWA (working on enriched encodings of forests) is constructed for every macro forest walking transducer. The notion of strict b -boundedness is also generalized for 2MFTs.

Chapter 9

Tree Walking Transducers

Nondeterministic tree transducers describe transformations τ from trees to sets of trees over a ranked alphabet Σ , i.e., $\tau : \mathcal{T}_\Sigma \rightarrow 2^{\mathcal{T}_\Sigma}$. Consider for example a transformation, which translates documents describing company structures as the one in Figure 1.2 into a collection of all employees, which are listed under a new root node labeled **staff**. Besides a **data** element, these new **employee** elements now contain an element **boss** if the employee is the subordinate of someone. For our example document, the transformation produces:

```
<staff>
  <employee>
    <data> <name> Charles Montgomery Burns </name> ... </data>
  </employee>
  <employee>
    <data> <name> Waylon Smithers </name> ... </data>
    <boss> <name> Charles Montgomery Burns </name> ... </boss>
  </employee>
  <employee> ... </employee> ...
  <employee> ... </employee> ...
</staff>
```

The corresponding tree is referred as v_B and its binary encoding (cf. Page 14) as v_B' . The latter is depicted on the right in Figure 9.1. Remember that the input tree of this example is abbreviated by u_B and its encoding by u_B' , illustrated in Figures 1.2(b) and 2.1, respectively.

A tree walking transducer starts at the root of the input tree. Depending on the label of the current node, the direction and the state, it produces a tree with leaves, which again may contain state calls for nodes of the input tree. These recursively accessed nodes are determined according to the directives specified in the right-hand side of the applied rule: On directive *up*, the father of the current node is processed, on directive *down_i*, the *i*-th child and on directive *stay* the current node itself. Tree walking transducers can be considered as generalizations of *top-down* tree transducers. While top-down tree transducers

are only allowed to move downward in the input tree, tree walking transducers may also stay at the current node or move upward in the tree.

Example 9.1. Using our representation of forests by binary trees (Page 14), the transformation of our example is realized by a tree walking transducer M_{staff} with the following rules.

- 1 $q_I(\text{department}) \rightarrow \text{staff}(q(\text{down}_1), \mathbf{e})$,
- 2 $q(\text{employee}) \rightarrow \text{employee}(\text{data}(q_{\text{data}}(\text{down}_1), q_{\text{boss}}(\text{stay})),$
 $q_{\text{sub}}(\text{down}_1))$,
- 3 $q(\mathbf{e}) \rightarrow q_{\text{up}}(\text{up})$,

together with a state q_{data} for copying the personal data

- 4 $q_{\text{data}}(\text{data}) \rightarrow \text{copy}(\text{down}_1)$,

as well as a state q_{boss} to find the boss

- 5 $q_{\text{boss}}(\text{employee}) \rightarrow q_{\text{boss}}(\text{up})$,
- 6 $q_{\text{boss}}(\text{department}) \rightarrow \mathbf{e}$,
- 7 $q_{\text{boss}}(\text{subordinates}) \rightarrow \text{boss}(q_{\text{data}}(\text{up}), \mathbf{e})$,

and a state q_{sub} , which processes the subordinates

- 8 $q_{\text{sub}}(\text{data}) \rightarrow q_{\text{sub}}(\text{down}_2)$,
- 9 $q_{\text{sub}}(\text{subordinates}) \rightarrow q(\text{down}_1)$,
- 10 $q_{\text{sub}}(\mathbf{e}) \rightarrow q_{\text{next}}(\text{up})$.

The state q_{next} searches (in dfs-manner) the next employee

- 11 $q_{\text{next}}(\text{data}) \rightarrow q_{\text{next}}(\text{up})$,
- 12 $q_{\text{next}}(\text{employee}) \rightarrow q(\text{down}_2)$,

together with a state q_{up} for going to the boss if there is no further subordinate

- 13 $q_{\text{up}}(\text{employee}) \rightarrow q_{\text{up}}(\text{up})$,
- 14 $q_{\text{up}}(\text{subordinates}) \rightarrow q_{\text{next}}(\text{up})$,
- 15 $q_{\text{up}}(\text{department}) \rightarrow \mathbf{e}$

where state *copy* in Line 4 is meant to copy the content of **data** (i.e., the left child in the binary representation). The initial state is q_I , which means that we start with state q_I at the root of the tree. The output trees of this transformation are binary representations of the lists of all members of staff. The root, which is labeled **staff**, has a right child with label **e**. The left child of **staff** has label **employee** whose left child is a **data**-node (with the personal data and the boss) and whose right child is a chain of **employee**-nodes. Figure 9.1 illustrates this transformation for the binary example tree u_B' resulting in the tree v_B' . \triangleleft

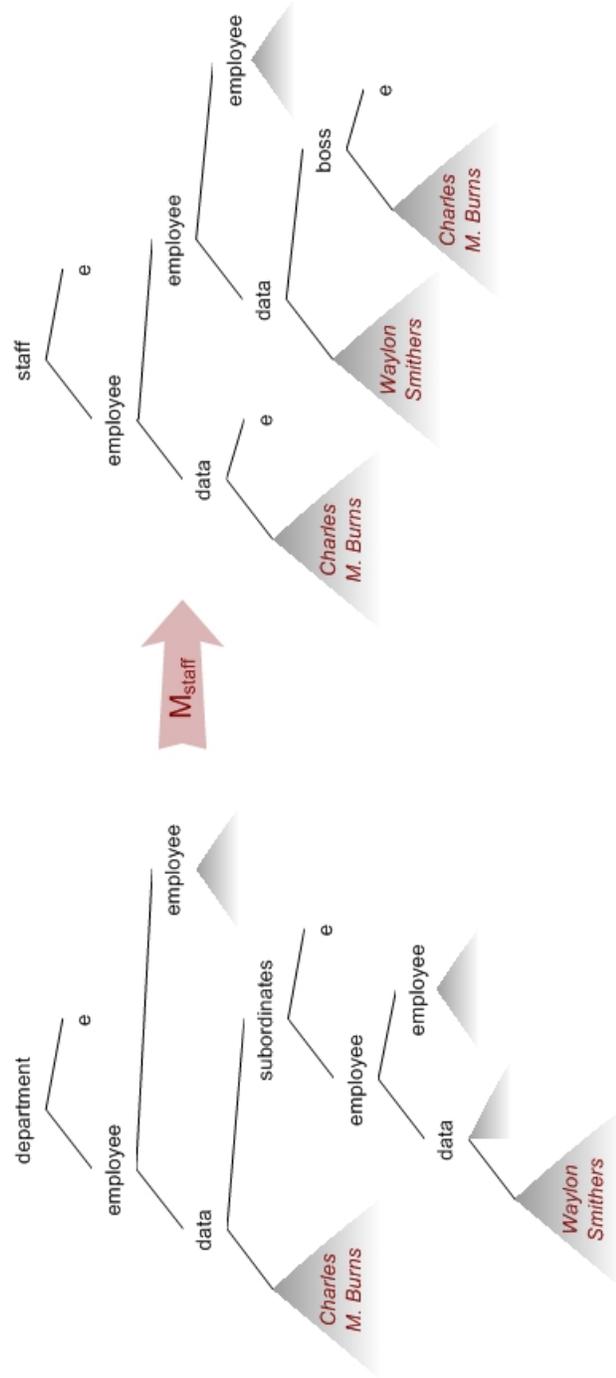


Figure 9.1: The tree u_B' and its output tree v_B' of the transformation of the 2TT M_{staff} .

The example illustrates that the “first-child next-sibling” encoding of forests implies that the *up*-operation of the tree walking transducer may not necessarily access directly the father in the forest representation but may instead reach the *left* sibling – depending whether or not the current node is a left or right child (i.e., has direction 1 or 2). In the example this was no problem: The state q_{boss} simply proceeds upwards in the tree representation until a node with the right label is reached. A direct construction of forest walking transducers, which provides the operations *up*, *down*, *left*, and *right* will be presented in Chapter 13. For the moment, we restrict ourselves to tree walking transducers on ranked trees (which perhaps are encodings of unranked forests).

Formally, the rules of a tree walking transducer are slightly more general than the ones shown in Example 9.1: Additional to the label of the current node, the left-hand side of a rule also checks the direction of the current node, i.e., whether the current node is the root node (direction is 0), or whether it is the i -th child of its parent node. It is well-known that in the case of tree walking automata (viz. tree walking transducers with output symbols $\{0, 1\}$ of rank 0), such direction tests (or “child number” test) are crucial: Without them, the automaton cannot even realize a depth-first left-to-right traversal over the input tree, i.e., it cannot systematically search through every node of the input. For some translations, however, direction tests are not needed (such as in our Example 9.1). In that example, we must think of every rule as existing in (at most) three incarnations, for direction 0 (root node), direction 1 (left child), and direction 2 (right child). For instance, the q -rule for **employee**-nodes (rule number 2 of the example) is needed in the following two incarnations:

$$\begin{aligned} 2a \quad & q(\mathbf{employee}, 1) \rightarrow \mathbf{employee}(\mathbf{data}(q_{data}(\mathit{down}_1), q_{boss}(\mathit{stay})), q_{sub}(\mathit{down}_1)) \\ 2b \quad & q(\mathbf{employee}, 2) \rightarrow \mathbf{employee}(\mathbf{data}(q_{data}(\mathit{down}_1), q_{boss}(\mathit{stay})), q_{sub}(\mathit{down}_1)) \end{aligned}$$

Recall from the Preliminaries (Chapter 2) that the maximal rank of symbols in a ranked alphabet Σ is denoted by $mr(\Sigma)$.

Definition 9.1 (2TT). A *tree walking transducer* M (2TT for short) is a tuple (Q, Σ, R, Q_0) where Q is a set of states, Σ is a ranked alphabet, $Q_0 \subseteq Q$ is a set of initial states, and R is a finite set of rules. A rule is of the form $q(\mathbf{a}, \eta) \rightarrow \zeta$ where $q \in Q$, $\mathbf{a} \in \Sigma^{(m)}$, $m \geq 0$, $\eta \geq 0$ and ζ is a tree generated by the grammar

$$\zeta ::= \mathbf{b}(\underbrace{\zeta, \dots, \zeta}_{m' \text{ times}}) \mid q'(\mathit{op}),$$

with $\mathbf{b} \in \Sigma^{(m')}$, $m' \geq 0$, $q' \in Q$, and $\mathit{op} \in \{\mathit{stay}, \mathit{up}\} \cup \{\mathit{down}_i \mid 1 \leq i \leq m\}$.

Tree walking transducers are also called 2-way tree transducers because they generalize to trees the well known concept of 2-way finite state transducer on words (see, e.g., [Gre78]).

Conventionally, tree transducers are defined over two ranked alphabets of input and output symbols. In Definition 9.1 of a 2TT M , we only use one alphabet Σ , which contains input and output symbols. If we want to distinguish

the two, we say that $\mathbf{a} \in \Sigma$ is an *input symbol* if \mathbf{a} appears on the left-hand side of a rule of M ; we say that it is an *output symbol* if it appears in the right-hand side of a rule of M . In Example 9.1, **data** is an input and output symbol and **boss** is an output symbol of M_{staff} .

In practice, transducers also have to cope with *unknown* labels in the input such as, e.g., portions of text, which then either are ignored or copied into the output. In order to deal with this, we could simply extend our formalism by an extra symbol \bullet of any given rank which serves as a placeholder for unknown labels of this rank. This idea can be extended to placeholders for unknown elements of different atomic types, for instance **String**, **Number**, or **Date**. Thus, we can describe the so called “Simple Types” of XML Schema (cf. [FW04]).

For a right-hand side ζ , we also write $\zeta = z[q_1(op_1), \dots, q_c(op_c)]$ to refer to all occurrences of state calls in the right-hand side; there, $z \in \mathcal{T}_\Sigma(X)$ is a tree, which contains exactly one occurrence of the variable x_i for $i = 1, \dots, c$. Note that z does not contain state calls. For example the right-hand side of the rule in Line 2 in the Example 9.1 can be written as

$$z[q_{\text{data}}(\text{down}_1), q_{\text{boss}}(\text{stay}), q_{\text{sub}}(\text{down}_1)]$$

where $z = \text{employee}(\text{data}(x_1, x_2), x_3)$.

A 2TT is called *deterministic* iff there is at most one state in the set Q_0 and for every triple (q, \mathbf{a}, η) of a state, a symbol, and a direction, there is at most one rule with $q(\mathbf{a}, \eta)$ as left-hand side. The example 2TT M_{staff} is deterministic.

Intuitively, the meaning of the expressions of a right-hand side is as follows: The output can either be an element \mathbf{b} whose content is recursively determined, or a recursive call to some state q' on the current input node, on its father or on its i -th subtree. The match patterns in the left-hand side of the rules are restricted to the form “ \mathbf{a}, η ”, i.e., it is only allowed to check the label of the current input node and its direction. Thus, the transformation of a 2TT M starts at the root node of the input u with one of the initial states. A state q can be applied to an input node ϑ with label $\text{lab}(\vartheta) = \mathbf{a}$ and direction $\eta = \eta(\vartheta)$ if there is a rule with left-hand side $q(\mathbf{a}, \eta)$. The evaluation continues on a child ϑ_i of ϑ for each occurrence of a state call $q'(\text{down}_i)$, at ϑ itself for each occurrence of a state call $q'(\text{stay})$, and at the parent of ϑ , for each occurrence of a state call $q'(\text{up})$.

Hence, the *meaning* $\llbracket q \rrbracket_u$ of a state q of M with respect to an input tree u can be defined as a function from the nodes (of the input tree) to sets of trees, i.e., $\llbracket q \rrbracket_u : \text{Nodes}(u) \rightarrow 2^{\mathcal{T}_\Sigma}$. The values $\llbracket q \rrbracket_u$ for all q are jointly defined as the *least* functions satisfying:

$$\llbracket q \rrbracket_u(\vartheta) \supseteq (\llbracket \zeta \rrbracket_u(\vartheta)) \quad \text{for rule } q(\mathbf{a}, \eta) \rightarrow \zeta$$

where ϑ is a node of u with $\text{lab}(\vartheta) = \mathbf{a}$ and $\eta = \eta(\vartheta)$ with

$$\begin{aligned} \llbracket \mathbf{b}(\zeta_1, \dots, \zeta_m) \rrbracket_u(\vartheta) &= \{\mathbf{b}(u'_1, \dots, u'_m) \mid u'_i \in \llbracket \zeta_i \rrbracket_u(\vartheta)\} \\ \llbracket q'(op) \rrbracket_u(\vartheta) &= \llbracket q' \rrbracket_u(\llbracket op \rrbracket_u(\vartheta)) \end{aligned}$$

where op stands for $stay$, up or $down_i$ for $1 \leq i \leq rank(\mathbf{a})$, and $\llbracket op \rrbracket_u$ is defined by:

$$\begin{aligned}\llbracket stay \rrbracket_u(\vartheta) &= \vartheta \\ \llbracket down_i \rrbracket_u(\vartheta) &= \vartheta i \\ \llbracket up \rrbracket_u(\vartheta i) &= \vartheta.\end{aligned}$$

The transformation τ_M realized by the 2TT M on an input tree u and sets U of input trees, respectively, is defined by

$$\tau_M(u) = \bigcup \{ \llbracket q_0 \rrbracket_u(\epsilon) \mid q_0 \in Q_0 \} \quad \text{and} \quad \tau_M(U) = \bigcup \{ \tau_M(u) \mid u \in U \}.$$

For a deterministic tree walking transducer M , the transformation τ_M is a partial function $\tau_M : \mathcal{T}_\Sigma \dashrightarrow \mathcal{T}_\Sigma$. The *domain* of the transducer is the domain of the transformation, i.e., $dom(M) = dom(\tau_M) = \{u \mid \tau_M(u) \neq \emptyset\}$. As usual, the *size* $|M|$ of a 2TT M is the sum of the sizes of all its rules where the size of a rule $q(\mathbf{a}, i) \rightarrow \zeta$ is defined as $3 + |\zeta|$. Recall that $|\zeta|$ equals the number of nodes of ζ .

Applying the 2TT M_{staff} from before to $u_{B'}$, we obtain the tree $v_{B'}$. The right-hand sides of rules in a 2TT may be arbitrarily large and contain arbitrarily many state calls. Dealing with such rules increases the complexity of some algorithms on 2TTs. Thus, we give a normal form for 2TTs where the number of state calls in right-hand sides is bounded by the maximal rank of output symbols. In the particular case where we consider binary representations of forests, the number of state calls in right-hand sides can be restricted to 2.

Lemma 9.1. *For every 2TT M , a 2TT M' can be constructed in time $\mathcal{O}(|M|)$ such that*

1. $\tau_{M'} = \tau_M$ and
2. *the right-hand side of each rule of M' contains at most k occurrences of states*

where k is the maximal rank of the output symbols of M .

Proof. Let $M = (Q, \Sigma, R, Q_0)$. Intuitively, the idea of the construction is to introduce auxiliary states for all proper subtrees which contain more than one state call. For a symbol $\mathbf{a} \in \Sigma$ and direction η , let $Z_{\mathbf{a}, \eta}$ denote the set of all subterms with more than one state call in right-hand sides of rules for \mathbf{a}, η . For each $\zeta \in Z_{\mathbf{a}, \eta}$, we introduce a fresh state $q_{\mathbf{a}, \eta, \zeta}$. Assume that $\zeta = \mathbf{b}(\zeta_1, \dots, \zeta_m)$. Then we introduce the new rule

$$q_{\mathbf{a}, \eta, \zeta}(\mathbf{a}, \eta) \rightarrow \mathbf{b}(\zeta'_1, \dots, \zeta'_m)$$

where $\zeta'_j = \zeta_j$ if ζ_j contains at most one occurrence of a state, and otherwise $\zeta'_j = q_{\mathbf{a}, \eta, \zeta_j}(\text{stay})$. We construct $M' = (Q', \Sigma, R', Q_0)$ as follows. The set of rules R' of the new transducer consists of all these newly constructed rules.

Additionally, we add for every rule $q(\mathbf{a}, \eta) \rightarrow \mathbf{b}(\zeta_1, \dots, \zeta_m)$ of M a new rule $q(\mathbf{a}, \eta) \rightarrow \mathbf{b}(\zeta'_1, \dots, \zeta'_m)$ where for every j , $\zeta'_j = \zeta_j$ if ζ_j contains at most one occurrence of a state, and $\zeta'_j = q_{\mathbf{a}, \eta, \zeta_j}(\text{stay})$ otherwise. The set of states Q' contains all states of Q and, additionally, the new states $q_{\mathbf{a}, \eta, \zeta}$ for every symbol $\mathbf{a} \in \Sigma$, direction η , and every term $\zeta \in Z_{\mathbf{a}, \eta}$.

The resulting transducer M' has a new state at most for every non-leaf node of a right-hand side of a rule in M . Thus, in the worst case, we have at most $|M|$ new states. In the new rules, the right-hand side of the original rule of M is split in its subtrees. Thereby, we have $|M'| \in \mathcal{O}(|M|)$. \square

In order to describe the behavior of the 2TT $M = (Q, \Sigma, R, Q_0)$ on a fixed input tree u , we are also going to define *runs* of M . A run can itself be described by a ranked tree over the set of rules. Here, the rank of a rule $q(\mathbf{a}, \eta) \rightarrow \zeta$ is given by the number of occurrences of calls $q'(\text{op})$ in ζ to states q' in Q .

Definition 9.2 (Run). Let q denote a state of M and ϑ a node in the input tree u of direction η , which is labeled with \mathbf{a} . Assume that $r : q(\mathbf{a}, \eta) \rightarrow \zeta$ is a rule in R with $\zeta = z[q_1(\text{op}_1), \dots, q_m(\text{op}_m)]$. Then the tree

$$\rho = r(\rho_1, \dots, \rho_m) \in \mathcal{T}_R$$

is a (q, ϑ) -run of the 2TT M on the tree u if for every $i \in [m]$, ρ_i is a (q_i, ϑ_i) -run of M on u where ϑ_i is obtained from ϑ by operation op_i . The *output* $\tau(\rho)$ produced by a run ρ is defined by

$$\tau(\rho) = z[\tau(\rho_1), \dots, \tau(\rho_m)].$$

A (q_0, ϵ) -run for an initial state q_0 is also called *accepting run* of M on u .

If M is deterministic, there exists at most one accepting run on every tree.

Example 9.2. Figure 9.2(a) shows the behavior of the (deterministic) example 2TT M_{staff} on the tree u_e , which describes a department with one employee:

$$u_e = \text{department}(\text{employee}(\text{data}(\dots, \mathbf{e}), \mathbf{e}), \mathbf{e})$$

All states in an oval around a node are applied to this node. The picture includes the dependences of the states. For example, consider the employee node $\vartheta = 1$ (with $\text{lab}(\vartheta) = \text{employee}$). There, we have the state q . In the 2TT, there is just one rule with left-hand side $q(\text{employee}, 1)$:

$${}_{2a} \quad q(\text{employee}, 1) \rightarrow \text{employee}(\text{data}(q_{\text{data}}(\text{down}_1), q_{\text{boss}}(\text{stay})), q_{\text{sub}}(\text{down}_1))$$

Thus, we have a $(q, 1)$ -run $\rho = r_{2a}(\rho_1, \rho_2, \rho_3)$ where ρ_1, ρ_3 are $(q_{\text{data}}, 1.1)$ - and $(q_{\text{sub}}, 1.1)$ -runs, respectively, and ρ_2 is a $(q_{\text{boss}}, 1)$ -run. This is illustrated by the three arrows starting at q at node ϑ . The Figure 9.2(b) shows a (q_I, ϵ) -run $\rho' = r_1(\rho)$. The state *copy* was not detailed in Example 9.1. Accordingly, the $(\text{copy}, 1.1.1)$ -run here is not complete. The output $\tau(\rho)$ of this run is the tree $\text{staff}(\text{employee}(\text{data}(\dots, \mathbf{e}), \mathbf{e}), \mathbf{e})$. \triangleleft

Another approach to define the semantics of a 2TT is via accepting runs. Indeed, this operational semantics of a 2TT coincides with the denotational semantics provided first.

Theorem 9.2. *For a tree u and a 2TT M , the following two statements are equivalent.*

1. *There is an accepting run ρ of M for u with $\tau(\rho) = v$*
2. *$v \in \tau_M(u)$.*

Theorem 9.2 can be proved by fixpoint induction. The denotational view on the semantics of a 2TT allows us to use fixpoint arguments for proving the correctness of constructions, whereas the operational view is better suited for combinatorial arguments.

