

Foundational Analysis Techniques for High-Level Transformation Programs

Ahmad Salim Al-Sibahi

Advisors:

Andrzej Wąsowski

Aleksandar Dimovski

Submitted: September 2017

IT UNIVERSITY OF COPENHAGEN

ليس التّقدم بتحسين ما كان
بل بالسّير نحو ما سيكون
جبران خليل جبران

*Progress is beyond improving what is;
Progress is moving towards what will be.*

GIBRAN KHALIL GIBRAN

Abstract

High-level transformation languages like TXL and Rascal, simplify writing and maintaining transformations such as refactorings and optimizations, by providing specialized language constructs that include powerful pattern matching, generic traversals, and fixed-point iteration. This expressivity unfortunately makes it hard to reason about the semantics of these languages, and there exists few automated validation and verification tools for these languages.

The goal of