9.1 Notes and References

Top-down tree transducers were invented by Rounds and Thatcher [Rou70, Tha69]. Top-down tree transducers terminate for every input tree because they process the input tree strictly top-down. While the height increase of a top-down tree transducer is at most linear, the size increase is at most exponential (viz. the translation of a monadic tree with n nodes into a full binary tree of height n). A nondeterministic top-down tree transducer can associate at most double exponentially many output trees to a given input tree; e.g., the transducer with the five rules $q(\mathbf{a}, \eta) \rightarrow \mathbf{b}(q(\text{down}_1), q(\text{down}_1))$, $q(\mathbf{a}, \eta) \rightarrow \mathbf{c}(q(\text{down}_1), q(\text{down}_1))$ for $\eta \in \{0, 1\}$, and $q(\mathbf{e}, 1) \rightarrow \mathbf{e}$. Tree walking transducers with output strings were invented in [AU71]; by adding the ability to generate output trees rather than strings, we obtain the tree walking transducer of this thesis. It can be seen as the k -pebble tree transducer of [MSV03], for the case that $k = 0$. In [KS81], it was shown that tree walking transducers without child number test are not useful: They cannot even check whether all leaves of input trees are labeled by some symbol \mathbf{a} . As mentioned in [EM03b], in the total deterministic case the tree walking transducer is essentially the same as the attribute grammar [Knu68]. Similar to the fact that circularity of attribute grammars is decidable, it is possible to change any deterministic tree walking transducer in such a way that all runs are terminating [EM03b]. This is not possible for nondeterministic tree walking transducers because they can associate infinitely many output trees to a given input tree (viz. the transducer with the two rules $q(\mathbf{a}, 0) \rightarrow \mathbf{b}(q(\text{stay}))$ and $q(\mathbf{a}, 0) \rightarrow \mathbf{e}$). The normal form of Lemma 9.1 is similar to the one for pebble macro tree transducers given in Theorem 16 of [EM03b]. Attribute grammars with tree output are also called “attributed tree transducers” [Fül81]; for total deterministic such transducers (which coincide with our tree walking transducers when they are total deterministic) it is known that the size-to-height relationship of input tree to output tree is linear, and that the number of different output subtrees in an output tree is linear in the size of the corresponding input tree (see, e.g., [FV98]).

Chapter 10

Type Checking

In this chapter, we present general techniques for certifying that all outputs produced by a transducer M for trees of a given input type are *well-formed*, i.e., comply with some given output type O . This problem is called *type checking* of the transducer M . Here, a type is just a set of trees, i.e., a *tree language*. Clearly, the tractability of type checking heavily depends on the class of languages used as types, and the class of transformations.

The binary tree v_B' is an example for the output language of the 2TT M_{staff} (cf. Figure 9.1). Such output trees are binary trees with a root labeled with **staff** and a right-comb of **employee** nodes as left subtree. It is the first-child next-sibling representation of a tree, which has a root labeled with **staff** and arbitrary many **employee** nodes as children. The following DTD (Document Type Definition, [BPS08]) is describing this type (not the binary representation) where **content** stands for further personal data which are not specified here:

```
<!ELEMENT staff (employee)*>
<!ELEMENT employee (data, boss)>
<!ELEMENT data (name, content)>
<!ELEMENT name (#PCDATA)>
<!ELEMENT boss (name, content)>
<!ELEMENT content ... >
```

10.1 Type Checking by Forward Type Inference

Type checking a transducer M means to verify that all trees produced by M for input trees in the given input type I are necessarily contained in the given output type O . If τ is the transformation induced by the transducer M , we want to check whether or not

$$\tau(I) \subseteq O$$

where $\tau(I) = \{\tau(u) \mid u \in I\}$. If this check succeeds, then we say that M *type checks w.r.t. I and O* . We solve this problem by forward inference, i.e., we

determine whether

$$\tau(I) \cap \overline{O} = \emptyset$$

where \overline{O} is the complement of the type O . In order to decide emptiness of this intersection, we proceed in two steps. First, we construct from transducer M a transducer $M_{\overline{O}}$ which produces only those outputs of M , which are from \overline{O} . For 2TTs, this construction is presented in Section 11.1. Second, we present methods for deciding emptiness of transducers (w.r.t. I).

10.2 Tree Automata

There are several specification formalisms for XML types, such as DTD, XML Schema, or RELAX NG. For our purpose, the particular type formalisms is not essential, as all of these formalisms can be abstracted by recognizable (or regular) tree languages. Thus, each type definition can be translated into a finite tree automaton. XML Schema specifications, e.g., can be considered as simple classes of deterministic top-down automata. We briefly recall crucial definitions of finite tree automata as needed in this chapter. Several kinds of tree automata and basic constructions are presented, e.g., in [Löd11, CDG⁺07]. Recall that we denote the maximal rank of a ranked alphabet Σ by $mr(\Sigma)$.

Definition 10.1 (BTA). Let Σ be a ranked alphabet. A *bottom-up finite state tree automaton* A (over Σ), *BTA* for short, is a tuple (P, Σ, δ, F) where

- P is a finite set of states,
- $F \subseteq P$ is a set of accepting states, and
- $\delta \subseteq \bigcup_{m=0}^{mr(\Sigma)} (P \times \Sigma^{(m)} \times P^m)$ is a finite set of transitions of the form $(p, \mathbf{a}, p_1 \dots p_m)$ where $\mathbf{a} \in \Sigma^{(m)}$ and $p, p_1, \dots, p_m \in P$.

A transition $(p, \mathbf{a}, p_1 \dots p_m)$ denotes that, if for all $i \in [m]$, BTA A arrives in state p_i after processing some tree u_i , then it can assign state p to the tree $\mathbf{a}(u_1, \dots, u_m)$. Technically, a *p-run* ρ of A on a tree $u = \mathbf{a}(u_1, \dots, u_m) \in \mathcal{T}_\Sigma$ is a tree

$$\rho = r(\rho_1, \dots, \rho_m) \in \mathcal{T}_\delta$$

where r is a transition $(p, \mathbf{a}, p_1 \dots p_m) \in \delta$ and ρ_i is a p_i -run of A for u_i . For some applications, it suffices to represent p -runs by trees from \mathcal{T}_P with root p . The tree language $\mathcal{L}(A)$, accepted by A , consists of the trees $u \in \mathcal{T}_\Sigma$ by which A can reach an accepting state, i.e. it exists a p -run ρ of A for u with $p \in F$; the latter run is called *accepting run of A on u* . A bottom-up tree automaton $A = (P, \Sigma, \delta, F)$ is *deterministic* (DBTA) if for each symbol $\mathbf{a} \in \Sigma^{(m)}$ and every sequence $p_1 \dots p_m$ of states, there is at most one state p with $(p, \mathbf{a}, p_1 \dots p_m) \in \delta$, i.e., δ induces a partial function of type $\Sigma \times P^* \dashrightarrow P$. A BTA is called *complete* if there is at least one rule $(p, \mathbf{a}, p_1 \dots p_m) \in \delta$ for all $m \geq 0$, $\mathbf{a} \in \Sigma^{(m)}$, and $p_1, \dots, p_m \in P$.

We may also interpret the transitions of a BTA in a *top-down* fashion. Then we obtain the known top-down tree automaton (TTA), which starts at the input root node and assigns states to the children of a node, depending on the label of the node and the current state.

Definition 10.2 (DTTA). A BTA is called *deterministic top-down* (DTTA for short) if the set of final states is a singleton set, and the transition relation δ induces a partial function $P \times \Sigma \dashrightarrow P^*$, i.e., for each state $p \in P$ and each symbol $\mathbf{a} \in \Sigma^{(m)}$, there is at most one sequence of states $p_1 \dots p_m \in P^m$ with $(p, \mathbf{a}, p_1 \dots p_m) \in \delta$.

As usual, the *size* $|A|$ of a finite state tree automaton A is the sum of sizes of all its transitions and the number of states where a transition $(p, \mathbf{a}, p_1 \dots p_m)$ has size $m+2$. Let BTA, DBTA, and DTTA denote the classes of all languages definable by BTAs, DBTAs, and DTTAs, respectively. It is known that BTA = DBTA equals the class of regular tree languages, and that DTTA is properly contained in this class.

Example 10.1. Coming back to the transformation from Example 9.1, the set of valid output documents should be lists of staff members (cf. the DTD on Page 149). More precisely:

- **staff** should contain a possibly empty sequence of **employee** elements;
- Each **employee** element should contain a **data** element and optionally, a **boss** element.

A bottom-up tree automaton describing (the binary representations of) this set is given by $A_{\text{staff}} = (P, \Sigma, \delta, F)$ where $P = \{r_{\text{staff}}, r_{\text{empl}}, r_{\text{data}}, r_{\text{name}}, r_{\text{boss}}, r_{\mathbf{e}}, \dots\}$ and δ contains (amongst others) the transitions

$$\begin{array}{lll}
 (r_{\text{staff}}, & \mathbf{staff}, & r_{\text{empl}} r_{\mathbf{e}}), \\
 (r_{\text{empl}}, & \mathbf{employee}, & r_{\text{data}} r_{\text{empl}}, \quad (r_{\text{empl}}, \quad \mathbf{e}), \\
 (r_{\text{data}}, & \mathbf{data}, & r_{\text{name}} r_{\text{boss}}), \\
 (r_{\text{boss}}, & \mathbf{boss}, & r_{\text{name}} r_{\mathbf{e}}), \quad (r_{\text{boss}}, \quad \mathbf{e}), \\
 (r_{\text{name}}, & \mathbf{name}, & r_{\text{content}} r_{\mathbf{e}}), \\
 (r_{\mathbf{e}}, & \mathbf{e} &)
 \end{array}$$

where r_{content} is the state characterizing valid personal data of employees. The set of accepting states is, thus, given by $F = \{r_{\text{staff}}\}$. Note that this BTA is, in fact, deterministic top-down. Note further that the documents adhering to the DTD on the beginning of Chapter 10 are, in general, no ranked trees but forests. \triangleleft

In the previous example, the BTA ran on the first-child next-sibling encoding of forests as binary ranked trees. For convenience, we call such a finite tree automaton also *finite forest automaton* (short: BFA). Again, if it is deterministic or deterministic top-down, we abbreviate it DBFA and DTFA, respectively.

10.3 Basic Properties of BTAs

The approach, which we advocate here, is called *forward type checking* (cf. Section 10.1). Assume that O is the type of all valid output trees. In order to check that the transducer M produces only outputs in O , we construct a transducer, which, for every input u , only produces those output trees of transducer M , which are *not* valid, i.e., which are in \overline{O} (the complement of O). Thus, type correctness for M is reduced to emptiness of the auxiliary transducer $M_{\overline{O}}$. For this idea to work, it is useful to have effective constructions which take the specification of a type and return a specification for its complement. For a *complete* DBTA $A = (P, \Sigma, \delta, F)$, this construction is simple: We need to exchange accepting and non-accepting states, i.e., replace F with $P \setminus F$. Since every regular tree language can be accepted by a complete DBTA, this construction implies that the complement of a regular tree language is a regular tree language, too. The complement of a type described by a deterministic top-down tree automaton is a regular language as well, but not necessarily in DTTA. The obvious technique for constructing an automaton for the complement is, therefore, to transform the deterministic top-down automaton into a complete deterministic bottom-up automaton and then apply the complement construction for complete DBTAs. This first construction, however, possibly incurs an exponential blow-up in the number of states. Therefore, we approve a different approach: Instead of constructing a *deterministic* automaton for the complement, we construct a *non-deterministic* automaton. The latter can be achieved by only moderately increasing the size.

Lemma 10.1. *For a DTTA A over the ranked alphabet Σ , there is a BTA A' over Σ with $\mathcal{L}(A') = \mathcal{T}_\Sigma \setminus \mathcal{L}(A)$ and $|A'| \in \mathcal{O}((|A| + |P| \cdot |\Sigma|) \cdot mr(\Sigma))$ where P is the set of states of A .*

Proof. Intuitively, the automaton A' guesses a path in the input tree to some node where the original automaton A fails. Formally, let $A = (P, \Sigma, \delta, \{p_0\})$ be a DTTA and define BTA $A' = (P', \Sigma, \delta', \{p'_0\})$ with $P' = \{p' \mid p \in P\} \cup \{\bullet\}$ for a new state $\bullet \notin P$. A state p' is meant to generate only trees for which there is *no* p -run of A . The state \bullet describes arbitrary trees, i.e., the language \mathcal{T}_Σ . The set δ' of transitions of the new BTA is defined as follows:

- For every transition $(p, \mathbf{a}, p_1 \dots p_m) \in \delta$ with $m = rank(\mathbf{a}) \geq 1$, and for every $i \in [m]$ let $(p', \mathbf{a}, \bullet^{i-1} p'_i \bullet^{m-i}) \in \delta'$.
- For every state $p \in P$, $0 \leq m \leq mr(\Sigma)$, and $\mathbf{a} \in \Sigma^{(m)}$, $(p', \mathbf{a}, \bullet^m) \in \delta'$ whenever $\forall p_1, \dots, p_m \in P : (p, \mathbf{a}, p_1 \dots p_m) \notin \delta$.
- For every $\mathbf{a} \in \Sigma^{(m)}$, $(\bullet, \mathbf{a}, \bullet^m) \in \delta'$.

For the correctness of the construction, we claim that for every state p of the DTTA A , and every input tree u , the BTA A' has a p' -run on u iff DTTA A has no p -run on u . This claim can be proven by induction on the height of input trees.

Now let $k = mr(\Sigma)$ be the maximal rank of symbols in Σ . For each transition in δ , we get at most k new transitions in δ' (one for each successor state). Additionally, we possibly require a new rule of length at most $k + 2$ for each pair of a state in P' and a symbol in Σ . Thus, the size of the new automaton A' is in $\mathcal{O}((|A| + |P| \cdot |\Sigma|) \cdot mr(\Sigma))$. \square

Example 10.2. For the DTTA $A_{\text{staff}} = (P, \Sigma, \delta, \{r_{\text{staff}}\})$ in Example 10.1, the bottom-up tree automaton $A'_{\text{staff}} = (P', \Sigma, \delta', \{r'_{\text{staff}}\})$ for the complement has the following transitions for the label **staff**:

$$(r'_{\text{staff}}, \text{staff}, r'_{\text{empl}}\bullet), \quad (r'_{\text{staff}}, \text{staff}, \bullet r'_e), \quad (r', \text{staff}, \bullet\bullet)$$

for all $r' \in P' \setminus \{r'_{\text{staff}}\}$. For the label **boss** it has the transitions:

$$(r'_{\text{boss}}, \text{boss}, r'_{\text{name}}\bullet), \quad (r'_{\text{boss}}, \text{boss}, \bullet r'_e), \quad (r', \text{boss}, \bullet\bullet)$$

for all $r' \in P' \setminus \{r'_{\text{boss}}\}$. The set of accepting states is $\{r'_{\text{staff}}\}$. \triangleleft

10.4 Notes and References

XML type definition languages such as DTDs [BPS08], XML Schema [FW04], or RELAX NG [CM01] are closely related to the regular tree languages [MLM00, Nev02] that is, to the class of tree languages recognized by finite tree automata.

Tree automata are a well studied formalism in computer science, dating back to the late 1960s. For surveys on tree automata, please see [GS84, GS97, CDG⁺07]. Tree automata inherit most of the good properties of finite automata on strings, such as effective closure under Boolean operations and decidability of emptiness. An important property, which will be used later for type checking, is that emptiness of BTAs can be decided in linear time (see, e.g., [Löd11], Theorem 1.7.4 in [CDG⁺07]).

Theorem 10.2. *Given a BTA A it can be decided in linear time whether or not $\mathcal{L}(A) = \emptyset$.*

Just as in the string case, nondeterministic bottom-up tree automata can be determinized (with a potential and sometimes unavoidable exponential blow up in automaton size). This is not the case for top-down tree automata: The class DTTA of languages accepted by deterministic top-down tree automata is a strict subclass of BTA, which does not even contain all finite languages. A famous example of a language, which is not in DTTA, is the set $U = \{\mathbf{f}(\mathbf{a}, \mathbf{b}), \mathbf{f}(\mathbf{b}, \mathbf{a})\}$. Note that for a given BTA, it is decidable if its language is in DTTA. This is due to the fact that DTTA languages can be characterized by the “path-closed” property [Cou78, Vir80]; the latter means that the trees in the languages are exactly obtained by combining all paths of the corresponding path language. The language U for instance is not path-closed. Using a similar example, it is easily shown that DTTA is not closed under complementation (and neither under union).

Chapter 11

Type Checking of Tree Walking Transducers

In this chapter, we present techniques to type check 2TTs against regular tree languages. We use the approach of forward type inference presented in the previous chapter. Thereto, we start with intersecting tree walking transducers with complements of output types where the output types are given by different types of tree automata.

11.1 Intersecting 2TTs with Output Types

For a given tree walking transducer, we build a second 2TT, which produces only output trees in the *complement* of the output type, and otherwise realizes the same transformation as the original transducer. If the output type is described by a complete DBTA $A = (P, \Sigma, \delta, F)$, the complement will be recognized by the complete DBTA $\bar{A} = (P, \Sigma, \delta, P \setminus F)$. Furthermore, for a given DTTA exists a BTA describing the complement (cf. Lemma 10.1). Thus, it is sufficient to construct a 2TT M_A for a 2TT M and a BTA A (which may be the complement automaton of a complete DBTA or of a DTTA) with $\tau_{M_A}(u) = \tau_M(u) \cap \mathcal{L}(A)$ for every tree u .

Theorem 11.1. *For every 2TT M and every BTA A , there is a 2TT M_A with*

$$\tau_{M_A}(u) = \tau_M(u) \cap \mathcal{L}(A)$$

for all $u \in \mathcal{T}_\Sigma$. The size $|M_A|$ of M_A is in $\mathcal{O}(|M| \cdot |A|^{d+1})$ where d is the maximal number of occurrences of states in right-hand sides of M .

Proof. Let $M = (Q, \Sigma, R, Q_0)$ and $A = (P, \Sigma, \delta, F)$. For each state q in Q and all states $p \in P$, we generate new states for M_A of the form $\langle q, p \rangle$. Such a state is meant to generate only trees $u \in \mathcal{T}_\Sigma$ for which there is a run of A starting at

the leaves and reaching the root of u in state p . The rules of the new 2TT M_A are

$$\langle q, p \rangle(\mathbf{a}, \eta) \rightarrow \zeta'$$

for every rule $q(\mathbf{a}, \eta) \rightarrow \zeta$ of M and $\zeta' \in \tau^p[\zeta]$. The sets $\tau^p[\cdot]$ are inductively defined by:

$$\begin{aligned} \tau^p[\mathbf{b}(\zeta_1, \dots, \zeta_m)] &= \{\mathbf{b}(\zeta'_1, \dots, \zeta'_m) \mid (p, \mathbf{b}, p'_1 \dots p'_m) \in \delta \wedge \forall i : \zeta'_i \in \tau^{p'_i}[\zeta_i]\} \\ \tau^p[q'(op)] &= \{\langle q', p \rangle(op)\} . \end{aligned}$$

The set of initial states of M_A is $Q'_0 = Q_0 \times F$. By fixpoint induction, we verify for every state q , every input tree $u \in \mathcal{T}_\Sigma$, every node $\vartheta \in \text{Nodes}(u)$, and every state p that

$$\llbracket \langle q, p \rangle \rrbracket_u(\vartheta) = \llbracket q \rrbracket_u(\vartheta) \cap \{v \in \mathcal{T}_\Sigma \mid \exists \text{ run } \rho \text{ on } v \text{ with } \rho(\epsilon) = p\} .$$

For each state in M we have at most $|A|$ new states in M_A . If we have c occurrences of state calls in the right-hand side of a rule r of M , with the state on the left-hand side, we obtain at most $|A|^{c+1}$ new rules for r in M_A . Therefore, the new 2TT is of size $\mathcal{O}(|M| \cdot |A|^{d+1})$ where d is the maximal number of occurrences of state calls in right-hand sides in M . \square

Considering only binary trees, we obtain size $\mathcal{O}(|M| \cdot |A|^3)$ for the intersection 2TT (with Lemma 9.1). The last step is to decide whether $\tau_{M_A} \neq \emptyset$. Thereto, we build a BTA describing the domain of M_A . This will be done after completing the example.

Example 11.1. Let us try to type check the 2TT $M_{\text{staff}} = (Q, \Sigma, R, Q_0)$ via forward type inference. According to Lemma 9.1, we restrict the maximal number of state calls in right-hand sides to 2. In our example 2TT, the rule

$$q(\text{employee}, \eta) \rightarrow \text{employee}(\text{data}(q_{\text{data}}(\text{down}_1), q_{\text{boss}}(\text{stay})), q_{\text{sub}}(\text{down}_2))$$

has three state calls. We obtain the new rules:

$$\begin{aligned} q(\text{employee}, \eta) &\rightarrow \text{employee}(q'(\text{stay}), q_{\text{sub}}(\text{stay})) \\ q'(\text{employee}, \eta) &\rightarrow \text{data}(q_{\text{data}}(\text{down}_1), q_{\text{boss}}(\text{stay})) . \end{aligned}$$

According to the proof of Lemma 9.1, the new state q' in these rules is the state

$$q_{\text{employee, data}(q_{\text{data}}(\text{down}_1), q_{\text{boss}}(\text{stay}))} .$$

As output type, consider again the DTTA $A_{\text{staff}} = (P, \Sigma, \delta, \{r_{\text{staff}}\})$. The complement BTA $A'_{\text{staff}} = (P', \Sigma, \delta', \{r'_{\text{staff}}\})$ is given in Example 10.2. Then, the intersection transducer $(M_{\text{staff}})_{A'_{\text{staff}}}$ is given by $(Q \times P', \Sigma, R', Q_0 \times \{r'_{\text{staff}}\})$. In what follows, we show how the first few rules in R' are constructed from R and δ' .

For $q_I(\text{department}, 0) \rightarrow \text{staff}(q(\text{down}_1), \mathbf{e})$ and r'_{staff} , we obtain the rule

$$\langle q_I, r'_{\text{staff}} \rangle(\text{department}, 0) \rightarrow \text{staff}(\langle q, r'_{\text{empl}} \rangle(\text{down}_1), \mathbf{e}) .$$

For the rule $q(\mathbf{employee}, \eta) \rightarrow \mathbf{employee}(q'(stay), q_{sub}(down_2))$ and r'_{empl} , we obtain

$$\begin{aligned} \langle q, r'_{empl} \rangle(\mathbf{employee}, \eta) &\rightarrow \mathbf{employee}(\langle q', r'_{data} \rangle(stay), \langle q_{sub}, \bullet \rangle(down_2)) \\ \langle q, r'_{empl} \rangle(\mathbf{employee}, \eta) &\rightarrow \mathbf{employee}(\langle q', \bullet \rangle(stay), \langle q_{sub}, r'_{empl} \rangle(down_2)) . \end{aligned}$$

For $q(\mathbf{e}, \eta) \rightarrow q_{up}(up)$ and all $r \in P'$, we obtain

$$\langle q, r \rangle(\mathbf{e}, \eta) \rightarrow \langle q_{up}, r \rangle(up) . \quad \triangleleft$$

11.2 Deciding Emptiness of 2TTs

In order to check the emptiness of a tree walking transducer w.r.t. a given input type, we construct a nondeterministic finite state automaton (cf. Section 10.2) which then is checked for emptiness. First, for a 2TT M and an input type I , we define an *alternating tree walking automaton* M' , which ignores the output of the 2TT, but apart from that imitates the behavior of M on trees in I . For M' , we then construct a nondeterministic BTA $A_{M'}$ accepting all trees u such that $\tau_M(u) \neq \emptyset$. The right-hand sides of transitions of an alternating tree walking automaton are conjunctions. Whereas the empty conjunction, i.e., $\bigwedge \emptyset$, equals true.

Definition 11.1 (ATWA). An *alternating tree walking automaton* (ATWA for short) is a tuple $M = (Q, \Sigma, \delta_M, Q_0)$ where

- Q is a finite set of states,
- Σ a ranked input alphabet,
- $Q_0 \subseteq Q$ a set of initial states, and
- δ_M a finite set of rules of the form

$$q(\mathbf{a}, \eta) \rightarrow q_1(op_1) \wedge \dots \wedge q_c(op_c)$$

with $c \geq 0$, $q, q_1, \dots, q_c \in Q$, $\mathbf{a} \in \Sigma^{(m)}$, $m \geq 0$, $\eta \geq 0$, and for $1 \leq i \leq c$, $op_i \in \{stay, up\} \cup \{down_j \mid j \in [m]\}$.

A transition $q(\mathbf{a}, \eta) \rightarrow H$ is also called q -rule. An ATWA traverses a tree like a 2TT, but produces no output. The language $\mathcal{L}(M)$ of an ATWA M is defined as the set of all trees for which an accepting run of M exists.

Definition 11.2 (Run). For an ATWA $M = (Q, \Sigma, \delta_M, Q_0)$, assume that the rule $r : q(\mathbf{a}, \eta) \rightarrow H$ is in δ_M with $H = q_1(op_1) \wedge \dots \wedge q_c(op_c)$. Then the tree

$$\rho = r(\rho_1, \dots, \rho_c) \in \mathcal{T}_{\delta_M}$$

is a (q, ϑ) -run of the ATWA M on the tree u if ϑ is a node of u with direction η and label \mathbf{a} and for all i , ρ_i is a (q_i, ϑ_i) -run of M on u where ϑ_i is obtained from ϑ by operation op_i . A (q_0, ϵ) -run ρ with $q_0 \in Q_0$ is also called *accepting*.

A subtree $\rho[w]$ of a run ρ on a node w , which is a (q, ϑ) -run for some state q and some node ϑ of u is called (q, ϑ) -*subrun*. The set of rules, which are applied to one node ϑ during a run ρ on u , is the set

$$\text{rules}_\rho(\vartheta) = \{r \mid \exists \rho_1, \dots, \rho_c : r(\rho_1, \dots, \rho_c) \text{ is a } (q, \vartheta)\text{-subrun of } \rho\}.$$

Lemma 11.2. *Assume that M is a 2TT. Then an ATWA M' can be constructed in linear time such that for every input tree u , M' has an accepting run for u iff M has an accepting run for u .*

Proof. The ATWA M' has the same set of states as M . The rules of M' are obtained from those of M by replacing every right-hand side ζ of M with the conjunction of all $q(op)$ occurring in ζ . Formally, let $M = (Q_M, \Sigma, R, Q_0)$ be a 2TT. Then, the ATWA M' is defined by $M' = (Q_M, \Sigma, \delta', Q_0)$ for a set δ' of transitions $\delta' = \{[r] \mid r \in R\}$ where

$$\begin{aligned} [q(\mathbf{a}, \eta) \rightarrow \zeta] &= q(\mathbf{a}, \eta) \rightarrow [\zeta] \\ [\zeta] &= q_1(op_1) \wedge \dots \wedge q_c(op_c) \\ &\text{for a right-hand side } \zeta = z[q_1(op_1), \dots, q_c(op_c)]. \end{aligned}$$

In order to prove the correctness of this construction, we first extend the translation $[\cdot]$ from rules to trees of rules. For $\rho = r(\rho_1, \dots, \rho_c) \in \mathcal{T}_R$, $[\rho]$ is inductively defined as the tree $[\rho] = [r]([\rho_1], \dots, [\rho_c])$. Then we claim that ρ is a (q, ϑ) -run of M iff $[\rho]$ is a (q, ϑ) -run of M' . The proof is by structural induction on ρ . Since for every (q, ϑ) -run ρ' of M' , some $\rho \in \mathcal{T}_R$ exists with $[\rho] = \rho'$, we conclude that $\tau_M(u) \neq \emptyset$ iff $u \in \mathcal{L}(M')$. \square

As an example for this translation of tree walking transducers into ATWAs, consider the 2TT M_{staff} (Example 9.1). For the second rule, we get:

$$r'_2 : q(\text{employee}, \eta) \rightarrow q_{\text{data}}(\text{down}_1) \wedge q_{\text{boss}}(\text{stay}) \wedge q_{\text{sub}}(\text{down}_1).$$

In order to accept only trees of an input type I , we enlarge an ATWA M . Let I be given by a finite state tree automaton $A = (P, \Sigma, \delta_A, F)$ and ATWA $M = (Q, \Sigma, \delta_M, Q_0)$. For every rule $q_0(\mathbf{a}, \eta) \rightarrow H$ of M with $q_0 \in Q_0$ and every $p_f \in F$, we enhance the ATWA with the rule $q_0(\mathbf{a}, \eta) \rightarrow p_f(\text{stay}) \wedge H$. Additionally, for every transition $(p, \mathbf{a}, p_1 \dots p_m) \in \delta_A$ of the automaton A , we add the rule $p(\mathbf{a}, \eta) \rightarrow \bigwedge_{i \leq m} p_i(\text{down}_i)$.

In an accepting run of an ATWA, subruns for the same state and node are interchangeable. Thus, we define *uniform* runs as runs where for each state q and each node ϑ , the (q, ϑ) -subruns are the same.

Definition 11.3 (Uniform Run). A (q, ϑ) -run ρ of an ATWA M on a tree u is called *uniform* if, for every state q , every node ϑ in u , and every two (q, ϑ) -subruns ρ_1, ρ_2 of ρ , $\rho_1 = \rho_2$.

If an ATWA is deterministic, i.e., for each state q , direction η and label \mathbf{a} , there is at most one rule of the form $q(\mathbf{a}, \eta) \rightarrow H$, then every (q, ϑ) -run is uniform. In general, this may not be the case. However, we have:

Lemma 11.3. *For an ATWA M and an input tree u , the following two statements are equivalent.*

1. M has an accepting run for u .
2. M has a uniform accepting run for u .

Proof. Every uniform accepting run is an accepting run. For the reverse direction, it suffices to prove that for every (q, ϑ) -run ρ of M on u , there is also a uniform (q, ϑ) -run of M on u . For that, we consider the set $B(\rho)$ of all pairs (q', ϑ') for which ρ contains more than one (q', ϑ') -run as a subtree. We proceed by induction on the cardinality of the set $B(\rho)$. If the set $B(\rho)$ is empty, ρ is already uniform. Now assume $B(\rho)$ is non-empty. Then ρ contains a subtree ρ_1 with the following two properties:

1. ρ_1 is a (q_1, ϑ_1) -run with $(q_1, \vartheta_1) \in B(\rho)$;
2. all subtrees ρ'_1 of ρ_1 are (q', ϑ') -runs with $(q', \vartheta') \notin B(\rho)$.

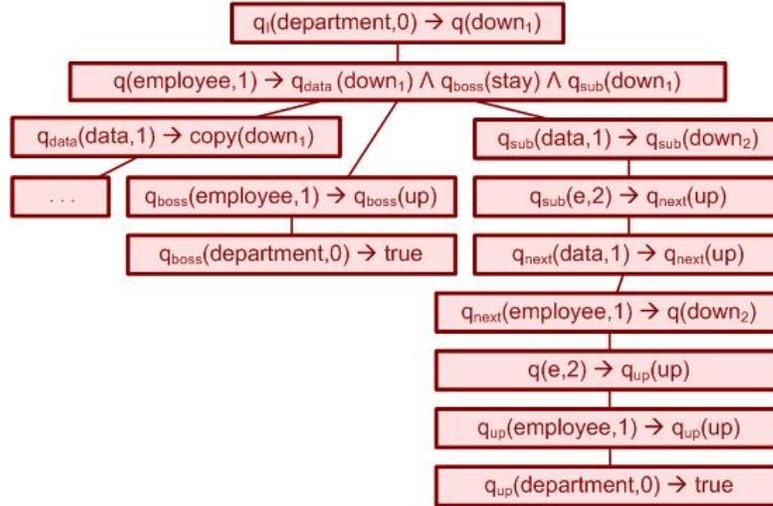
Then we construct from the run ρ a tree ρ' by replacing every occurrence of a subtree in ρ , which is a (q_1, ϑ_1) -run, with ρ_1 . Then ρ' is again a (q, ϑ) -run on u but now, the set $B(\rho') \subseteq B(\rho) \setminus \{(q_1, \vartheta_1)\}$ contains at least one element less. Thus, by induction hypothesis applied to ρ' , there is a (q, ϑ) -run on u , which is uniform. \square

Figure 9.2(a) shows the behavior of the 2TT M_{staff} on the tree u_e . The corresponding ATWA M'_{staff} yields the same behavior. The Figure 9.2(b) shows an accepting run ρ of M_{staff} on u_e . The corresponding run $[\rho]$ of M'_{staff} (Figure 11.1(a)) is uniform. In order to decide emptiness of an ATWA M and, accordingly, of a 2TT, we construct a *nondeterministic* bottom-up finite state tree automaton A_M . In order to accept the domain of M , the BTA A_M guesses uniform accepting runs. Since A_M visits each node in the input tree at most once, it *guesses* at every node all transitions, which are applied at this node during a uniform run of M .

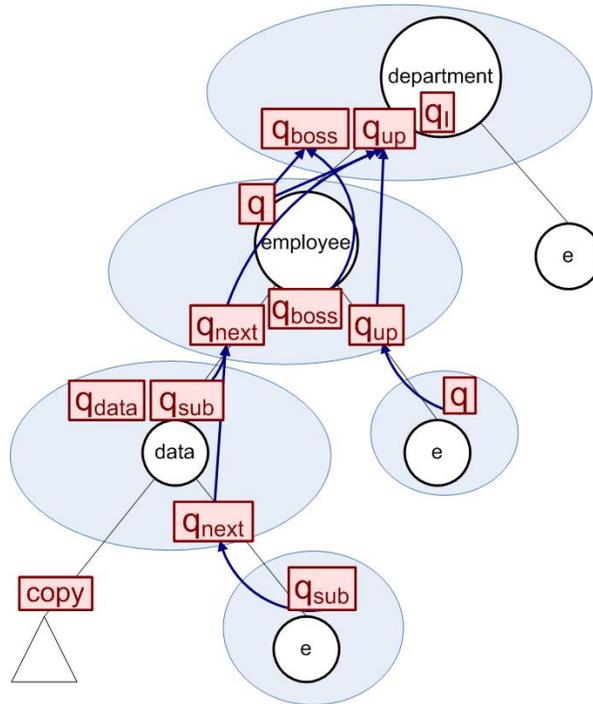
Technically, the states of the BTA A_M consist of guessed directions together with *partial* mappings $\mu : Q \dashrightarrow 2^Q$ of states to sets of states: $\mu(q) = B$ at a given node ϑ means that the (q, ϑ) -run on the input tree will cause calls $q'(up)$ at ϑ only for states q' from B . For a mapping μ , we refer to the domain as $\text{dom}(\mu) = \{q \mid \mu(q) \text{ is defined}\}$. Furthermore, to each node ϑ in the input tree we implicitly attach $\mathcal{R}_\vartheta = \text{rules}_\rho(\vartheta)$ collecting the ATWA rules which are applied to the node ϑ in an accepting run ρ of the ATWA on the input tree u . Since ρ is uniform, the set $\text{rules}_\rho(\vartheta)$ contains at most one q -rule for every q .

Example 11.2. Consider again the run ρ in Figure 11.1(a) on the tree u_e . To simulate this run on node 1.2 (with label **e** and direction 2), a run of the BTA A_M assigns to 1.2 the partial mapping $\mu_{1.2} = \{q \mapsto \{q_{up}\}\}$. In addition, we attach $\mathcal{R}_{1.2} = \{q(\mathbf{e}, 2) \rightarrow q_{up}(up)\}$ to this node. At node 1.1 with label **data** and direction 1, we obtain the partial mapping

$$\mu_{1.1} = \{q_{data} \mapsto \emptyset, q_{next} \mapsto \{q_{next}\}, q_{sub} \mapsto \{q_{next}\}\}$$



(a) A uniform run of the ATWA M'_{staff} on the tree u_e .



(b) The mappings μ for that run.

Figure 11.1: Behavior of M'_{staff} on u_e .

because q_{next} at node 1 also depends on q_{sub} at node 1.1. For the run ρ we get

$$\begin{aligned} \mathcal{R}_{1.1} = \{ & q_{data}(\mathbf{data}, 1) \rightarrow copy(down_1), \\ & q_{sub}(\mathbf{data}, 1) \rightarrow q_{sub}(down_2), \\ & q_{next}(\mathbf{data}, 1) \rightarrow q_{next}(up)\}. \end{aligned}$$

The mappings for the run ρ are illustrated in Figure 11.1(b). \triangleleft

A pair $\langle \mu, \eta \rangle$ is accepting if $\eta = 0$, for all $q \in \text{dom}(\mu)$ holds $\mu(q) = \emptyset$, and $\text{dom}(\mu)$ contains an accepting state $q_0 \in Q_0$ of the ATWA M .

Now assume that \mathbf{a} is a label of arity m , η is a direction, and $\mu, \mu_i : Q \dashrightarrow 2^Q$ are partial mappings ($i = 1, \dots, m$). Then

$$\langle \langle \mu, \eta \rangle, \mathbf{a}, \langle \mu_1, 1 \rangle, \dots, \langle \mu_m, m \rangle \rangle$$

is a transition of the BTA A_M iff there is a set \mathcal{R} of rules of the ATWA M with the following properties. Let Q_s, Q_u , and $Q_{d,i}$ ($i \in [m]$) denote the set of states q for which there is a q -rule in \mathcal{R} , the set of states q' with a recursive call $q'(up)$, and the sets of states q' with a recursive call $q'(down_i)$ in some right-hand side of rules in \mathcal{R} , respectively. Then the set \mathcal{R} of rules should have the following properties:

1. All rules in \mathcal{R} have a left-hand side of the form $q(\mathbf{a}, \eta)$ where $q \in Q$.
2. Assume $q(\mathbf{a}, \eta) \rightarrow q_1(op_1) \wedge \dots \wedge q_c(op_c) \in \mathcal{R}$. Then we have for all $j \leq c$:
 - If $op_j = stay$, then \mathcal{R} also contains a q_j -rule, i.e., $q_j \in Q_s$;
 - If $op_j = down_i$, then $q_j \in \text{dom}(\mu_i)$.
3. Whenever $q' \in \text{dom}(\mu_i)$ for i , then $\mu_i(q') \subseteq Q_s$.
4. Consider the following graph G with set of vertices

$$V = \{q(stay) \mid q \in Q_s\} \cup \{q(down_i) \mid i \in [m], q \in Q_{d,i}\} \cup \{q(up) \mid q \in Q_u\}$$

and the following set E of edges:

- If a q -rule in \mathcal{R} contains a call $q'(op)$, then $(q(stay), q'(op)) \in E$;
- If $\mu_i(q) = B_i$ is defined, then $(q(down_i), q'(stay)) \in E$ for all $q' \in B_i$.

The resulting directed graph $G = (V, E)$ should be acyclic, and the mapping μ is obtained from G as follows:

- $q \in \text{dom}(\mu)$ iff $q \in Q_s$ and
- $\mu(q) = B$ iff the set B equals the set of all vertices $q'(up)$ which are reachable in G from $q(stay)$.

The size of A_M is exponential in the size of M . We give a detailed example for this construction in Example 11.3. Now we state the correlation of runs of A_M and of M .

Lemma 11.4. *For a tree u the following statements are equivalent.*

1. *There is a uniform accepting run of M on u .*
2. *There is an accepting run of A_M on u .*

Proof. (1) \Rightarrow (2) : Let ρ be a uniform accepting run of the ATWA M on a tree u . For a node ϑ of u with label \mathbf{a} and direction η , let \mathcal{R}_ϑ denote the set of all ATWA rules applied at the root in subruns of ρ starting at ϑ , i.e., $\mathcal{R}_\vartheta = \text{rules}_\rho(\vartheta)$. We then construct for every node ϑ of u with label $\mathbf{a} \in \Sigma^{(m)}$ and direction η a state μ_ϑ and a transition $r_\vartheta = (\langle \mu_\vartheta, \eta \rangle, \mathbf{a}, \langle \mu_{\vartheta_1}, 1 \rangle \dots \langle \mu_{\vartheta_m}, m \rangle)$ of the BTA A_M .

The sets \mathcal{R}_ϑ allow us to construct a directed graph G_u . The set V_u of vertices of G_u are given by the set of all pairs (q, ϑ) for nodes ϑ of u and states q for which there is a q -rule in \mathcal{R}_ϑ . The set E_u of edges consists of:

- all pairs $((q, \vartheta), (q', \vartheta))$ where the q -rule in \mathcal{R}_ϑ contains a call $q'(\text{stay})$;
- all pairs $((q, \vartheta), (q', \vartheta_i))$ where the q -rule in \mathcal{R}_ϑ contains a call $q'(\text{down}_i)$;
- all pairs $((q, \vartheta), (q', \vartheta'))$ where the q -rule in \mathcal{R}_ϑ contains a call $q'(\text{up})$ and $\vartheta = \vartheta' i$ for some i .

The graph G_u is acyclic. Moreover, since the uniform run ρ is accepting, every vertex in V_u is reachable from some vertex (q_0, ϵ) with $q_0 \in Q_0$.

The graph G_u allows to construct partial mappings μ_ϑ for every node ϑ . The state q is in $\text{dom}(\mu_\vartheta)$ iff (q, ϑ) is a vertex of G_u . Assume (q, ϑ) is a vertex in the graph G_u . We consider two cases. If $\vartheta = \epsilon$, then \mathcal{R}_ϵ cannot contain any q -rule, which has an up -call. In this case, $\eta(\vartheta) = 0$, and we set $\mu_\epsilon(q) = \emptyset$. Otherwise, assume that $\vartheta = \vartheta' i$. Then $\eta(\vartheta) = i$ and $q' \in \mu_{\vartheta'}(q)$ iff there is an edge $((q_1, \vartheta'), (q', \vartheta'))$ in G_u where (q_1, ϑ') is reachable from (q, ϑ) by a path, which contains only vertices (q_2, ϑ_2) referring to nodes ϑ_2 from the subtree at node ϑ , i.e., to nodes, which have ϑ as a prefix. If no such state q' exists, then $\mu_{\vartheta'}(q) = \emptyset$.

It now can be verified for every node ϑ with label $\mathbf{a} \in \Sigma^{(m)}$ and direction η that $(\langle \mu_\vartheta, \eta \rangle, \mathbf{a}, \langle \mu_{\vartheta_1}, 1 \rangle \dots \langle \mu_{\vartheta_m}, m \rangle)$ constitutes a transition of A_M (with \mathcal{R}_ϑ as set of rules of the ATWA). Since by construction, $\langle \mu_\epsilon, 0 \rangle$ is an accepting state of A_M , we have, thus, constructed an accepting run of A_M for u .

(2) \Rightarrow (1) : Let ρ' be an accepting run of A_M on the tree u , and let $\langle \mu_\vartheta, \eta \rangle$ and r_ϑ denote the state and transition of A_M attained for the node ϑ in u . We can find sets \mathcal{R}_ϑ conforming to the properties of the transition relation of M . These allow to construct a graph G'_u analogously to the graph G_u above. By definition of the transition relation of A_M , G'_u is acyclic. This allows us to define for every vertex (q, ϑ) in G'_u , the number $h(q, \vartheta)$ as the maximal length of a path in G'_u to a leaf, i.e., a vertex with out-degree 0. Using the sets \mathcal{R}_ϑ of ATWA rules, we now construct for every node ϑ and ATWA rule $r : q(\mathbf{a}, \eta) \rightarrow H$ from \mathcal{R}_ϑ with $H = q_1(\text{op}_1) \wedge \dots \wedge q_c(\text{op}_c)$, a tree $\rho[q, \vartheta]$ by

$$\rho[q, \vartheta] = r(\rho[q_1, \vartheta_1], \dots, \rho[q_c, \vartheta_c])$$

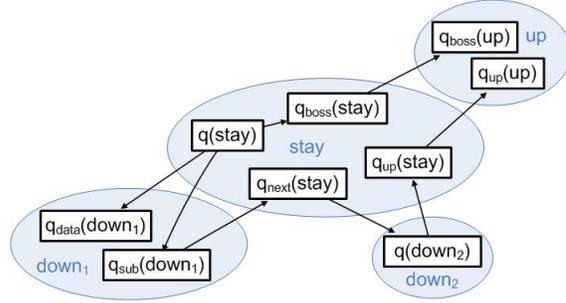


Figure 11.2: Graph G for transition $(\langle \mu_1, 1 \rangle, \text{employee}, \langle \mu_{1.1}, 1 \rangle, \langle \mu_{1.2}, 2 \rangle)$.

where $\vartheta_j = \llbracket op_j \rrbracket(\vartheta)$. Note that all these trees are well-defined, since the height of $\rho[q, \vartheta]$ precisely equals $h(q, \vartheta)$. Moreover, the tree $\rho[q, \vartheta]$ is a (q, ϑ) -run of the ATWA M on u . Since every (q', ϑ') -subrun of this tree equals the (q', ϑ') -run $\rho[q', \vartheta']$, this run is also uniform. In particular, the tree $\rho[q_0, \epsilon]$ constitutes a uniform accepting run of the ATWA M . \square

By Lemmas 11.2, 11.3, and 11.4, the BTA A_M recognizes the domain of the given 2TT M , which gives us Theorem 11.5. Note that this implies, by Theorem 10.2, that also emptiness of M 's domain (and hence of M 's translation τ_M) can be decided in exponential time.

Theorem 11.5. *Assume M is a 2TT. Then a BTA A can be constructed in exponential time such that $\mathcal{L}(A) = \text{dom}(M)$. Thus, emptiness for a 2TT can be decided in deterministic exponential time.* \square

We continue with our example to illustrate the construction of the bottom-up tree automaton:

Example 11.3. We consider again the tree walking transducer M_{staff} and its corresponding ATWA M'_{staff} . The size of the BTA $A_{M'_{\text{staff}}}$ is exponential in the size of the ATWA or 2TT. Therefore, we only construct states occurring in a run of $A_{M'_{\text{staff}}}$ on the tree in Figure 9.2(a):

$$u_e = \text{department}(\text{employee}(\text{data}(\dots, e), e), e)$$

The uniform accepting run ρ in Figure 11.1(a) yields the sets \mathcal{R}_ϑ . For instance, for the node labeled with **employee**, we get

$$\begin{aligned} \mathcal{R}_1 = \{ & q(\text{employee}, 1) \rightarrow q_{\text{boss}}(\text{stay}) \wedge q_{\text{sub}}(\text{down}_1) \wedge q_{\text{data}}(\text{down}_1), \\ & q_{\text{boss}}(\text{employee}, 1) \rightarrow q_{\text{boss}}(\text{up}), \\ & q_{\text{up}}(\text{employee}, 1) \rightarrow q_{\text{up}}(\text{up}), \\ & q_{\text{next}}(\text{employee}, 1) \rightarrow q(\text{down}_2) \}. \end{aligned}$$

The graph G_{u_e} in the proof of Lemma 11.4 is similar to the graph in Figure 9.2(a). There, it spans the tree u_e . To obtain the graph G_{u_e} , we have to replace a vertex q , which is located at a node ϑ of u_e , by (q, ϑ) .

The mappings μ_ϑ are illustrated in Figure 11.1(b). For each state q fixed at a node ϑ , $\mu_\vartheta(q)$ is defined. If q has no outgoing edges then $\mu_\vartheta(q) = \emptyset$. Otherwise, it is the set of all direct successors of q in this Figure 11.1(b). For example:

$$\begin{aligned} \mu_\epsilon &= \{q_I \mapsto \emptyset, q_{up} \mapsto \emptyset, q_{boss} \mapsto \emptyset\} \\ \mu_1 &= \{q \mapsto \{q_{up}, q_{boss}\}, q_{next} \mapsto \{q_{up}\}, q_{boss} \mapsto \{q_{boss}\}, q_{up} \mapsto \{q_{up}\}\} \\ \mu_{1.1} &= \{q_{data} \mapsto \emptyset, q_{sub} \mapsto \{q_{next}\}, q_{next} \mapsto \{q_{next}\}\} \\ \mu_{1.2} &= \{q \mapsto \{q_{up}\}\} \end{aligned}$$

Note that $q_{up} \notin \mu_{1.1}(q_{next})$ although $(q_{up}, 1)$ is reachable from $(q_{next}, 1.1)$ in G_{u_e} , but the path contains $(q, 1.2)$ and 1.1 is not a prefix of 1.2. In order to illustrate a transition of BTA $A_{M'_{staff}}$, consider, e.g., the transition

$$\langle \mu_1, 1 \rangle, \mathbf{employee}, \langle \mu_{1.1}, 1 \rangle, \langle \mu_{1.2}, 2 \rangle \in \delta_A$$

with the set \mathcal{R}_1 . All rules in \mathcal{R}_1 agree in the input label **employee** and the direction 1 (Condition 1, Page 161). Also it contains a q_{boss} -rule for the call $q_{boss}(stay)$. For all states of occurring $down_1$ -calls, i.e., states q_{sub} and q_{data} , the mapping $\mu_{1.1}$ is defined. Likewise, for q with a $down_2$ -call, $q \in \text{dom}(\mu_{1.2})$. Thus, \mathcal{R}_1 also conforms with Condition 2 (Page 161). For Condition 3, we verify that \mathcal{R}_1 has rules both for q_{next} and q_{up} . For the last property, we construct the graph G . The set of vertices is

$$\begin{aligned} V = \{ & q(stay), q_{boss}(stay), q_{next}(stay), q_{up}(stay), \\ & q_{data}(down_1), q_{sub}(down_1), q(down_2), \\ & q_{boss}(up), q_{up}(up) \} \end{aligned}$$

The edges are illustrated in Figure 11.2. As the last condition (Page 161) requires, the graph G is acyclic, and we can read off the mapping μ_1 .

We also verify that $\langle \mu_\epsilon, 0 \rangle$ is an accepting state. According to Lemma 11.4, the resulting run ρ' is an accepting run of $A_{M'_{staff}}$ on u_e . \triangleleft

11.3 Efficient Subcases

In the previous section, we have provided an algorithm for deciding emptiness of ATWAs and thus, also of 2TTs, which runs in exponential time. This algorithm is indeed worst-case optimal. Notwithstanding that, this algorithm allows us to identify subclasses of transducers where emptiness can be decided in polynomial time.

We call a tree walking transducer M *b-bounded* if every accepting run ρ of M has at most b subruns starting at node ϑ . We call M *strictly b-bounded* if every accepting run ρ of M visits each node ϑ in the input tree at most b times, i.e.,

has at most b subrun *occurrences* starting at node ϑ . The same definitions are also employed for alternating tree walking automata.

Note that the definition of b -boundedness does not exclude that the same node ϑ is traversed arbitrarily often: If so, however, these traversals will be copies of at most b distinct traversals. Note further that for a given transducer M , it is decidable whether or not there exists a b such that M is b -bounded; the same holds for strict b -boundedness. To see this, we add for each input symbol a new marked symbol of the same rank. We then consider input trees in which exactly one node is labeled by a marked symbol (this is a regular input tree language). Finally, we change the transducer M in such a way that it produces a specific output tree for each subrun that starts at the marked input node (resp. for each time the marked node is visited), and other than that does not produce any output. The output of the new transducer, when applied to input trees with exactly one node marked, is finite if and only if the transducer is b -bounded for some b (resp. strictly b -bounded for some b). The finiteness is decidable for a very large class of tree transformations [DE98], which contains the pebble tree transducers (and hence also the 2TTs) by [EM03b].

Example 11.4. Our current example 2TT M_{staff} is b -bounded — but not strictly b -bounded. Let ρ be an arbitrary run of M_{staff} on a tree u . Because the transducer is deterministic, ρ is uniform. Hence, for every node ϑ in u , there is at most one (q, ϑ) -subrun for every state q . For a node ϑ with label **employee**, there are 4 different states q' with rules of the form $q'(\text{employee}, \eta) \rightarrow \zeta$ in M_{staff} . For every other label, there are 3 or 2 different states. Thus, M_{staff} is 4-bounded.

Note the difference between subruns and occurrences of subruns. There may be different occurrences of the same subrun. Consider for example a tree u with a node ϑ , which represents the boss of n subordinates. Then, the transducer searches the boss for every subordinate again. Hence, for an accepting run of M_{staff} on u we have n occurrences of $(q_{\text{boss}}, \vartheta)$ -subruns. Thus, the 2TT M_{staff} is *not* strictly b -bounded for every b . \triangleleft

Consider the construction from the last section of a BTA A_M which accepts the same language as an ATWA M . If the alternating tree walking automaton M is b -bounded, it suffices to consider sets \mathcal{R} of ATWA rules of size b . Also, this means that partial mappings μ need to be taken into account, which are of the form: $B \dashrightarrow 2^{B'}$ for subsets B and B' of states of cardinalities at most b . Note that the number of subsets of size at most b of a set with n elements is bounded by $\frac{1}{b!}(n+1)^b$. Thus, the number of the partial mappings μ can be bounded by:

$$(n+1)^{2b} \frac{2^{b^2}}{b!}$$

if n is the number of states of M . We will not provide an explicit estimation of the number of transitions of the BTA since it crucially also depends on other parameters such as the number of rules of M that agree in input symbol and direction (which is typically small). We just note that for b -bounded M , the size of the BTA A_M is polynomial in the size of M . The occurring exponent, though, is bounded by $\mathcal{O}(kb^2)$ where k is the maximal rank of an input symbol.

Summarizing, we find that emptiness for b -bounded 2TTs can be done in polynomial time. Note, however, that neither b -boundedness or strict b -boundedness is preserved by our construction to reduce the number of state calls in right-hand sides. For an efficient method for type checking, we also require that the bound on the number of visits to every node in the input tree is preserved under the intersection construction (Section 11.1). In this respect, we observe:

Lemma 11.6. *If M is strictly b -bounded and A is a BTA, then M_A is also strictly b -bounded. If M is just b -bounded, this need not be the case.*

We thus obtain a polynomial-time algorithm for the class of strictly b -bounded 2TTs where the number of occurrences of state calls in right-hand sides is also bounded.

11.4 Conclusion

In this chapter, we presented techniques to type check tree walking transducers against regular tree languages. Our approach is forward type inference. For that purpose, for a given 2TT M , we build a second tree walking transducer, which produces only output trees in the complement of the output type, and otherwise realizes the same transformation as the original transducer M . For a BTA A describing the complement of the output type, the size of the new 2TT M_A is in $\mathcal{O}(|M| \cdot |A|^{d+1})$ where d is the maximal number of occurrences of states in right-hand sides in M ; for binary trees and with Lemma 9.1, it is in $\mathcal{O}(|M| \cdot |A|^3)$.

For this intersection 2TT M_A , we build an alternating tree walking automaton M' , which imitates the behavior of M_A , but does not produce any output. This construction can be done in linear time. Then, in order to decide emptiness of the ATWA M' and, accordingly, of the transducer M_A , we construct a nondeterministic bottom-up finite state tree automaton $A_{M'}$. In general, this construction is exponential in the size of M' . Hence, emptiness of a 2TT can be decided in exponential time — a result, which has already been known for a long time, see the notes at the end of Section 11.5. The general approach, however, allowed us to identify more efficient subclasses that have been discussed in Section 11.3. If M' is b -bounded, emptiness can be decided in polynomial time where the exponent of the polynomial only depends on b^2 if the transducer is two-way, i.e., uses *up*-operations. A closer inspection of the construction of a bottom-up tree automaton from an ATWA, though, reveals that the exponent can be reduced to b if the transducer is stay top-down, i.e., uses no *up*-operations (but possibly *stay*-operations). The construction for the intersection, on the other hand, is polynomial in the sizes both of the 2TT and the BTA — but may be exponential in the number of occurrences of state calls in right-hand sides. Also, if we start with a b -bounded 2TT M , the construction may not preserve b -boundedness. Instead, the bound on the number of visits to an input node may be increased as much as by a factor of the number of

states of the BTA. If the tree walking transducer M is *strictly* b -bounded, this property will be retained and also the alternating tree walking automaton M' is strictly b -bounded.

As a last step of verifying whether a tree walking transducer M type checks w.r.t. input and output types $\mathcal{L}(A_I)$ and $\mathcal{L}(A)$ (for a BTA A_I), we construct a further bottom-up tree automaton C with $\mathcal{L}(C) = L(A_{M'}) \cap \mathcal{L}(A_I)$ using the obvious product construction (see, e.g., Section 1.3 in [CDG⁺07]) such that

$$|C| \in \mathcal{O}(|A_{M'}| \cdot |A_I|) .$$

According to Theorem 10.2, we can test whether $\mathcal{L}(C) = \emptyset$ (which means that tree walking transducer M type checks w.r.t. $\mathcal{L}(A_I)$ and $\mathcal{L}(A)$) in time linear in $|C|$.

Theorem 11.7. *Deciding whether a strictly b -bounded 2TT M type checks w.r.t. regular tree languages I and O , given by BTAs A_I and A_O , is polynomial in the size of M , A_I , and A_O , but exponential in $b^2 \cdot (d + 1)$ where d is the maximal number of occurrences of state calls in right-hand sides. If M has no up-operations, the exponent can be improved to $b \cdot (d + 1)$.*

11.5 Notes and References

Our definition of ATWAs is equivalent to the alternating two-way finite tree automata of [Slu85]. Note that alternating tree automata have recently been used in the context of a practical implementation of type checking for tree transducers [FH07].

The intersection of a 2TT with a given output type (Theorem 11.1) can be seen as a sequential composition of the 2TT with a translation in FTA; the latter is the class of partial identity mappings for regular tree languages. With this in mind, we can, for instance, obtain that top-down tree transducers allow a similar result as the one in Theorem 11.1: By Corollary 2(1) of [Bak79], top-down tree transducers are closed under composition with linear and nondeleting top-down tree transducers; since FTA is included in the latter class, we obtain the desired result for top-down tree transducers. It is an interesting open problem whether a similar composition result holds for 2TTs, i.e., whether 2TTs are closed under composition with linear and nondeleting 2TTs.

The notion of b -boundedness is similar to the notion of finite-copying in tree transducers, see, e.g., [ERS80, EM99]. Similar to the results of [EM03c], it probably holds that b -bounded transformations are of linear size increase. A more static version of b -boundedness is the single-use restriction known for attribute grammars [Gie88]. According to [EM99], it can probably be shown that total, deterministic, strictly b -bounded 2TTs are equivalent to single-use restricted attribute grammars. In Section 5 of [MPS07], a similar result as Theorem 11.7 has been shown for stay-macro tree transducers (cf. also discussion in Section 12.4)

Engelfriet et al. show in [EHS07] (Theorem 5) that for every 2TT M (TT in [EHS07]), a regular tree grammar G can be constructed in exponential time

such that G generates the domain of τ_M . They refer to the relationship between 2TTs and attributed tree transducers explained in [EM03b] and a result of [Bar82] — giving the Theorem 11.5 above. The result of Theorem 11.5 has also been stated in Theorem A.2 of [CGKV88] with a proof sketch that uses a game theoretic interpretation of acceptance due to Muller and Schupp [MS87]. The result also appears as Theorem 1 in [Eng09], where a finite state automaton for the domain of a tree-walking tree transducer (twtt) is constructed in exponential time; Theorem 2 of that paper states that inverse type inference is in k -fold exponential time, for k -fold compositions of twtts.

Chapter 12

Macro Tree Walking Transducers

In our current example, we have considered the 2TT M_{staff} which lists the staff members of a department. Although in general, several employees may have the same boss, the transducer spawns for every employee a separate computation to determine the corresponding boss. Conceptually as well as technically, it would be more convenient to determine the boss first, store it in some accumulating parameter and then propagate it to each of his employees. For this reason, we enhance tree transducers with accumulating parameters. A tree transducer with accumulating parameters is also called *macro tree transducer*.

Example 12.1. We omit the state q_{boss} and store the data of the boss in the first parameter y_b . The transformation of the next employee, which is not a subordinate of the current, is then stored in the second parameter (y_n). By this construction, we completely omit the states q_{sub} , q_{next} , and q_{up} . The transducer consists of the following rules, for all $\eta \in \{1, 2\}$:

$$\begin{array}{ll} 1 & q_I(\text{department}, 0) \quad \rightarrow \text{staff}(q(\text{down}_1, \mathbf{e}, \mathbf{e}), \mathbf{e}) \\ 2 & q(\text{employee}, \eta, y_b, y_n) \quad \rightarrow \text{employee}(\text{data}(q_{\text{data}}(\text{down}_1), y_b), \\ 3 & \quad \quad \quad q(\text{down}_1, \text{boss}(q_{\text{data}}(\text{down}_1), \mathbf{e}), \\ 4 & \quad \quad \quad q(\text{down}_2, y_b, y_n))) \\ 5 & q(\mathbf{e}, \eta, y_b, y_n) \quad \rightarrow y_n \\ 6 & q(\text{data}, 1, y_b, y_n) \quad \rightarrow q(\text{down}_2, y_b, y_n) \\ 7 & q(\text{subordinates}, 2, y_b, y_n) \rightarrow q(\text{down}_1, y_b, y_n) \\ 8 & q_{\text{data}}(\text{data}, 1) \quad \rightarrow \text{copy}(\text{down}_1) \end{array}$$

Here state *copy* is meant to copy the content of **data** (i.e., the left child in the binary representation). ◁

For the formal definition of macro tree walking transducer, we assume that every

state $q \in Q$ has a fixed rank, i.e., $Q = \bigsqcup_{n \in \mathbb{N}} Q^{(n)}$ where $Q^{(n)}$ is the set of all states with rank n .

Definition 12.1 (2MTT). A *macro tree walking transducer* M (2MTT for short) is a tuple (Q, Σ, R, Q_0) where Q is a set of ranked states, Σ is a ranked alphabet, $Q_0 \subseteq Q^{(1)}$ is a set of initial states, and R is a finite set of rules of the form $q(\mathbf{a}, \eta, y_1, \dots, y_n) \rightarrow \zeta$, where $q \in Q^{(n+1)}$, $\mathbf{a} \in \Sigma^{(m)}$, $n, m \geq 0$, $\eta \geq 0$ is a direction and y_1, \dots, y_n are the accumulating parameters of q . Possible right-hand sides are described by the grammar

$$\zeta ::= \mathbf{b}(\overbrace{\zeta, \dots, \zeta}^{m' \text{ times}} \mid y_j \mid \overbrace{q'(op, \zeta, \dots, \zeta)}^{n' \text{ times}}),$$

with $m', n' \geq 0$, symbol $\mathbf{b} \in \Sigma^{(m')}$, $j \in [n]$, state $q' \in Q^{(n'+1)}$, and operation $op \in \{\text{stay}, \text{up}\} \cup \{\text{down}_i \mid i \in [m]\}$.

In practice, states q may differ in their *rank*, i.e., the numbers of their accumulating parameters. Let $X = \{x_1, x_2, \dots\}$ denote a countable set of variables of rank (not necessarily 0), and assume that Σ, X and Y are disjoint. For a right-hand side ζ , we write also

$$\zeta = z[q_1(op_1), \dots, q_c(op_c)]$$

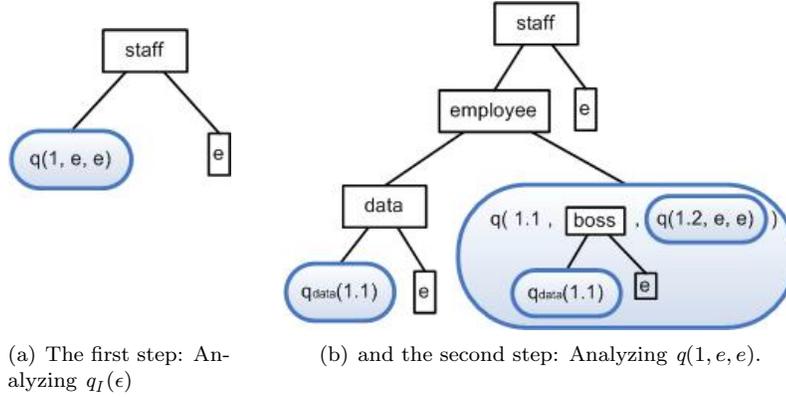
to refer to all occurrences of (maybe nested) state calls in the right-hand side. Here, $z \in \mathcal{T}_{\Sigma \cup X}(Y)$ is a tree, which contains each variable x_1, \dots, x_c exactly once with $\zeta = z[q_1(op_1)/x_1, \dots, q_c(op_c)/x_c]$ where $z[q_i(op)/x_i]$ denotes the substitution of the state call $q_i(op, z_1, \dots, z_n)$ for the subtree $x_i(z_1, \dots, z_n)$ in z where n is the rank of x_i and $n+1$ is the rank of q_i . Note that in z no state call occurs anymore. The right-hand side of the rule in Lines 2-4 in the Example 12.1 can, for example, be written as

$$z[q_{data}(\text{down}_1), q(\text{down}_1), q_{data}(\text{down}_1), q(\text{down}_2)]$$

where $z = \text{employee}(\text{data}(x_1, y_b), x_2(\text{boss}(x_3, \mathbf{e}), x_4(y_b, y_n)))$.

Intuitively, the meaning of the expressions of a right-hand side is as follows: The produced output is defined analogously to the output of a 2TT up to the accumulating parameters. Here, we consider *call-by-value* parameter passing only. Thus, the expression ζ_j in parameter position j is evaluated first; then the result (which is a tree without state calls) may be copied to the various uses of the formal parameter y_j . This evaluation strategy is also called *inside-out* (IO for short). Note that we slightly abuse Definition 12.1 and use accumulating parameters with names other than y_1, y_2, \dots (e.g., in Example 12.1 where we use y_b and y_n). Clearly, this is without loss of generality, as parameters can easily be renamed according to the definition.

Example 12.2. The rules in the beginning of this chapter with the initial state q_I form a 2MTT $M_{y, \text{staff}}$. To transform the tree u_B' (the binary representation of our common example tree, see the left tree in Figure 9.1), the macro

Figure 12.1: Applying the 2MTT $M_{y,staff}$ to u_B .

tree walking transducer starts at the root and applies the initial state. Thus, for the first step, we get $\mathbf{staff}(q(\mathit{down}_1, \mathbf{e}, \mathbf{e}), \mathbf{e})$ where down_1 refers to the node 1. Applying the state q to this $\mathbf{employee}$ -node, we get several state calls. These state calls are partially nested (Figure 12.1). In the left figure is the first output with one state call. The right figure shows the tree after processing $q(1, \mathbf{e}, \mathbf{e})$. There, we get the state call

$$q(\mathit{down}_1, \mathbf{boss}(q_{\mathit{data}}(\mathit{down}_1), \mathbf{e}), q(\mathit{down}_2, y_b, y_n))$$

with nested calls. The first parameter accumulates a tree with root \mathbf{boss} , whereas the second parameter is a further state call. \triangleleft

The order in which nested state calls are evaluated indeed matters. Consider, e.g., a transducer with the rules

$$\begin{aligned} q_I(\mathbf{a}, 0) &\rightarrow p(\mathit{stay}, q'(\mathit{stay})) \\ p(\mathbf{a}, 0, y) &\rightarrow \mathbf{a} \\ q'(\mathbf{b}, 0) &\rightarrow \mathbf{b} . \end{aligned}$$

If we evaluate the outermost calls first, the tree $u = \mathbf{a}(u_1, \dots, u_k)$ will be transformed into \mathbf{a} . In this case, the accumulating parameter of p need not to be evaluated. If we start with the innermost calls, there is no rule to evaluate the state call $q'(\mathit{stay})$ in the right-hand side of the first rule. Thus, the output is empty.

We specify the translation induced by a 2MTT using a denotational formulation. Later, we will also consider an operational semantics based on runs. In the denotational semantics, the meaning $\llbracket q \rrbracket_u$ of state q of transducer M with n accumulating parameters (with respect to an input tree u) is defined as a mapping from nodes in the input tree to sets of trees with parameters in $Y_n = \{y_1, \dots, y_n\}$, i.e.,

$$\llbracket q \rrbracket_u : \mathbf{Nodes}(u) \rightarrow 2^{\mathcal{T}_{\Sigma}(Y)} .$$

When we evaluate an innermost call $q(\vartheta, v_1, \dots, v_n)$ on a node ϑ during a computation, it suffices to substitute actual parameters v_j for the formal parameters y_j of all terms from $\llbracket q \rrbracket_u(\vartheta)$ to obtain the set of produced outputs. The values $\llbracket q \rrbracket_u$ for all q are jointly defined as the least mappings satisfying:

$$\llbracket q \rrbracket_u(\vartheta) \supseteq \llbracket \zeta \rrbracket_u \quad \text{for rule } q(\mathbf{a}, \eta, \bar{y}) \rightarrow \zeta$$

where \bar{y} denotes the sequence y_1, \dots, y_n of parameters and ϑ is a node of u with $\text{lab}(\vartheta) = \mathbf{a}$, $\eta(\vartheta) = \eta$, and $\llbracket \zeta \rrbracket_u$ is defined by:

$$\begin{aligned} \llbracket y_j \rrbracket_u &= \{y_j\} \\ \llbracket \mathbf{b}(\zeta_1, \dots, \zeta_m) \rrbracket_u &= \{\mathbf{b}(v_1, \dots, v_m) \mid v_i \in \llbracket \zeta_i \rrbracket_u\} \\ \llbracket q'(op, \zeta_1, \dots, \zeta_{n'}) \rrbracket_u &= \{z[v_1/y_1, \dots, v_{n'}/y_{n'}] \mid z \in \llbracket q' \rrbracket_u(\llbracket op \rrbracket_u(\vartheta)), v_i \in \llbracket \zeta_i \rrbracket_u\} \end{aligned}$$

Again, op stands for $down_i$, $stay$ or up . Recall that the meaning $\llbracket op \rrbracket$ is defined by

$$\llbracket stay \rrbracket_u(\vartheta) = \vartheta, \quad \llbracket down_i \rrbracket_u(\vartheta) = \vartheta i, \quad \text{and} \quad \llbracket up \rrbracket_u(\vartheta i) = \vartheta.$$

Also, $z[v_1/y_1, \dots, v_n/y_n]$ denotes the simultaneous substitution of the trees v_j for all occurrences of the variables y_j in the tree z . Note that the call-by-value semantics is reflected in the last equation: The same trees v_j are used for all occurrences of a variable y_j in the tree z corresponding to a potential evaluation of the state q' . The transformation τ_M , realized by the 2MTT M on an input tree u and sets U of input trees, respectively, is, thus, defined by

$$\tau_M(u) = \bigcup \{ \llbracket q_0 \rrbracket_u(\epsilon) \mid q_0 \in Q_0 \} \quad \text{and} \quad \tau_M(U) = \bigcup \{ \tau_M(u) \mid u \in U \}.$$

For the operational semantics, runs of a 2MTT M on a tree u may be similarly defined as for a 2TT. It is a ranked tree over the set of rules. Here, the rank of a rule $q(\mathbf{a}, \eta, y_1, \dots, y_n) \rightarrow \zeta$ is given by the number of occurrences of recursive calls $q'(op)$ in ζ to states q' in Q . These calls may be nested. Figure 12.2 shows an accepting run ρ of the (deterministic) example 2MTT $M_{y, \text{staff}}$ on the tree $u_e = \text{department}(\text{employee}(\text{data}(\dots, \mathbf{e}), \mathbf{e}), \mathbf{e})$, which describes a department with one employee.

The denotational view on the semantics of a 2MTT allows us to use fixpoint arguments for proving the correctness of constructions, whereas the operational view is better suited for more combinatorial arguments. In particular, we can show that the number of occurrences of states in right-hand sides can be restricted to the maximum of the ranks of output symbols and states. We have:

Lemma 12.1. *For every 2MTT M , there exists a 2MTT M' with*

1. $\tau_{M'} = \tau_M$
2. *the number of states occurring in each right-hand side is bounded by k*
3. $|M'| \in \mathcal{O}(|M| \cdot k^2)$

where k is the maximum of the ranks of output symbols and states of M .

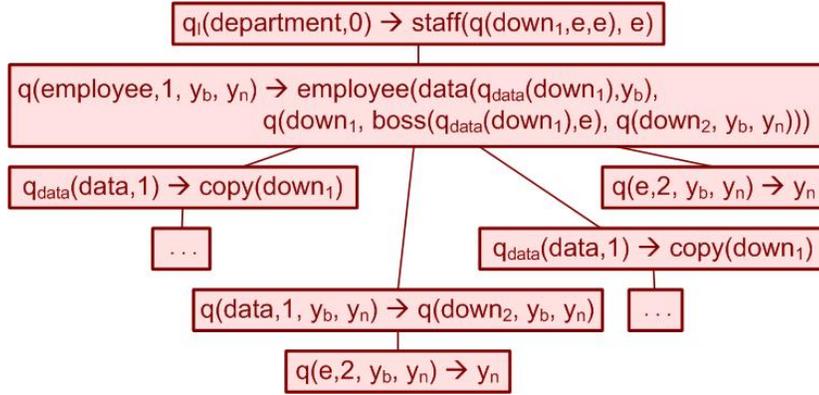


Figure 12.2: A run of the 2MTT $M_{y,staff}$ on the tree u_e .

Proof. The construction proceeds in two phases. In the first phase, we replace every *complicated* call $q'(op, \zeta_1, \dots, \zeta_n)$ in the right-hand side of a rule $q(\mathbf{a}, \eta, y_1, \dots, y_{n'}) \rightarrow \zeta$ by the simple call

$$\langle q, op, \zeta_1, \dots, \zeta_n \rangle(stay, y_1, \dots, y_{n'})$$

for a new state $\langle q, op, \zeta_1, \dots, \zeta_n \rangle$. Let $[\zeta]$ denote the resulting tree. For the new state $\langle q, op, \zeta_1, \dots, \zeta_n \rangle$, we introduce the rule

$$\begin{aligned} &\langle q, op, \zeta_1, \dots, \zeta_n \rangle(\mathbf{a}, \eta, y_1, \dots, y_{n'}) \\ &\rightarrow q(op, \langle \zeta_1, n' \rangle(stay, y_1, \dots, y_{n'}), \dots, \langle \zeta_n, n' \rangle(stay, y_1, \dots, y_{n'})) \end{aligned}$$

for again fresh states $\langle \zeta_j, n' \rangle$, which are meant to produce the output of ζ_j using n' parameters. For these states, we introduce the rules:

$$\langle \zeta_j, n' \rangle(\mathbf{a}, \eta, y_1, \dots, y_{n'}) \rightarrow [\zeta_j].$$

As a result of this first transformation phase, we achieve that all right-hand sides either are of the form $q(op, q_1(stay, y_1, \dots, y_{n'}), \dots, q_n(stay, y_1, \dots, y_{n'}))$ or contain only non-nested calls, i.e., calls of the form $q(op, y_1, \dots, y_{n'})$. In order to restrict the number of calls in right-hand sides of the second type, we essentially proceed as in the proof of Lemma 9.1, i.e., we introduce extra auxiliary states for every proper subtree of right-hand sides of the second kind, which contain more than one call.

The resulting transducer has at most one fresh state for every node of a right-hand side while the total sum of sizes of right-hand sides, may increase by a factor of k^2 in order to spell out all the auxiliary lists of parameters for the new states. \square

12.1 Type Checking 2MTTs

As for 2TTs, we now consider type checking for *macro* tree walking transducer. For a 2MTT M and a regular language \mathcal{L} , we again construct a transducer M' with $\tau_{M'}(t) = \tau_M(t) \cap \mathcal{L}$. As in Section 11.1, the language \mathcal{L} consists of all erroneous outputs. In our application scenario of type checking, the language \mathcal{L} is the complement of the output type which is either specified by a complete DBTA or by a DTTA. Beyond the case of 2TTs, we now additionally have to deal with accumulating parameters. The macro tree walking transducer M' has to keep track of the states of an automaton for \mathcal{L} on the current values of the respective parameters. We start with a general construction for deterministic bottom-up automata.

Theorem 12.2. *For every 2MTT M and every DBTA A , a 2MTT M_A can be constructed with*

$$\tau_{M_A}(u) = \tau_M(u) \cap \mathcal{L}(A)$$

for all $u \in \mathcal{T}_\Sigma$. The 2MTT M_A is of size $\mathcal{O}(|M| \cdot |A|^{l \cdot (d+1)})$ where l is the maximal rank of a state in M and d is the maximal number of occurrences of states in right-hand sides of M .

Proof. Let $M = (Q, \Sigma, R, Q_0)$ and $A = (P, \Sigma, \delta, F)$. For each state q in $Q^{(n+1)}$ and all states $p_0, \dots, p_n \in P$, the 2MTT M_A has a state $\langle q, p_0 p_1 \dots p_n \rangle$, which is meant to generate all trees z (possibly with variables from $\{y_1, \dots, y_n\}$) that could be produced by M , and for which, additionally, there is a run of A starting at the leaves y_j with state p_j , and reaching the root of z in state p_0 . The rules of M_A are:

$$\langle q, p_0 p_1 \dots p_n \rangle(\mathbf{a}, \eta, y_1, \dots, y_n) \rightarrow \zeta'$$

for every rule $q(\mathbf{a}, \eta, y_1, \dots, y_n) \rightarrow \zeta$ of M and $\zeta' \in \tau^{p_0, p_1, \dots, p_n}[\zeta]$ where the sets $\tau^{p_0, p_1, \dots, p_n}[\cdot]$ are inductively defined by:

$$\begin{aligned} \tau^{p_j, p_1, \dots, p_n}[y_j] &= \{y_j\} \\ \tau^{p_0, p_1, \dots, p_n}[\mathbf{b}(\zeta_1, \dots, \zeta_m)] &= \{\mathbf{b}(\zeta'_1, \dots, \zeta'_m) \mid \exists p'_1, \dots, p'_m \in P : \\ &\quad (p_0, \mathbf{b}, p'_1 \dots p'_m) \in \delta \wedge \forall j : \zeta'_j \in \tau^{p'_j p_1 \dots p_n}[\zeta_j]\} \\ \tau^{p_0, p_1, \dots, p_n}[q'(op, \zeta_1, \dots, \zeta_{n'})] &= \{\langle q', p_0 p'_1 \dots p'_{n'} \rangle(op, \zeta'_1, \dots, \zeta'_{n'}) \mid \\ &\quad \exists \zeta'_1, \dots, \zeta'_{n'} \in Z : \forall j : \zeta'_j \in \tau^{p'_j p_1 \dots p_n}[\zeta_j]\} \end{aligned}$$

where Z denotes the set of all subterms of possible right-hand sides of rules of M_A . The set of initial states of M_A is given by $Q'_0 = Q_0 \times F$. By fixpoint induction, we verify for every state q of rank $n + 1$, every input tree $u \in \mathcal{T}_\Sigma$, every node $\vartheta \in \text{Nodes}(u)$ and states p_0, \dots, p_n of A that

$$\llbracket \langle q, p_0, \dots, p_n \rangle \rrbracket_u(\vartheta) = \llbracket q \rrbracket_u(\vartheta) \cap \{z \in \mathcal{T}_\Sigma(Y) \mid \delta^*(z, p_1 \dots p_n) = p_0\} \quad (12.1)$$

where $Y = \{y_1, \dots, y_n\}$ and δ^* is the extension of the transition function of A to trees containing variables from Y , namely, for $\bar{p} = p_1 \dots p_n$:

$$\begin{aligned}\delta^*(y_j, \bar{p}) &= p_j \\ \delta^*(\mathbf{a}(t_1, \dots, t_m), \bar{p}) &= \delta(\mathbf{a}, \delta^*(t_1, \bar{p}) \dots \delta^*(t_m, \bar{p}))\end{aligned}$$

The correctness of the construction follows from Equation 12.1.

For each state in M , we have at most $|A|^l$ new states in M_A where l is the maximal rank of states in M . Assume that d is the maximal number of occurrences of states in right-hand sides of rules of M . Then each rule of M gives rise to at most $|A|^{l \cdot (d+1)}$ new rules in M_A . Therefore, the new 2MTT is of size $\mathcal{O}(|M| \cdot |A|^{l \cdot (d+1)})$. \square

Lemma 10.1 provides a BTA describing the complement of a DTTA — which, however, is not necessarily deterministic. Theorem 12.2, on the other hand, only holds for deterministic BTAs. A similar construction is also possible if the BTA A is nondeterministic — but then only for transducers, which are output-linear. Here, we call a 2MTT output-linear if every accumulating parameter occurs at most once in a right-hand side. Nonetheless, we are able to handle complements of output types described by DTTAs directly. For that, however, we introduce a dedicated construction of a 2MTT $M_{\bar{A}}$.

Theorem 12.3. *For every 2MTT M and every DTTA A , a 2MTT $M_{\bar{A}}$ can be constructed with*

$$\tau_{M_{\bar{A}}}(u) = \tau_M(u) \cap \overline{\mathcal{L}(A)}$$

for all $u \in \mathcal{T}_\Sigma$. The 2MTT $M_{\bar{A}}$ is of size $\mathcal{O}(|M| \cdot (h \cdot |A|)^{d+2})$ where $h + 1$ is the maximum of the maximal rank of a state in M and the maximal rank of an output symbol, and d is the maximal number of occurrences of state calls in right-hand sides of M .

Proof. Let $M = (Q, \Sigma, R, Q_0)$ and $A = (P, \Sigma, \delta, p_0)$. Our goal is to construct a 2MTT $M_{\bar{A}}$, which simulates the behavior of M , while at the same time guessing a path in the output tree, which proves non-containment in the set $\mathcal{L}(A)$. For that, the set Q' is defined as:

$$\begin{aligned}Q' &= Q \cup \{\langle q, p \rangle \mid q \in Q, p \in P\} \\ &\cup \{\langle q, p, j, p' \rangle \mid q \in Q^{(n)}, p, p' \in P, j \in [n]\}\end{aligned}$$

Here, a state $q \in Q$ of $M_{\bar{A}}$ behaves like the state q of M . States $\langle q, p \rangle$ or $\langle q, p, j, p' \rangle$ behave like q in M but, additionally, make sure that there is a path in the generated output starting from a state p of A at the root, which verifies that there is no p -run of A on the output. Thereby, a state $\langle q, p \rangle$ will directly generate the end point of such a path whereas state $\langle q, p, j, p' \rangle$ will only generate a path with p at the root reaching a parameter y_j with state p' . Accordingly, the 2MTT $M_{\bar{A}}$ has the following rules:

$$\begin{aligned}
q(\mathbf{a}, \eta, y_1, \dots, y_n) &\rightarrow \zeta \\
\langle q, p \rangle(\mathbf{a}, \eta, y_1, \dots, y_n) &\rightarrow \zeta' \text{ with } \zeta' \in [\zeta]^p \\
\langle q, p, j, p' \rangle(\mathbf{a}, \eta, y_1, \dots, y_n) &\rightarrow \zeta' \text{ with } \zeta' \in [\zeta]^{p, j, p'}
\end{aligned}$$

for every rule $q(\mathbf{a}, \eta, y_1, \dots, y_n) \rightarrow \zeta$ of M . The sets $[\cdot]^x$ are inductively defined in Figure 12.3.

First, we verify that for every tree $z \in \mathcal{T}_\Sigma(Y)$, the sets $[z]^p$ and $[z]^{p, j, p'}$ either are empty or equal $\{z\}$ where following holds:

1. $[z]^p = \{z\}$ iff z contains a node $\vartheta = i_1 \dots i_r$ such that there are transitions $(p_0^{(j)}, \mathbf{a}_j, p_1^{(j)} \dots p_{m_j}^{(j)}) \in \delta$ for every $j \in [r-1]$, such that
 - (a) the label of the node $i_1 \dots i_j$ equals \mathbf{a}_j ;
 - (b) $p_0^{(1)} = p$ and for every $j \in [r-2]$, $p_0^{(j+1)} = p_{i_j}^{(j)}$;
 - (c) there is no $p_{i_r}^{(r-1)}$ -transition of A for $\mathbf{a}_r = \text{lab}(\vartheta)$.
2. $z \in [z]^{p, j, p'}$ iff z contains a node $\vartheta = i_1 \dots i_{r+1}$, which is labeled with $y_{j'}$ and there are transitions $(p_0^{(j)}, \mathbf{a}_j, p_1^{(j)} \dots p_{m_j}^{(j)}) \in \delta$ for every $j \in [r]$, such that:
 - (a) For every $j = 1, \dots, r$, the label of the node $i_1 \dots i_j$ equals \mathbf{a}_j ;
 - (b) $p_0^{(1)} = p$ and $p_{i_{r+1}}^{(r)} = p'$ and for every $j \in [r-1]$, $p_0^{(j+1)} = p_{i_j}^{(j)}$.

Note in particular that by this definition, $z \notin \mathcal{L}(A)$ iff $[z]^{p_0} = \{z\}$ for the initial state p_0 of A . Let us extend the operators $[\cdot]^p$ and $[\cdot]^{p, j, p'}$ by:

$$[Z]^p = \bigcup \{[z]^p \mid z \in Z\} \quad [Z]^{p, j, p'} = \bigcup \{[z]^{p, j, p'} \mid z \in Z\}$$

for $Z \subseteq \mathcal{T}_\Sigma(Y)$ with $Y = \{y_1, \dots, y_n\}$. By fixpoint induction, we verify for every state q of rank $n+1$, every $j \in [n]$, every input tree $u \in \mathcal{T}_\Sigma$ and states $p, p' \in P$ that:

$$\begin{aligned}
\llbracket \langle q, p, j, p' \rangle \rrbracket(u) &= \llbracket [q] \rrbracket(u)^{p, j, p'} \\
\llbracket \langle q, p \rangle \rrbracket(u) &= \llbracket [q] \rrbracket(u)^p
\end{aligned}$$

For each state p and right-hand side ζ of a rule, we assign states of the deterministic top-down tree automaton to the nodes of the tree. Either we immediately hit a node certifying the non-existence of a p -run of the DTTA on the output generated from ζ , or we hit an occurrence of a state call $q(op, \dots)$. If $n \geq 1$ is the rank of q , we have n choices here: Either we expect a certificate for the failure of the DTTA A inside the evaluation of q or in one of the parameters of q . Overall, we find that every rule of M , thus, gives rise to $(h \cdot |A|)^{d+2}$ rules. Thus, we have $|M_{\bar{A}}| \in \mathcal{O}(|M| \cdot (h \cdot |A|)^{d+2})$. \square

Assume that we have a binary ranked alphabet and that at least one state has an accumulating parameter, i.e., the maximal rank l of states is at least 2. By

$$\begin{aligned}
[y_k]^p &= \emptyset \\
[\mathbf{b}(\zeta_1, \dots, \zeta_m)]^p &= \{\mathbf{b}(\zeta_1, \dots, \zeta_m) \mid \forall \bar{p} \in P^m : (p, \mathbf{b}, \bar{p}) \notin \delta\} \\
&\cup \{\mathbf{b}(\zeta_1, \dots, \zeta_{\nu-1}, \zeta'_\nu, \zeta_{\nu+1}, \dots, \zeta_m) \mid \nu \geq 1, (p, \mathbf{b}, p_1 \dots p_m) \in \delta, \zeta'_\nu \in [\zeta_\nu]^{p_\nu}\} \\
[q(op, \zeta_1, \dots, \zeta_n)]^p &= \{(q, p, \nu, p_\nu)(op, \zeta_1, \dots, \zeta_{\nu-1}, \zeta'_\nu, \zeta_{\nu+1}, \dots, \zeta_n) \mid \zeta'_\nu \in [\zeta_\nu]^{p_\nu}\} \\
&\cup \{(q, p)(op, \zeta_1, \dots, \zeta_n)\} \\
[y_j]^{p, i, p'} &= \{y_j \mid p = p'\} \\
[\mathbf{b}(\zeta_1, \dots, \zeta_m)]^{p, i, p'} &= \{\mathbf{b}(\zeta_1, \dots, \zeta_{\nu-1}, \zeta'_\nu, \zeta_{\nu+1}, \dots, \zeta_m) \mid \nu \geq 1, (p, \mathbf{b}, p_1 \dots p_m) \in \delta, \zeta'_\nu \in [\zeta_\nu]^{p_\nu, i, p'}\} \\
[q(op, \zeta_1, \dots, \zeta_n)]^{p, i, p'} &= \{(q, p, \nu, p_\nu)(op, \zeta_1, \dots, \zeta_{\nu-1}, \zeta'_\nu, \zeta_{\nu+1}, \dots, \zeta_n) \mid \zeta'_\nu \in [\zeta_\nu]^{p_\nu, i, p'}\}
\end{aligned}$$

Figure 12.3: The definition of the sets $[\cdot]^x$ in the proof of Theorem 12.3.

Lemma 12.1 it is then possible to restrict the number of occurrences of state calls in right-hand sides of the 2MTT to l . This implies that the size of the intersection 2MTT M_A in Theorem 12.2 for a DBTA describing the output language is in $\mathcal{O}(|M| \cdot |A|^{l \cdot (l+1)})$. Furthermore, the size of the 2MTT $M_{\bar{A}}$ in Theorem 12.3 for a deterministic top-down tree automaton describing the output language is in $\mathcal{O}(|M| \cdot (l \cdot |A|)^{l+2})$.

12.2 Deciding Emptiness of 2MTTs

To decide emptiness of a 2MTT M , we follow the approach taken for 2TTs: We construct an alternating tree walking automaton A_M , which is then tested for emptiness. The ATWA A_M has the same set of states as M (but they are not ranked anymore now) where the initial states of M and A_M coincide. For every rule $q(\mathbf{a}, \eta, y_1, \dots, y_n) \rightarrow \zeta$ of the 2MTT M , the ATWA A_M has a rule of the form

$$q(\mathbf{a}, \eta) \rightarrow q_1(op_1) \wedge \dots \wedge q_c(op_c)$$

if $q_1(op_1, \dots), \dots, q_c(op_c, \dots)$ is the sequence of calls to states of M (possibly nested inside each other), in any order. Since we use 2MTTs with call-by-value semantics, M has an accepting run on some input tree u iff A_M has an accepting run on u . Note that this construction is wrong for call-by-need semantics because M could have an accepting run on a tree $u \notin \mathcal{L}(A_M)$; for instance, the 2MTT with the rules $q_I(\mathbf{a}, 0) \rightarrow q(stay, q'(stay))$ and $q(\mathbf{a}, 0, y) \rightarrow \mathbf{a}$ on the tree \mathbf{a} .

Theorem 12.4. *For every 2MTT M , an ATWA A_M can be constructed in polynomial time such that $\mathcal{L}(A_M) = \text{dom}(M)$. Thus, it can be decided in deterministic exponential time whether the translation of a 2MTT is empty or not.*

12.3 Input-Linear 2MTTs

The notions of b -boundedness and strict b -boundedness, which we have defined for 2TTs stay meaningful also in presence of accumulating parameters, i.e., for 2MTTs. Analogously, we find that emptiness for b -bounded macro tree walking transducers is decidable in polynomial time — independent on the number of accumulating parameters of states.

In order to identify classes of 2MTTs where full type checking is tractable, we therefore take a closer look at the construction for the intersection of 2MTTs with (complements of) output types. For simplicity, we first consider macro tree walking transducers, which are *strictly 1-bounded*. This notion is only meaningful for *top-down* MTTs, i.e., 2MTTs without operations *up* or *stay*. A top-down transducer M is guaranteed to visit each node of the input tree at most once if for the same i , the operation $down_i$ does not occur twice in the same right-hand side of M . This property can easily be checked syntactically. Tree transducers satisfying this restriction are called *input-linear*.

Note that input-linearity for a tree transducer implies that the number of state calls in right-hand sides is bounded by the maximal rank of input symbols. Moreover, the output language can be described by rules that are obtained by simply deleting all directives from the transducer's rules. The resulting rules no longer specify a transformation but constitute a *context-free tree grammar* (short: CFTG) for generating output trees.

Example 12.3. As an example of an input-linear MTT consider the following macro tree transducer, which produces the same output as the common example 2MTT $M_{y, \text{staff}}$ without the **boss**-subtrees. The transducer only needs one parameter (for the next employee) and has a new state q_{empl} . For $\eta \in \{1, 2\}$, it has the rules:

- 1 $q_I(\text{department}, 0) \rightarrow \text{staff}(q(\text{down}_1, \mathbf{e}), \mathbf{e})$
- 2 $q(\text{employee}, \eta, y_n) \rightarrow q_{\text{empl}}(\text{down}_1, q(\text{down}_2, y_n))$
- 3 $q_{\text{empl}}(\text{data}, 1, y_n) \rightarrow \text{employee}(\text{data}(q_{\text{data}}(\text{down}_1), \mathbf{e}), q(\text{down}_2, y_n))$
- 4 $q(\text{subordinates}, 2, y_n) \rightarrow q(\text{down}_1, y_n)$
- 5 $q(\mathbf{e}, \eta, y_n) \rightarrow y_n$
- 6 $q_{\text{data}}(\text{data}, 1) \rightarrow \text{copy}(\text{down}_1).$

The grammar characterizing its output language looks as follows:

- 1 $q_I \rightarrow \text{staff}(q(\mathbf{e}), \mathbf{e})$
- 2 $q(y_n) \rightarrow q_{\text{empl}}(q(y_n)) \mid q(y_n) \mid y_n$
- 3 $q_{\text{empl}}(y_n) \rightarrow \text{employee}(\text{data}(q_{\text{data}}, \mathbf{e}), q(y_n))$
- 4 $q_{\text{data}} \rightarrow \text{copy}$

where $q_I, q, q_{\text{empl}}, q_{\text{data}}$, and copy are nonterminals. Selection of rules depending on input symbols and directions now has been replaced with nondeterministic choice. \triangleleft

Context-free tree grammars generalize context-free grammars to trees. Formally, a CFTG G can be represented by a tuple (E, Σ, P, E_0) where

- E is a finite ranked set of function symbols or nonterminals,
- $E_0 \subseteq E$ is a set of initial symbols of rank 0,
- Σ is the ranked alphabet of terminal nodes, and
- P is a set of rules of the form $q(y_1, \dots, y_n) \rightarrow \zeta$ where $q \in E$ is a nonterminal of rank $n \geq 0$. The right-hand side ζ is a tree built up from variables y_1, \dots, y_n by means of application of nonterminal and terminal symbols.

As for 2MTTs, inside-out (IO) and outside-in evaluation order for nonterminal symbols must be distinguished. Here, we use the IO or call-by-value evaluation

order. The least fixpoint semantics for the CFTG G is obtained straightforwardly along the lines for 2MTTs — simply by removing the corresponding directive components, i.e., by removing in the last line of the definition of $\llbracket \zeta \rrbracket$ for 2MTTs (Page 171) the op and $\llbracket op \rrbracket_t(v)$. In particular, this semantics assigns to every nonterminal q of rank $n \geq 0$, a set $\llbracket q \rrbracket \subseteq \mathcal{T}_\Sigma(Y)$ for $Y = \{y_1, \dots, y_n\}$. The language generated by G is $\mathcal{L}(G) = \bigcup \{ \llbracket q_0 \rrbracket \mid q_0 \in E_0 \}$.

It is easy to see that the output language of an input-linear MTT M can be characterized by a CFTG G_M , which can be constructed from M in linear time. During this construction, every rule $q(\mathbf{a}, \eta, y_1, \dots, y_n) \rightarrow \zeta$ is rewritten as a production $q(y_1, \dots, y_n) \rightarrow \zeta'$ where ζ' is obtained from ζ by deleting all occurrences of navigation operators.

The characterization of output languages for input-linear MTTs by CFTGs is useful because emptiness for (IO-)CFTGs is decidable using a similar algorithm as the one for ordinary context-free (word) grammars, and hence, can be done in linear time.

Theorem 12.5. *It can be decided in linear time for a CFTG G whether or not $\mathcal{L}(G) = \emptyset$.*

Here, we are interested in testing whether a given input-linear MTT M type checks w.r.t. input and output types I and O . Assume that I is given by a (possibly non-deterministic) BTA B with only productive states, i.e., for every state p of B , there exists a p -run of B on a tree. As a first step, we construct a new input-linear MTT M_B such that M_B 's range, i.e., the set $\tau_{M_B}(\mathcal{T}_\Sigma)$, is equal to $\tau_M(I)$. This is done by a straightforward product construction of the BTA B and the input-linear MTT M . Note that it may happen that M does not visit a certain subtree u of the input tree. In such a case, the checking of u w.r.t. B cannot be done by the new transducer M_B . This does not affect the corresponding output language though. We now construct the intersection 2MTT for M_B and the complement of the output type O . In case, O is given by a DBTA, this can be done along the lines of the proof of Theorem 12.2. If O is given by a DTTA, we rely on the construction from Theorem 12.3. Since M is input-linear, the intersection 2MTT is again input-linear — meaning that its range can be described by a CFTG (thus generating all “illegal outputs” of M w.r.t. I and O). Therefore, Theorem 12.5 gives us:

Theorem 12.6. *Assume M is an input-linear MTT where the ranks of input symbols are bounded, and let I and O denote input and output types for M where I is given by a BTA.*

1. *Assume that the output type O is specified with a DBTA and the maximal rank of states of M is bounded. Then M can be type checked relative to I and O in polynomial time.*
2. *Assume that the output type O is specified with a DTTA. Then M can be type checked relative to I and O in polynomial time — even in presence of unbounded ranks of states.*

The worst-case complexity bounds for the construction of Theorem 12.6 are exponential in $l \cdot (k + 1)$ (for output types given through DBTAs) or $k + 2$ (for output types given through DTTAs) where l is the maximal rank of states and k is the maximal rank of an input symbol of M . In practical applications, both k and l may be moderately small. Still, we want to point out that in case of input-linear MTTs, the intersection construction can be organized in such a way that only “useful” states are constructed. In order to see this, consider again an input-linear MTT M and a DBTA A (representing the incorrect output trees). The idea is to introduce for every q of M of rank $n + 1$, a *Datalog* predicate $q/n + 1$. Every rule $q(\mathbf{a}, \eta, y_1, \dots, y_n) \rightarrow \zeta$ of M then gives rise to the *Datalog* implication:

$$q(Y_0, \dots, Y_n) \Leftarrow \mathcal{D}[\zeta]_{Y_0}$$

where $\mathcal{D}[\zeta]_X$ (X is a variable) is defined by

$$\begin{aligned} \mathcal{D}[y_j]_X &= X = Y_j \\ \mathcal{D}[\mathbf{b}(\zeta_1, \dots, \zeta_m)]_X &= \delta(X, \mathbf{b}, X_1 \dots X_m) \wedge \mathcal{D}[\zeta_1]_{X_1} \wedge \dots \wedge \mathcal{D}[\zeta_m]_{X_m} \\ \mathcal{D}[q'(\zeta_1, \dots, \zeta_m)]_X &= q'(X, X_1, \dots, X_m) \wedge \mathcal{D}[\zeta_1]_{X_1} \wedge \dots \wedge \mathcal{D}[\zeta_m]_{X_m} \end{aligned}$$

and the variables X_1, \dots, X_m in the last two rows are fresh. For subsets $X, X_1, \dots, X_{m'}$ of the set of states of A , $\delta(X, \mathbf{a}, X_1 \dots X_{m'})$ denotes the fact that $(x, \mathbf{a}, x_1 \dots x_{m'}) \in \delta_A$ for all $x \in X$ and $x_j \in X_j$, $j = 1, \dots, m'$. A bottom-up evaluation of the resulting *Datalog* program computes for every $q/(n + 1)$, the set of all tuples (p_0, \dots, p_n) such that the translation of $\langle q, p_0 \dots p_n \rangle$ is non-empty. If we, additionally, want to restrict these predicates only to tuples, which may contribute to a terminal derivation of initial nonterminals $\langle q_0, p_f \rangle$, we may top-down query the program with queries $\Leftarrow q_0(p_f)$. Practically, top-down solving organizes the construction such that only useful nonterminals of the intersection grammar are considered. Using this approach, the number of newly constructed nonterminals often will be much smaller than the bounds stated in the theorem. A similar construction is also possible for the intersection of MTTs with the complements of DTTA languages.

The algorithm for *input-linear* MTTs can also be applied to *non-input-linear* 2MTTs. Then, the constructed *Datalog* program does no longer precisely characterize the non-empty functions of the intersection 2MTT because dependencies on input subtrees (viz. several transformations of the same input node) have been lost. Accordingly, a *superset* is returned. By means of CFTGs, we can express this observation as follows:

Theorem 12.7. *Let G_M be the CFTG constructed for a 2MTT M . Then*

$$\tau_M(\mathcal{T}_\Sigma) \subseteq \mathcal{L}(G_M) .$$

Since the CFTG still provides a safe *superset* of produced outputs, type checking based on context-free tree grammars is sound in the sense that if it does not flag an error, the transformation also will not go wrong. On the other hand, a flagged error may be possibly spurious, i.e., due to the over-approximation of

the output language through the CFTG. Consider a top-down transducer M with the rules

$$\begin{aligned} q_0(\mathbf{a}, 0) &\rightarrow \mathbf{c}(p(\mathit{down}_1), p(\mathit{down}_1)) \\ p(\mathbf{a}, \eta) &\rightarrow \mathbf{a}(p(\mathit{down}_1)) \\ p(\mathbf{b}, \eta) &\rightarrow \mathbf{b}(p(\mathit{down}_1)) \\ p(\mathbf{e}, \eta) &\rightarrow \mathbf{e} \end{aligned}$$

where p realizes the identity. In this case, the corresponding approximating CFTG G_M is rather coarse: It generates $\mathbf{c}(u, v)$ with $u, v \in \{\mathbf{a}, \mathbf{b}\}^* \mathbf{e}$ (seen as monadic trees). Exact tree copying, however, can be realized through the use of parameters: The transducer with the rules

$$\begin{aligned} q_0(\mathbf{a}, 0) &\rightarrow q(\mathit{down}_1, p(\mathit{down}_1)) \\ q(\sigma, \eta, y_1) &\rightarrow \mathbf{c}(y_1, y_1) \quad \text{for all } \sigma \in \Sigma \end{aligned}$$

and the same p -rules as M will realize the same translation as M . The context-free tree grammar for the resulting transducer now does not provide an over-approximation but precisely captures the output language of the top-down transducer M .

When approximating the output languages of general 2MTTs with CFTGs, we no longer can assume that the maximal number d of occurrences of nonterminals in a right-hand side of this grammar is bounded by a small constant. If d turns out to be unacceptably large, we still can apply Lemma 12.1 to limit the maximal number of occurrences of nonterminals in each right-hand side of the transducer to a number k which is the maximal rank of output symbols and states. This construction, however, introduces *stay*-moves and thus destroys input-linearity.

Example 12.4. For our example 2MTT $M_{y, \text{staff}}$, we construct for each rule of the transducer a production of the context-free tree grammar $G_{M_{y, \text{staff}}}$ and obtain:

$$\begin{aligned} q_I &\rightarrow \mathbf{staff}(q(\mathbf{e}, \mathbf{e}), \mathbf{e}) \\ q(y_b, y_n) &\rightarrow \mathbf{employee}(\mathbf{data}(q_{data}, y_b), q(\mathbf{boss}(q_{data}, \mathbf{e}), q(y_b, y_n))) \\ &\quad | y_n \\ &\quad | q(y_b, y_n) \\ q_{data} &\rightarrow \mathbf{copy} \end{aligned}$$

◁

12.4 Notes and References

Macro tree transducers [EV85] are a combination of top-down tree transducers and macro grammars [Fis68]. Macro grammars are just like context-free tree grammars (CFTGs), but produce strings (a CFTG can be seen as a special

macro grammar because terms are particular strings). In [Fis68], Fischer already distinguishes IO and OI for macro grammars, and proves that the corresponding classes of languages are incomparable (which also holds in the tree case). Our normal form of Lemma 12.1 can be seen as a variant of Fischer’s IO standard form, which in fact is very similar to Chomsky normal form of context-free grammars: there are exactly 2 or 0 nonterminals in every right-hand side of a grammar in IO standard form. A similar normal form might be possible for macro tree walking transducers, too, but will cause a larger size increase of the transducer (note that Fischer does not report on grammar sizes in his constructions). Fischer remarks, just after Corollary 3.1.6, that emptiness of IO macro languages can be reduced to emptiness of context-free languages, by simply dropping all parentheses, commas, argument, and terminal symbols. The resulting context-free (word) grammar generates the empty word if and only if the original language is empty. Since emptiness for a context-free grammar can be decided in linear time (see, e.g., [HMU01]), we obtain a linear time procedure for checking emptiness of IO context-free tree languages, as stated in Theorem 12.5. Context-free tree grammars were considered in [Rou70] and extensively studied in [ES77, ES78].

The fact that output languages of input-linear macro tree transducers are IO context-free tree languages is mentioned in Corollary 5.7 of [EV85] (the class of input-linear MTTs is called LMT_{IO} there). A 2MTT without up-moves is called “stay-MTT”. Results similar as the ones obtained in this chapter for 2MTTs (Chapter 12) have been obtained already in [MPS07] for the restricted case of stay-MTTs. For instance, Proposition 3 of that paper is similar to our Theorem 12.1, Theorem 2 of [MPS07] corresponds to our Theorem 12.2, and Theorem 5 of that paper corresponds to our Theorem 12.6. In Lemma 34 of [EM03b] it is shown that every transformation definable by a 2TT (0-ptt in that paper) can be realized by a stay-MTT (sMTT). Note however, that this construction leads to a blow-up of parameters, which keep track of all possible runs on subtrees. Thus, the algorithm of [MPS07] falls short on our efficient subclasses of 2TTs (Section 11.3), the polynomial complexity bounds would not hold true for 2TTs using stay-MTTs.

Just before Theorem 12.6, we describe how to incorporate an input type into an input-linear transducer, so that the corresponding output language is preserved. The technical details are exactly as in the proof of Theorem 3.2.1 of [ERS80] where this result was proved for top-down tree transducers. 2MTTs are essentially the same as the k -pebble macro tree transducers (k -PMTT) of [EM03b] for the case $k = 0$. For k -PMTTs, a normal form similar to our Lemma 12.1 was shown in Theorem 16 of [EM03b].

Chapter 13

Macro Forest Walking Transducers

Conceptually, XML documents are not trees but forests. Therefore, we extend the concept of tree walking transducers (without or with parameters) to a transformation formalism of forests. Forests are introduced in Definition 2.1.

Example 13.1. We consider again a transformation from company structures (cf. Figure 1.2) to lists of employees under new root nodes labeled **staff** (cf. Chapter 9). In contrast to tree walking transducers, forest transducers do not depend on a ranked alphabet. They can deal with arbitrarily many subtrees of nodes. Here, we define a transformation, which returns trees of the form **staff** $\langle f \rangle$ where f is a forest composed of **employee**-trees, i.e., trees of the form **employee**(**data**(...), **boss**(...)). The input trees are described by the DTD in the beginning of Chapter 10. Additionally to the operations *up*, *stay* as in tree walking transducers, forest transducers may use a directive *down* for proceeding to the first child as well as directives *left* and *right* for proceeding to the left or right sibling, respectively.

1	$q_I(\text{department}, 0)$	$\rightarrow \text{staff}\langle q(\text{down}, \epsilon) \rangle$
2	$q(\text{employee}, \eta, y_b)$	$\rightarrow \text{employee}\langle \text{data}\langle q_{\text{data}}(\text{down}) \rangle y_b \rangle$
3		$q(\text{down}, \text{boss}\langle q_{\text{data}}(\text{down}) \rangle)$
4		$q(\text{right}, y_b)$
5	$q(\text{data}, 2, y_b)$	$\rightarrow q(\text{right}, y_b)$
6	$q(\text{subordinates}, 3, y_b)$	$\rightarrow q(\text{down}, y_b)$
7	$q(\epsilon, \eta, y_b)$	$\rightarrow \epsilon$
8	$q_{\text{data}}(\text{data}, 2)$	$\rightarrow \text{copy}(\text{down})$

where state *copy* in Line 8 is meant to copy the forest f of a subtree **data** $\langle f \rangle$. The right-hand side of the second rule is a composition of three forests (Lines 2–4). The initial state is q_I , which means that we start with state q_I at the root

of the first tree of a forest. Here, the transducer walks only the first tree of an input forest. Note also that now rules may be selected depending on the current label of a node in the forest together with its forest direction. Thus, the transducer can check whether the current node is the leftmost node on the top-level (value 0), is on the top-level, but not leftmost (value 1), is leftmost but not on the top-level (value 2), or is neither leftmost nor on the top-level (value 3). \triangleleft

Definition 13.1 (2MFT). A *macro forest walking transducer* (2MFT for short) is a tuple $M = (Q, \Sigma, Q_0, R)$ where

- Q is a finite ranked set of states,
- Σ is a finite alphabet with $Q \cap \Sigma = \emptyset$,
- $Q_0 \subseteq Q^{(1)}$ is the set of initial states, and
- R is a finite set of rules of the form:

$$q(\epsilon, \eta, y_1, \dots, y_n) \rightarrow \zeta \quad \text{or} \quad q(\mathbf{a}, \eta, y_1, \dots, y_n) \rightarrow \zeta$$

with $\mathbf{a} \in \Sigma$, direction $\eta \in \{0, \dots, 3\}$ and $q \in Q^{(n+1)}$ where the right-hand sides are forests ζ of the following form:

$$\zeta ::= \epsilon \mid y_j \mid q'(op, \zeta_1, \dots, \zeta_{n'}) \mid \mathbf{b}(\zeta_1) \mid \zeta_1 \zeta_2$$

where $q' \in Q^{(n'+1)}$, $\mathbf{b} \in \Sigma$, $op \in \{up, stay, down, left, right\}$ and $j \in [n]$. Moreover, the right-hand sides for empty input forests ϵ must not contain occurrences of the operations *down*, *right* or *left*.

In case of several rules for the same q , the same direction η and the same symbol \mathbf{a} (or ϵ), we also write $\zeta_1 \mid \dots \mid \zeta_k$ to list all occurring right-hand sides. In case, no operation *up* is used, the 2MFT is also called *top-down* (short: *1MFT* or *MFT*). Likewise, if all states are of rank 1, i.e., have no accumulating parameters, the 2MFT is an (ordinary) *forest walking* transducer (short: *2FT*). Finally, a 1MFT without parameters is also called *forest transducer* (short: *1FT* or *FT*). As for macro tree walking transducers, in practice, states q may differ in their *ranks*, i.e., the numbers of their accumulating parameters plus 1. The set R of rules in Example 13.1 constitutes the transducer $M_{y, \text{staff}, f} = (Q, \Sigma, Q_0, R)$ with $Q = \{q_I, q, q_{data}, copy\}$ and $Q_0 = \{q_I\}$, which happens to be a 1MFT.

A forest transducer behaves similar to a corresponding tree transducer: While walking over the input forest, the transducer chooses rules corresponding to the current states, input symbols and directions at the respective current nodes in the input, and then evaluates the right-hand sides of the rules. Again, we just consider the *inside-out* (IO or call-by-value) strategy for evaluating parameters. There are two significant differences between forest walking and tree walking transducers: First, a forest walking transducer produces output *forests* and therefore, as an extra operation, also supports *concatenation* of output

forests. Second, the forest transducer has a different set of directions as well as a different set of navigational directives: *up* now means that the transducer moves to its ancestor in the input forest, *left* and *right* means that the transducer moves to its left or right sibling, whereas *down* means that the transducer moves to its first child. Formally, the semantics of these operations is defined by:

$$\begin{aligned} \llbracket up \rrbracket(vi) &= v \\ \llbracket left \rrbracket(vi) &= v(i-1) \quad \text{if } i > 0 \\ \llbracket right \rrbracket(vi) &= v(i+1) \\ \llbracket down \rrbracket(v) &= v0 \end{aligned}$$

Note that only the operations *down* and *right* have immediate equivalents in commands of a transducer on the *first-child next-sibling* encoding of forests as binary trees where they correspond to the commands $down_1$ and $down_2$, respectively. The *up*-command on the tree, on the other hand, may correspond to the forest commands *left* or *up* — depending on whether the current node is a right or left child. Analogously to Lemma 12.1, we find:

Lemma 13.1. *For every 2MFT M , there exists a 2MFT M' with*

- $\tau_{M'} = \tau_M$,
- *there are at most l occurrences of states on a right-hand side of rules in M' ,*
- $|M'| \in \mathcal{O}(|M| \cdot l^2)$

where l is the maximum of 2 and the maximal rank of states of M .

Note that it is unfortunately *not* possible to simulate 2MFTs by 2MTTs, which work on (possibly enriched) first-child next-sibling encodings of input and output trees. To see this, consider first the 1MFT case. As shown in [PS04], one can easily construct a 1MFT, which takes as input a binary tree with m nodes, and outputs a forest consisting of 2^m leaves, i.e., a string of length 2^m . Consider the height increase of the corresponding translation on encoded trees: it is *double*-exponential. However, the height-increase of MTTs is at most exponential (see [EV85]). Now consider the 2MFT case. Clearly, a 2MTT can translate a binary tree with m nodes into a monadic tree with 2^m nodes by doing a depth-first left-to-right traversal and at each step generating a duplicated state call in a parameter position. As intermediate sentential form, the transducer generates $q(\varepsilon, q(\varepsilon, \dots, q(\varepsilon, \mathbf{e})) \dots)$, which has 2^m -many occurrences of q ; it then replaces $q(\varepsilon, t)$ by $\mathbf{g}(t)$ where \mathbf{g} is an output symbol of rank 1. A 2MFT can generate the same sentential form, but can replace $q(\varepsilon, t)$ by tt , i.e., the forest of concatenating two copies of t . In this way, a forest consisting of 2^{2^m} -many \mathbf{e} 's is generated. On first-child next-sibling encodings, this corresponds to a tree of height 2^{2^m} . Thus, the translation on encodings has *double*-exponential size-to-height increase. However, it is not difficult to see that the size-to-height increase of 2MTTs is at most exponential.

13.1 Intersecting 2MFTs with Output Types

In this section, we consider general techniques for intersecting forest walking transducers with output types. Assume that we are given a regular forest language \mathcal{L} . Our goal is to construct for a given 2MFT M another 2MFT M' , which behaves similar to M but produces only outputs in \mathcal{L} . If \mathcal{L} describes the set of all invalid outputs, type checking for M , thus, reduces to checking emptiness of the transformation M' .

In order to provide a general construction for regular forest languages, let us first assume that \mathcal{L} is given as the language defined by a finite forest monoid A , i.e., $\mathcal{L} = \mathcal{L}(A)$. A *finite forest monoid* (short: FFM) can be considered as a deterministic bottom-up automaton, which combines the individual states for the trees u_i in a forest $f = u_1 \dots u_m$ by means of a monoid operation \circ (compare, e.g., the discussion in [BW05]). Formally, a *finite forest monoid* consists of a finite monoid G with a neutral element e , a finite subset $F \subseteq G$ of accepting elements, together with a function $up : \Sigma \times G \rightarrow G$ mapping a symbol of Σ together with a monoid element for its content to a monoid element representing a forest of length 1. A finite forest monoid accepts a forest f if $up^*(f) \in F$ where

$$\begin{aligned} up^*(f_1 f_2) &= up^*(f_1) \circ up^*(f_2) \\ up^*(\mathbf{a}(f')) &= up(\mathbf{a}, up^*(f')) \\ up^*(\epsilon) &= e \end{aligned}$$

Given a complete deterministic bottom-up forest automaton $A = (P, \Sigma, \delta, F_A)$, i.e., a DBTA operating on the first-child next-sibling representation of forests, we construct a finite forest monoid as follows. Let G be the monoid of functions $P \rightarrow P$ where the monoid operation is function composition. In particular, the neutral element of this monoid is the identity function. Moreover, the function up is defined by

$$up(\mathbf{a}, g) = p \mapsto \delta(\mathbf{a}, g(\delta(\mathbf{e})p)).$$

Finally, the set of accepting elements is given by

$$F = \{g \in G \mid g(\delta(\mathbf{e})) \in F_A\}.$$

On the other hand, every forest monoid G gives rise to a finite tree automaton A_G (running on first-child next-sibling representations) whose set of states is given by the elements of G . The transition function δ of A_G is defined by:

$$\delta(\mathbf{e}) = e \quad \text{and} \quad \delta(\mathbf{a}, g_1 g_2) = up(\mathbf{a}, g_1) \circ g_2.$$

Then the set of accepting states simply is given by the accepting elements of G . These constructions show that every recognizable forest language can be recognized by a finite forest monoid and vice versa. Although the FFM for a bottom-up tree automaton generally can be exponentially larger, this need not always be the case.

Example 13.2. For our current example, the BTA in Example 10.1 is not a complete deterministic bottom-up forest automaton. We get a complete DBFA by adding an extra error state \bullet . The new transition function δ' then is defined by:

$$\begin{aligned} \delta'(\mathbf{staff}, r_{\text{empl}} r_{\mathbf{e}}) &= r_{\text{staff}} & \delta'(\mathbf{staff}, r_{\mathbf{e}} r_{\mathbf{e}}) &= r_{\text{staff}} \\ \delta'(\mathbf{employee}, r_{\text{data}} r_{\text{empl}}) &= r_{\text{empl}} & \delta'(\mathbf{employee}, r_{\text{data}} r_{\mathbf{e}}) &= r_{\text{empl}} \\ \delta'(\mathbf{data}, r_{\text{name}} r_{\text{boss}}) &= r_{\text{data}} & \delta'(\mathbf{data}, r_{\text{name}} r_{\mathbf{e}}) &= r_{\text{data}} \\ \delta'(\mathbf{boss}, r_{\text{name}} r_{\mathbf{e}}) &= r_{\text{boss}} & \delta'(\mathbf{name}, r_{\text{content}} r_{\mathbf{e}}) &= r_{\text{name}} \\ \delta'(\mathbf{e}, \epsilon) &= r_{\mathbf{e}} \end{aligned}$$

and $\delta'(\mathbf{a}, r_1 r_2) = \bullet$ for all other combinations of a label $\mathbf{a} \in \Sigma$ and states $r_1, r_2 \in P \cup \{\bullet\}$. In the corresponding finite forest monoid $A = (G, \Sigma, \text{up}, F)$, the monoid G contains the following functions:

$$\begin{aligned} g_{\text{staff}} &= \{r_{\mathbf{e}} \mapsto r_{\text{staff}}\} \\ g_{\text{empl}} &= \{r_{\text{empl}} \mapsto r_{\text{empl}}, r_{\mathbf{e}} \mapsto r_{\text{empl}}\} \\ g_{\text{data}} &= \{r_{\text{boss}} \mapsto r_{\text{data}}, r_{\mathbf{e}} \mapsto r_{\text{data}}\} \\ g_{\text{boss}} &= \{r_{\mathbf{e}} \mapsto r_{\text{boss}}\} \\ g_{\text{name}} &= \{r_{\mathbf{e}} \mapsto r_{\text{name}}\} \\ g_{\text{dataBoss}} &= \{r_{\mathbf{e}} \mapsto r_{\text{data}}\} \\ g_{\text{content}} &= \{r_{\mathbf{e}} \mapsto r_{\text{content}}\} \\ g_{\bullet} &= \emptyset \\ Id &= \{r \mapsto r \mid r \in P\} \end{aligned}$$

where we have omitted all entries $r \mapsto \bullet$. Note that in this example, the forest monoid has only one element more than the underlying finite automaton. Also, the composition table of these functions is given by $Id \circ g = g \circ Id = g$ for all g and furthermore:

$$\begin{aligned} g_{\text{data}} \circ g_{\text{boss}} &= g_{\text{dataBoss}} \\ g_{\text{empl}} \circ g_{\text{empl}} &= g_{\text{empl}} \\ g \circ g' &= g_{\bullet} \quad \text{otherwise} \end{aligned}$$

For the function up , we find:

$$\begin{aligned} \text{up}(\mathbf{staff}, Id) &= g_{\text{staff}} & \text{up}(\mathbf{staff}, g_{\text{empl}}) &= g_{\text{staff}} \\ \text{up}(\mathbf{employee}, g_{\text{data}}) &= g_{\text{empl}} & \text{up}(\mathbf{employee}, g_{\text{dataBoss}}) &= g_{\text{empl}} \\ \text{up}(\mathbf{data}, g_{\text{name}}) &= g_{\text{data}} & & \\ \text{up}(\mathbf{boss}, g_{\text{name}}) &= g_{\text{boss}} & \text{up}(\mathbf{name}, g_{\text{content}}) &= g_{\text{name}} \\ \text{up}(\mathbf{a}, g) &= g_{\bullet} \quad \text{otherwise} \end{aligned}$$

where the set of accepting functions is given by $F = \{g_{\text{staff}}\}$. \triangleleft

Theorem 13.2. *For every 2MFT M and every finite forest monoid A , a 2MFT M_A can be constructed such that for all $f \in \mathcal{F}_\Sigma$,*

$$\tau_{M_A}(f) = \tau_M(f) \cap \mathcal{L}(A).$$

The size of M_A is in $\mathcal{O}(|M| \cdot |A|^{l \cdot (d+1)})$ where l is the maximal rank of a state q of M and d is the maximal number of occurrences of states in right-hand sides in M .

Proof. Let $M = (Q, \Sigma, R, Q_0)$ and $A = (G, \Sigma, up, F)$. For each state q in Q with rank $n + 1$ and all monoid elements $g_0, \dots, g_n \in G$, we generate new states for the intersection 2MTT M_A of the form $\langle q, g_0 g_1 \dots g_n \rangle$. Such a state is meant to generate all forests $f \in \mathcal{F}_\Sigma(\{y_1, \dots, y_n\})$ for which there is a run of A starting at the leaves y_i with monoid element g_i and reaching the root of f in g_0 . The rules of the new 2MFT M_A are:

$$\langle q, g_0 g_1 \dots g_n \rangle(\mathbf{a}, \eta, y_1, \dots, y_n) \rightarrow \zeta'$$

for every rule $q(\mathbf{a}, \eta, y_1, \dots, y_n) \rightarrow \zeta$ of M and $\zeta' \in \tau^{g_0 g_1 \dots g_n}[\zeta]$ where the sets $\tau^{g_0 g_1 \dots g_n}[\cdot]$ are inductively defined by:

$$\begin{aligned} \tau^{g_0 g_1 \dots g_n}[y_j] &= \{y_j \mid g_0 = g_j\} \\ \tau^{g_0 g_1 \dots g_n}[\mathbf{b}\langle \zeta \rangle] &= \{\mathbf{b}\langle \zeta' \rangle \mid up(\mathbf{b}, g') = g_0 \wedge \zeta' \in \tau^{g' g_1 \dots g_n}[\zeta]\} \\ \tau^{g_0 g_1 \dots g_n}[\epsilon] &= \{\epsilon \mid g_0 = e\} \\ \tau^{g_0 g_1 \dots g_n}[\zeta_1 \zeta_2] &= \{\zeta'_1 \zeta'_2 \mid g_0 = g'_1 \circ g'_2 \wedge \forall \nu : \zeta'_\nu \in \tau^{g'_\nu g_1 \dots g_n}[\zeta_\nu]\} \\ \tau^{g_0 g_1 \dots g_n}[q'(op, \zeta_1, \dots, \zeta_{n'})] &= \{\langle q', g_0 g'_1 \dots g'_{n'} \rangle(op, \zeta'_1, \dots, \zeta'_{n'}) \mid \forall \nu : \zeta'_\nu \in \tau^{g'_\nu g_1 \dots g_n}[\zeta_\nu]\} \end{aligned}$$

The set of initial states of M_A is $Q'_0 = Q_0 \times F$. By fixpoint induction, we verify for every state q of rank $n \geq 1$, every input forest $f \in \mathcal{F}_\Sigma$, every node $\vartheta \in \mathbf{Nodes}(f)$ and monoid elements g_0, \dots, g_n that:

$$\llbracket \langle q, g_0, \dots, g_n \rangle \rrbracket_f(\vartheta) = \llbracket q \rrbracket_f(\vartheta) \cap \{f' \in \mathcal{F}_\Sigma(Y) \mid up^*(f', g_1 \dots g_n) = g_0\} \quad (13.1)$$

where $Y = \{y_1, \dots, y_n\}$ and up^* is the extension of up to forests containing variables from Y , namely, for $\underline{g} = g_1 \dots g_n$, we have

$$\begin{aligned} up^*(y_i, \underline{g}) &= g_i \\ up^*(\epsilon, \underline{g}) &= e \\ up^*(\mathbf{a}\langle f' \rangle, \underline{g}) &= up(\mathbf{a}, up^*(f', \underline{g})) \\ up^*(f_1 f_2, \underline{g}) &= up^*(f_1, \underline{g}) \circ up^*(f_2, \underline{g}) \end{aligned}$$

The correctness of the construction follows from Equation 13.1.

For each state in M , we have at most $|A|^l$ new states in M_A , if l is the maximal rank of states in M . If we have d occurrences of states in the right-hand side of a rule r of M , we obtain $|A|^{l \cdot (d+1)}$ new rules for r in M_A . Therefore, the new 2MFT is of size $\mathcal{O}(|M| \cdot |A|^{l \cdot (d+1)})$ where l is the maximal rank of a state in M and d bounds the number of occurrences of states in right-hand sides in M . \square

forest f a *pair* of states $\langle p, p' \rangle$ so that the automaton A , when starting in p' to the right of f , possibly may arrive in state p to the left. Accordingly, the set Q' of M_A consists of all states

$$\langle q, p_0 p'_0 \dots p_n p'_n \rangle$$

where $q \in Q$ is of rank $n + 1$, i.e., has n accumulating parameters and $p_i, p'_i \in P$ for all i . Accordingly, the rules of the new 2MFT are of the form:

$$\langle q, p_0 p'_0 \dots p_n p'_n \rangle(\mathbf{a}, \eta, y_1, \dots, y_n) \rightarrow f' \quad \text{with } f' \in \tau^{p_0 p'_0 \dots p_n p'_n}[f]$$

for every rule $q(\mathbf{a}, \eta, y_1, \dots, y_n) \rightarrow f$ of M . The sets $\tau^{p_0 p'_0 \dots p_n p'_n}[\cdot]$ are defined by:

$$\begin{aligned} \tau^{p_j p'_j p_1 p'_1 \dots p_n p'_n}[y_j] &= \{y_j\} \\ \tau^{p_0 p'_0 p_1 p'_1 \dots p_n p'_n}[\mathbf{b}\langle \zeta' \rangle] &= \{\mathbf{b}\langle \zeta' \rangle \mid (p_0, \mathbf{b}, p'_1 p'_0) \in \delta \wedge (p'_2, \mathbf{e}) \in \delta \\ &\quad \wedge \zeta' \in \tau^{p'_1 p'_2 p_1 p'_1 \dots p_n p'_n}[\zeta]\} \\ \tau^{p_0 p'_0 p_1 p'_1 \dots p_n p'_n}[\epsilon] &= \{\epsilon \mid p_0 = p'_0\} \\ \tau^{p_0 p'_0 p_1 p'_1 \dots p_n p'_n}[\zeta_1 \zeta_2] &= \{\zeta'_1 \zeta'_2 \mid \exists p : \zeta'_1 \in \tau^{p_0 p p_1 p'_1 \dots p_n p'_n}[\zeta_1] \\ &\quad \wedge \zeta'_2 \in \tau^{p p'_0 p_1 p'_1 \dots p_n p'_n}[\zeta_2]\} \\ \tau^{p_0 p'_0 p_1 p'_1 \dots p_n p'_n}[q'(op, \zeta_1, \dots, \zeta_m)] &= \{q'(p_0 p'_0 p''_1 p'''_1 \dots p''_m p'''_m)(op, \zeta'_1, \dots, \zeta'_m) \mid \forall \nu : \zeta'_\nu \in \tau^{p''_\nu p'''_\nu p_1 p'_1 \dots p_n p'_n}[\zeta_\nu]\} \end{aligned}$$

The set of initial states of M_A then consists of all states $\langle q, p_0 p'_0 \rangle$ where $q \in Q_0$ and $p_0 \in P$ are accepting states of M and A , respectively, and $(p', \mathbf{e}) \in \delta$. The estimation of the size of the resulting transducer is similar to the case of forest monoids — only that we have to replace the number of monoid elements with the number of *pairs* of states. Thus, the new intersection transducer is of size $\mathcal{O}(|M| \cdot |A|^{2^{l \cdot (d+1)}})$. \square

13.2 Deciding Emptiness of 2MFTs

For deciding emptiness of a forest transducer M , we conceptually follow the approach taken for tree transducers. There, we first constructed an alternating tree walking automaton accepting the domain of M for which, in the second step, a nondeterministic tree automaton is constructed. In our case, this would mean that we first formally introduce the concept of alternating *forest walking* automata for which, in a separate construction, a nondeterministic forest automaton is constructed. In order to simplify this, we will not intermediately rely on forest walking automata. Instead, we consider for each forest f an *enriched* first-child next-sibling encoding through binary trees. This means that inside each node of the encoding, we additionally record whether or not the current tree node represents a node on the top-level of the forest. Let $\bar{\Sigma} = \{\bar{a} \mid a \in \Sigma\}$ denote a set of new symbols of rank 2. Then the ranked alphabet used by the encoding will be $\Sigma_2 = \Sigma \cup \bar{\Sigma} \cup \{\mathbf{e}, \bar{\mathbf{e}}\}$ where the barred symbols will only occur on the rightmost spine in the tree. A DTTA with two states $\{\mathbf{t}, \mathbf{n}\}$ can check whether a tree in \mathcal{T}_{Σ_2} is the enriched encoding of a forest or not.

Since the encoding is injective, it suffices for a forest transducer M to construct an ATWA M' , which defines the set of encodings of the domain of M . Then the set of states of the ATWA M' is given by $Q' = \{q'_0, \mathbf{t}, \mathbf{n}\} \cup Q \cup Q^{up}$ where $Q^{up} = \{q^{up} \mid q \in Q\}$ is a set of fresh copies of the states in Q and q'_0 serves as fresh initial state of M' . Assume that $q(\mathbf{a}, \eta, y_1, \dots, y_n) \rightarrow \zeta$ is a rule of M and $q_1(op_1, \dots), \dots, q_c(op_c, \dots)$ is the sequence of recursive calls in ζ . If $\eta \in \{2, 3\}$, i.e., if the rule is not applicable to nodes at the top-level of the input forest, then ATWA M' has the rule:

$$q(\mathbf{a}, \eta - 1) \rightarrow q'_1(op'_1) \wedge \dots \wedge q'_c(op'_c)$$

where:

$$q'_j(op'_j) = \begin{cases} q_j(stay) & \text{if } op_j = stay \\ q_j(down_1) & \text{if } op_j = down \\ q_j(down_2) & \text{if } op_j = right \\ q_j(up) & \text{if } op_j = left \text{ and } \eta = 3 \\ q_j^{up}(stay) & \text{if } op_j = up \end{cases}$$

Furthermore, for states q^{up} , the ATWA M' has the rules:

$$\begin{aligned} q^{up}(\mathbf{a}, 2) &\rightarrow q(up) \\ q^{up}(\mathbf{a}, 3) &\rightarrow q^{up}(up) \end{aligned}$$

If on the other hand $\eta \in \{0, 1\}$, i.e., the rule of the 2MFT refers to nodes at the top-level of the input forest, then the ATWA M' has the rule:

$$q(\bar{\mathbf{a}}, 2 \cdot \eta) \rightarrow q'_1(op'_1) \wedge \dots \wedge q'_c(op'_c)$$

where:

$$q'_j(op'_j) = \begin{cases} q_j(stay) & \text{if } op_j = stay \\ q_j(down_1) & \text{if } op_j = down \\ q_j(down_2) & \text{if } op_j = right \\ q_j(up) & \text{if } op_j = left \text{ and } \eta = 1 \end{cases}$$

For every rule $q(\mathbf{a}, 0) \rightarrow \zeta$ of M with $q \in Q_0$, we add the rule

$$q'_0(\bar{\mathbf{a}}, 0) \rightarrow \mathbf{t}(stay) \wedge q(stay)$$

where the rules for \mathbf{t} and \mathbf{n} simulate the computation of a top-down automaton to verify that the input tree is the enriched encoding of a forest. Thus, the rules of the ATWA for q'_0 are meant to spawn a subrun, which verifies the encoding and to spawn another subrun, which simulates an accepting run of the forest transducer on the binary encoding. In particular, the states q^{up} are auxiliary states to implement the operation up on the binary representation of the forest. More precisely, the rules for the state q^{up} perform the operation up as long as the current node is a right child. If the current node is a left child, a final up -operation is executed to arrive at the tree representation of the father node in state q . Overall, we find:

Theorem 13.4. *For every 2MFT M , an ATWA M' can be constructed in polynomial time such that $\mathcal{L}(M')$ is the set of enriched binary encodings of the set $\{f \mid \tau_M(f) \neq \emptyset\}$. In particular, $\mathcal{L}(M') \neq \emptyset$ iff $\tau_M \neq \emptyset$.*

We thus obtain an exponential algorithm for deciding emptiness of 2MFTs, which is optimal. Together with our intersection constructions, this algorithm then can be applied also for type checking 2MFT transducers w.r.t. regular input and output types.

In order to arrive at more tractable algorithms or sub-classes, we can apply the same ideas as for 2MTTs: in the first place, we can again approximate the set of output forests by means of a context-free forest grammar. A context-free forest grammar for the intersection with a regular forest language specified through a finite forest monoid is polynomial in the size of the grammar and the number of elements in the monoid and exponential only in $l \cdot (d + 1)$ where l is the maximal rank of nonterminals and d is the number of occurrences of nonterminals in right-hand sides. Emptiness for this forest grammar again can be checked in time linear in the size of the grammar. A practical implementation again may construct a *Datalog* program for the sets of useful nonterminals of the grammar. We can also generalize the notions of b -boundedness and strict b -boundedness for 2MFTs. While emptiness for b -bounded 2MFTs is decidable in polynomial time (where the exponent again depends on b^2), it is only strict b -boundedness, which is preserved by our intersection constructions. Here we only state the corresponding result for output types specified through finite forest monoids.

Theorem 13.5. *Assume M is a strictly b -bounded 2MFT and I and O are regular forest languages where I and O are given by a finite forest automaton and a finite forest monoid, respectively. Assume further that l is the maximal rank of a state of M and d is the maximal number of occurrences of state calls in right-hand sides. Then M can be type checked w.r.t. I and O in time polynomial in the sizes of M , the automaton for I and the automaton for O where the exponent linearly depends on $(b + 1)^2 \cdot l \cdot (d + 1)$.*

13.3 Notes and References

Top-down macro forest transducers have been introduced by Perst and Seidl in [PS04]. They are closely related to the top-down transducers of Maneth and Neven [MN99] but slightly more general. It was shown in [PS04] that, even though MFTs are more powerful than MTTs, they can be type checked with the same complexity bounds as macro tree transducers. This idea was extended to two-fold compositions of deterministic MTTs, in [MN08]. In [MBPS05], a general forest transformation language TL is introduced, which captures most features of XML transformation languages such as XSLT. The language TL supports full MSO pattern matching both for the selection of rules applicable at a node in the input tree and for navigation inside the input tree. Thus, 2MFTs can be considered as a sub-language of TL where rules are selected depending

on the current state and input label only, and where navigation is restricted to immediate neighbors in the input forest. The main contribution of that paper is to show how such transformations can be decomposed into three stay macro tree transducers running on the first-child next-sibling encoding of the XML documents in question. The semantics considered there was OI evaluation of nested calls but similar results can also be proven for IO evaluation, i.e., call-by-value parameter passing as considered here [Per07].

Chapter 14

Conclusion

In this part, we have reviewed basic constructions for tree walking transducers, which allow to obtain algorithms for type checking the transducers w.r.t. regular input and output types. There are three orthogonal variations in which the basic concept of a finite state machine can be made more expressive:

- top-down versus walking
- without parameters versus with parameters
- on ranked trees versus on unranked forests

At the very heart of our algorithms for type checking is to check whether or not a transducer realizes an empty translation. Already for the weakest, i.e., one-way top-down transducers, emptiness turns out to be complete for deterministic exponential time. Still, however, we were able to pin-point one major source for the complexity, namely the number of visits to the same input node. If the transducer visits the same node only constantly often, i.e., is b -bounded for some constant b , then emptiness becomes decidable in polynomial time.

The second ingredient of our algorithm are constructions for computing intersection transducers, i.e., transducers, which only produce outputs outside a specified regular set. Here, we considered regular sets as specified by bottom-up deterministic automata (or monoids, in case of forests) or by deterministic top-down automata. The latter construction for forest transducers was at least applicable to *output-linear* transducers, i.e., transducers, which use each of their accumulating parameters at most once. Two separate constructions are crucial, since translating top-down deterministic automata into bottom-up automata may incur an extra exponentiation in the number of states. Since these constructions preserve strict b -boundedness, we, thus, overall arrive at a general class of transducers for which type checking is polynomial.

Bibliography

- [ABS00] Serge Abiteboul, Peter Buneman, and Dan Suciu. *Data on the Web: From Relations to Semistructured Data and XML*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2000.
- [AD82] André Arnold and Max Dauchet. Morphismes et Bimorphismes d' Arbres. *Theoretical Computer Science*, 20:33–93, 1982.
- [ADB00] Hiyan Alshawi, Shona Douglas, and Srinivas Bangalore. Learning Dependency Translation Models as Collections of Finite-state Head Transducers. *Computational Linguistics*, 26:45–60, 2000.
- [AG68] Michael A. Arbib and Yehoshafat Give'on. Algebra Automata I: Parallel Programming as a Prolegomena to the Categorical Approach. *Information and Control*, 12(4):331–345, 1968.
- [AU71] Alfred V. Aho and Jeffrey D. Ullman. Translations on a Context-Free Grammar. *Information and Control*, 19(5):439–475, 1971.
- [Bak79] Brenda S. Baker. Composition of Top-Down and Bottom-Up Tree Transductions. *Information and Control*, 41(2):186–213, 1979.
- [Bar82] Miklós Bartha. An algebraic definition of attributed transformations. *Acta Cybernetica*, 5:409–421, 1982.
- [BCF⁺03] Scott Boag, Don Chamberlin, Mary F. Fernández, Daniela Florescu, Jonathan Robie, Jérôme Siméon, and Mugur Ştefănescu, editors. XQuery 1.0: An XML Query Language. W3C Working Draft. Available at <http://www.w3.org/TR/xquery/>, 2003.
- [BPS98] Tim Bray, Jean Paoli, and C. Michael Sperberg-McQueen, editors. *Extensible Markup Language (XML) 1.0*. W3C, first edition, 10 February 1998. Available at <http://www.w3.org/TR/1998/REC-xml-19980210>.
- [BPS08] Tim Bray, Jean Paoli, and C. Michael Sperberg-McQueen et. al., editors. *Extensible Markup Language (XML) 1.0*. W3C, fifth edition, 26 November 2008. Available at <http://www.w3.org/TR/2008/REC-xml-20081126>.

- [BR00] Srinivas Bangalore and Owen Rambow. Exploiting a Probabilistic Hierarchical Model for Generation. In *Proceedings of the 18th conference on Computational linguistics - Volume 1, COLING '00*, pages 42–48, Stroudsburg, PA, USA, 2000. Association for Computational Linguistics.
- [Bra68] Walter S. Brainerd. The Minimalization of Tree Automata. *Information and Control*, 13(5):484–491, 1968.
- [Bra69] Walter S. Brainerd. Tree Generating Regular Systems. *Information and Control*, 14(2):217–231, 1969.
- [BS01] Franz Baader and Wayne Snyder. Unification Theory. In John A. Robinson and Andrei Voronkov, editors, *Handbook of Automated Reasoning*, volume I, pages 447–533. Elsevier Science Publishers, 2001.
- [BW05] Mikolaj Bojańczyk and Igor Walukiewicz. Unranked Tree Algebra. Technical report, University of Warsaw, 2005.
- [CD99] James Clark and Steve DeRose, editors. XML Path Language (XPath) 1.0. W3C Recommendation, November 1999. Available at <http://www.w3.org/TR/xpath>.
- [CDG⁺07] Hubert Comon, Max Dauchet, Remi Gilleron, Christof Löding, Florent Jacquemard, Denis Lugiez, Sophie Tison, and Marc Tommasi. Tree Automata Techniques and Applications, 2007. Available at <http://www.grappa.univ-lille3.fr/tata> release October, 12th 2007.
- [CFZ82] Bruno Courcelle and Paul Franchi-Zannettacci. On the Equivalence Problem for Attribute Systems. *Information and Control*, 52(3):275–305, 1982.
- [CGKV88] Stavros Cosmadakis, Haim Gaifman, Paris Kanellakis, and Moshe Vardi. Decidable Optimization Problems for Database Logic Programs. In *Proceedings of the twentieth annual ACM symposium on Theory of computing, STOC '88*, pages 477–490, New York, NY, USA, 1988.
- [Cho03] Christian Choffrut. Minimizing Subsequential Transducers: A Survey. *Theoretical Computer Science*, 292:131–143, January 2003.
- [CM01] James Clark and Makoto Murata, editors. *RelaxNG Specification*. OASIS, 2001. Available at <http://www.oasis-open.org/committees/relax-ng>.
- [CNT04] Julien Carme, Joachim Niehren, and Marc Tommasi. Querying Unranked Trees with Stepwise Tree Automata. In Vincent van Oostrom, editor, *International Conference on Rewriting Techniques and*

- Applications*, volume 3091 of *Lecture Notes in Computer Science*, pages 105–118. Springer, June 2004.
- [Cou78] Bruno Courcelle. A Representation of trees by languages II. *Theoretical Computer Science*, 7(1):25–55, 1978.
- [DE98] Frank Drewes and Joost Engelfriet. Decidability of Finiteness of Ranges of Tree Transductions. *Information and Computation*, 145(1):1–50, 1998.
- [DS11] Deepak D’Souza and Priti Shankar, editors. *Modern Applications of Automata Theory*, volume 2 of *IISc Research Monographs*. World Scientific, 2011. To appear.
- [DT90] Max Dauchet and Sophie Tison. The Theory of Ground Rewrite Systems is Decidable. In *Logic in Computer Science*, LICS ’90, pages 242–248, June 1990.
- [EHS07] Joost Engelfriet, Hendrik Jan Hoogeboom, and Bart Samwel. XML Transformation by Tree-walking Transducers with Invisible Pebbles. In *Proceedings of the twenty-sixth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, PODS ’07, pages 63–72, New York, NY, USA, 2007.
- [Eis03] Jason Eisner. Learning Non-isomorphic Tree Mappings for Machine Translation. In *Proceedings of the 41st Annual Meeting on Association for Computational Linguistics*, ACL ’03, pages 205–208, 2003.
- [ELM08] Joost Engelfriet, Eric Lilin, and Andreas Maletti. Extended Multi Bottom-Up Tree Transducers. In *Proceedings of the 12th international conference on Developments in Language Theory*, volume 5257 of *DLT ’08*, pages 289–300. Springer-Verlag, Berlin, Heidelberg, 2008.
- [ELM09] Joost Engelfriet, Eric Lilin, and Andreas Maletti. Extended Multi Bottomup Tree Transducers. *Acta Informatica*, 46(8):561–590, October 2009.
- [EM99] Joost Engelfriet and Sebastian Maneth. Macro Tree Transducers, Attribute Grammars, and MSO Definable Tree Translations. *Information and Computation*, 154(1):34–91, October 1999.
- [EM03a] Abdessamad Echihabi and Daniel Marcu. A Noisy-channel Approach to Question Answering. In *Proceedings of the 41st Annual Meeting on Association for Computational Linguistics*, ACL ’03, pages 16–23, 2003.
- [EM03b] Joost Engelfriet and Sebastian Maneth. A Comparison of Pebble Tree Transducers with Macro Tree Transducers. *Acta Informatica*, 39(9):613–698, 2003.

- [EM03c] Joost Engelfriet and Sebastian Maneth. Macro Tree Translations of Linear Size Increase are MSO Definable. *SIAM Journal on Computing*, 32(4):950–1006, 2003.
- [EM06] Joost Engelfriet and Sebastian Maneth. The Equivalence Problem for Deterministic MSO Tree Transducers is Decidable. *Information Processing Letters*, 100(5):206–212, 2006.
- [EMS09] Joost Engelfriet, Sebastian Maneth, and Helmut Seidl. Deciding Equivalence of Top-down XML Transformations in Polynomial Time. *Journal of Computer and System Sciences*, 75(5):271–286, August 2009.
- [Eng75] Joost Engelfriet. Bottom-up and Top-down Tree Transformations - A Comparison. *Mathematical Systems Theory*, 9(2):198–231, 1975.
- [Eng77] Joost Engelfriet. Top-down Tree Transducers with Regular Look-ahead. *Mathematical Systems Theory*, 10(1):289–303, 1977.
- [Eng80] Joost Engelfriet. Some Open Questions and Recent Results on Tree Transducers and Tree Languages. In R.V. Book, editor, *Formal Language Theory; Perspectives and Open Problems*, pages 241–286. Academic Press, New York, 1980.
- [Eng09] Joost Engelfriet. The Time Complexity of Typechecking Tree-walking Tree Transducers. *Acta Informatica*, 46:139–154, March 2009.
- [ERS80] Joost Engelfriet, Grzegorz Rozenberg, and Giora Slutzki. Tree Transducers, L Systems, and Two-Way Machines. *Journal of Computer and System Sciences*, 20(2):150–202, 1980.
- [ES77] Joost Engelfriet and Erik Meineche Schmidt. IO and OI. I. *Journal of Computer and System Sciences*, 15(3):328–353, 1977.
- [ES78] Joost Engelfriet and Erik Meineche Schmidt. IO and OI. II. *Journal of Computer and System Sciences*, 16(1):67–99, 1978.
- [Ési80] Zoltán Ésik. Decidability Results Concerning Tree Transducers I. *Acta Cybernetica*, 5(1):1–20, 1980.
- [EV85] Joost Engelfriet and Heiko Vogler. Macro Tree Transducers. *Journal of Computer and System Sciences*, 31(1):71–146, 1985.
- [EW67] Samuel Eilenberg and Jesse B. Wright. Automata in General Algebras. *Information and Control*, 11(4):452–470, 1967.
- [FBCC⁺10] David A. Ferrucci, Eric W. Brown, Jennifer Chu-Carroll, James Fan, David Gondek, Aditya Kalyanpur, Adam Lally, J. William Murdock, Eric Nyberg, John M. Prager, Nico Schlaefer, and Christopher A. Welty. Building Watson: An Overview of the DeepQA Project. *AI Magazine*, 31(3):59–79, 2010.

- [FH07] Alain Frisch and Haruo Hosoya. Towards Practical Typechecking for Macro Tree Transducers. In *Proceedings of the 11th International Conference on Database Programming Languages*, DBPL '07, pages 246–260. Springer, Heidelberg, 2007.
- [Fis68] Michael J. Fischer. *Grammars with Macro-like Productions*. PhD thesis, Harvard University, Massachusetts, 1968.
- [Fri04] Alain Frisch. Regular Tree Language Recognition with Static Information. In *Programming Language Technologies for XML (PLAN-X)*, 2004.
- [FRR⁺10] Emmanuel Filiot, Jean-François Raskin, Pierre-Alain Reynier, Frédéric Servais, and Jean-Marc Talbot. Properties of Visibly Pushdown Transducers. In *Proceedings of the 35th International Conference on Mathematical Foundations of Computer Science*, MFCS '10, pages 355–367. Springer, Heidelberg, 2010.
- [FSM10] Sylvia Friese, Helmut Seidl, and Sebastian Maneth. Minimization of Deterministic Bottom-up Tree Transducers. In *Proceedings of the 14th International Conference on Developments in Language Theory*, DLT '10, pages 185–196. Springer, Heidelberg, 2010.
- [FSM11] Sylvia Friese, Helmut Seidl, and Sebastian Maneth. Earliest Normal Form and Minimization for Bottom-Up Tree Transducers. *International Journal of Foundations of Computer Science*, 2011. To appear.
- [Fül81] Zoltán Fülöp. On Attributed Tree Transducers. *Acta Cybernetica*, 5:261–279, 1981.
- [FV95] Zoltán Fülöp and Sándor Vágvölgyi. Attributed Tree Transducers Cannot Induce All Deterministic Bottom-up Tree Transformations. *Information and Computation*, 116(2):231–240, 1995.
- [FV98] Zoltán Fülöp and Heiko Vogler. *Syntax-Directed Semantics: Formal Models Based on Tree Transducers*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1998.
- [FV09] Zoltán Fülöp and Heiko Vogler. Weighted Tree Automata and Tree Transducers. In Manfred Droste, Werner Kuich, and Heiko Vogler, editors, *Handbook of Weighted Automata*, Monographs in Theoretical Computer Science, pages 313–403. Springer Berlin Heidelberg, 2009.
- [FW04] David C. Fallside and Priscilla Walmsley. XML Schema Part 0: Primer Second Edition. W3C Recommendation, October 2004. Available at <http://www.w3.org/TR/xmlschema-0/>.

- [Gaz06] Zsolt Gazdag. Decidability of the Shape Preserving Property of Bottom-up Tree Transducers. *International Journal of Foundations of Computer Science*, 17(2):395–414, 2006.
- [Gie88] Robert Giegerich. Composition and Evaluation of Attribute Coupled Grammars. *Acta Informatica*, 25(4):355–423, 1988.
- [GKM08] Jonathan Graehl, Kevin Knight, and Jonathan May. Training Tree Transducers. *Computational Linguistics*, 34(3):391–427, September 2008.
- [Gol78] E. Mark Gold. Complexity of Automaton Identification from Given Data. *Information and Control*, 37(3):302–320, 1978.
- [Gre78] Sheila A. Greibach. Hierarchy Theorems for Two-Way Finite State Transducers. *Acta Informatica*, 11:80–101, 1978.
- [GS84] Ferenc Gécseg and Magnus Steinby. *Tree Automata*. Akadémiai Kiadó, Budapest, Hungary, 1984.
- [GS97] Ferenc Gécseg and Magnus Steinby. Tree languages. In Grzegorz Rozenberg and Arto Salomaa, editors, *Handbook of Formal Languages, Volume 3: Beyond Words*, chapter 1, pages 1–68. Springer, Heidelberg, 1997.
- [HFC05] Haruo Hosoya, Alain Frisch, and Giuseppe Castagna. Parametric Polymorphism for XML. In *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '05, pages 50–62. ACM Press, New York, 2005.
- [HG01] Lynette Hirschman and Robert Gaizauskas. Natural Language Question Answering: The View from here. *Natural Language Engineering*, 7(4):275–300, December 2001.
- [HMU01] John E. Hopcroft, Rajeev Motwani, and Jeffrey D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, New York, second edition, 2001.
- [HP02] Haruo Hosoya and Benjamin C. Pierce. Regular Expression Pattern Matching for XML. *Journal of Functional Programming*, 13(6):961–1004, 2002.
- [HP03] Haruo Hosoya and Benjamin C. Pierce. XDuce: A Statically Typed XML Processing Language. *ACM Transactions on Internet Technology*, 3(2):117–148, 2003.
- [Iro61] Edgar T. Irons. A Syntax Directed Compiler for ALGOL 60. *Communications of the ACM*, 4:51–55, January 1961.

- [Kay03] Michael Kay, editor. XSL Transformations (XSLT) 2.0. W3C Working Draft, W3C, November 2003. Available at <http://www.w3.org/TR/xslt20>.
- [KB70] Donald E. Knuth and Peter B. Bendix. Simple Word Problems in Universal Algebras. In John Leech, editor, *Computational Problems in Abstract Algebra*, pages 263–297. Pergamon Press, 1970.
- [KG05] Kevin Knight and Jonathan Graehl. An Overview of Probabilistic Tree Transducers for Natural Language Processing. In Alexander F. Gelbukh, editor, *Computational Linguistics and Intelligent Text Processing, 6th International Conference, CICLing '05*, pages 1–24. Springer, Heidelberg, 2005.
- [KM02] Kevin Knight and Daniel Marcu. Summarization beyond sentence extraction: a probabilistic approach to sentence compression. *Artificial Intelligence*, 139(1):91–107, July 2002.
- [KMIS04] Christian Kirkegaard, Anders Møller, and Michael I. Schwartzbach. Static Analysis of XML Transformations in Java. *IEEE Transactions on Software Engineering (TSE)*, 30(3):181–192, March 2004.
- [Knu68] Donald E. Knuth. Semantics of Context-Free Languages. *Mathematical Systems Theory*, 2(2):127–145, June 1968.
- [Knu97] Donald E. Knuth. *Art of Computer Programming, Volume 1: Fundamental Algorithms (3rd Edition)*. Addison Wesley Longman Publishing Co., Inc., Redwood City, CA, USA, 1997.
- [KS81] Tsutomu Kamimura and Giora Slutzki. Parallel and Two-Way Automata on Directed Ordered Acyclic Graphs. *Information and Control*, 49(1):10–51, 1981.
- [LLN⁺11] Grégoire Laurence, Aurélien Lemay, Joachim Niehren, Slawomir Staworko, and Marc Tommasi. Normalization of Sequential Top-down Tree-to-word Transducers. In *5th International Conference on Language Automata Theory and Applications*, LNCS, Tarragona, Spain, 2011. Springer.
- [LMN10] Aurelien Lemay, Sebastian Maneth, and Joachim Niehren. A Learning Algorithm for Top-down XML Transformations. In *Proceedings of the twenty-ninth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems of Data*, PODS '10, pages 285–296. ACM, New York, NY, USA, 2010.
- [Löd11] Christof Löding. Basics on Tree Automata. In Deepak D’Souza and Priti Shankar, editors, *Modern Applications of Automata Theory*, volume 2 of *IISc Research Monographs*. World Scientific, 2011. To appear.

- [Mal11] Andreas Maletti. *Survey: weighted extended top-down tree transducers — part II: Application in machine translation*. *Fundamenta Informaticae*, 2011. submitted.
- [MBPS05] Sebastian Maneth, Alexandru Berlea, Thomas Perst, and Helmut Seidl. XML Type Checking with Macro Tree Transducers. In Chen Li, editor, *Proceedings of the twenty-fourth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, PODS '05, pages 283–294. ACM, New York, NY, USA, June 2005.
- [MFS11] Sebastian Maneth, Sylvia Friese, and Helmut Seidl. Type-Checking Tree Walking Transducers. In Deepak D'Souza and Priti Shankar, editors, *Modern applications of automata theory*, volume 2 of *IISc Research Monographs*. World Scientific, 2011. To appear.
- [MGHK09] Andreas Maletti, Jonathan Graehl, Mark Hopkins, and Kevin Knight. The Power of Extended Top-Down Tree Transducers. *SIAM Journal on Computing (SICOMP)*, 39:410–430, June 2009.
- [MLM00] Makoto Murata, Dongwon Lee, and Murali Mani. Taxonomy of XML Schema Languages using Formal Language Theory. In *Extreme Markup Languages*, 2000.
- [MM99] Inderjeet Mani and Mark T. Maybury. Advances in Automatic Text Summarization. pages 123–136. The MIT Press, 1999.
- [MN99] Sebastian Maneth and Frank Neven. Structured Document Transformations Based on XSL. In *Revised Papers from the 7th International Workshop on Database Programming Languages: Research Issues in Structured and Semistructured Database Programming*, DBPL '99, pages 80–98. Springer-Verlag, 1999.
- [MN04] Wim Martens and Frank Neven. Frontiers of Tractability for Type-checking Simple XML Transformations. In Alin Deutsch, editor, *Proceedings of the Twenty-third ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, PODS '04, pages 23–34. ACM Press, New York, 2004.
- [MN05] Wim Martens and Frank Neven. On the Complexity of Typechecking Top-down XML Transformations. *Theoretical Computer Science*, 336(1):153–180, 2005.
- [MN08] Sebastian Maneth and Keisuke Nakano. XML Type Checking for Macro Tree Transducers with Holes. In *Programming Language Technologies for XML (PLAN-X)*, 2008.
- [Moh00] Mehryar Mohri. Minimization Algorithms for Sequential Transducers. *Theoretical Computer Science*, 234(1–2):177–201, 2000.

- [MOS07] Anders Möller, Mads Østerby Olesen, and Michael I. Schwartzbach. Static Validation of XSL Transformations. *ACM Transactions on Programming Languages and Systems*, 29(4), 2007.
- [MPS07] Sebastian Maneth, Thomas Perst, and Helmut Seidl. Exact XML Type Checking in Polynomial Time. In Thomas Schwentick and Dan Suciu, editors, *11th International Conference Database Theory (ICDT'07)*, volume 4353 of *LNCS*, pages 254–268. Springer, 2007.
- [MS87] David E. Muller and Paul E. Schupp. Alternating Automata on Infinite Trees. *Theoretical Computer Science*, 54(2–3):267–276, 1987.
- [MS99] Christopher D. Manning and Hinrich Schütze. *Foundations of Statistical Natural Language Processing*. MIT Press, Cambridge, MA, USA, June 1999.
- [MS05] Anders Møller and Michael I. Schwartzbach. The Design Space of Type Checkers for XML Transformation Languages. In *Proc. 10th International Conference on Database Theory (ICDT '05)*, volume 3363 of *LNCS*, pages 17–36. Springer-Verlag, January 2005.
- [MSV00] Tova Milo, Dan Suciu, and Victor Vianu. Typechecking for XML Transformers. In *Proceedings of the nineteenth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS '00*, pages 11–22. ACM, New York, NY, USA, 2000.
- [MSV03] Tova Milo, Dan Suciu, and Victor Vianu. Typechecking for XML transformers. *Journal of Computer and System Sciences*, 66(1):66–97, 2003.
- [Myh57] John Myhill. Finite Automata and Representation of Events. Technical Report WADC-TR-57-624, Wright-Patterson Air Force Base, Ohio, 1957.
- [Ner58] Anil Nerode. Linear automaton transformations. *Proceedings of the American Mathematical Society*, 9(4):541–544, 1958.
- [Nev02] Frank Neven. Automata Theory for XML Researchers. *SIGMOD Record*, 31(3):39–46, September 2002.
- [OGV93] Jose Oncina, Pedro García, and Enrique Vidal. Learning Subsequential Transducers for Pattern Recognition Interpretation Tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(5):448–458, May 1993.
- [Per07] Thomas Perst. *Type Checking XML Transformations*. PhD thesis, Technische Universität München, Institut für Informatik, 2007.
- [PS04] Thomas Perst and Helmut Seidl. Macro Forest Transducers. *Information Processing Letters*, 89(3):141–149, 2004.

- [Rou70] William C. Rounds. Mappings and Grammars on Trees. *Mathematical Systems Theory*, 4(3):257–287, 1970.
- [Sei92] Helmut Seidl. Single-Valuedness of Tree Transducers is Decidable in Polynomial Time. *Theoretical Computer Science*, 106(1):135–181, January 1992.
- [Slu85] Giora Slutzki. Alternating Tree Automata. *Theoretical Computer Science*, 41(2-3):305–318, 1985.
- [Tak75] Masako Takahashi. Generalizations of Regular Sets and Their Application to a Study of Context-Free Languages. *Information and Control*, 27(1):1–36, 1975.
- [Tar72] Robert Endre Tarjan. Depth-First Search and Linear Graph Algorithms. *SIAM Journal on Computing*, 1(2):146–160, 1972.
- [Tha69] James W. Thatcher. Transformations and Translations from the Point of View of Generalized Finite Automata Theory. In *Proceedings of the 1st Annual ACM Symposium on Theory of Computing (STOC '69)*, pages 129–142. ACM, New York, NY, USA, 1969.
- [Tha70] James W. Thatcher. Generalized² sequential machine maps. *Journal of Computer and System Sciences*, 4(4):339–367, 1970.
- [Tha73] James W. Thatcher. Tree Automata: An Informal Survey. In Alfred V. Aho, editor, *Currents in the Theory of Computing*, pages 143–172. Prentice Hall, 1973.
- [Vir80] János Virágh. Deterministic Ascending Tree Automata I. *Acta Cybernetica*, 5(1):33–42, 1980.
- [Wu97] Dekai Wu. Stochastic Inversion Transduction Grammars and Bilingual Parsing of Parallel Corpora. *Computational Linguistics*, 23(3):377–403, September 1997.
- [Zac80] Z. Zachar. The Solvability of the Equivalence Problem for Deterministic Frontier-to-root Tree Transducers. *Acta Cybernetica*, 4:167–177, 1980.

Danke ...

Helmut Seidl für die Unterstützung und die vielen guten Ideen und Diskussionen.

Sebastian Maneth für die angenehme Zusammenarbeit und die hilfreichen Tipps von der anderen Seite der Erde.

Joachim Niehren für die äußerst interessanten und zielführenden Diskussionen sowie die detaillierten und umfangreichen Verbesserungsvorschläge.

Aurélien Lemay, merci pour notre discussion enrichissante à Lille.

Christian Pott, dass Du (als Nicht-Informatiker) die komplette Dissertation so detailliert Korrektur gelesen hast.

Alexander Herz für das ausführliche Korrekturlesen der Dissertation.

liebe Kollegen für die hilfreichen Diskussionen und, dass Ihr Euch immer wieder bereitwillig meine Vorträge über die Bäumchen angehört habt.

liebe Familie (Pott, Schulze-Eickenbusch und Friese) für Eure Unterstützung und Eure immer offenen Ohren.

Thomas Friese für Alles!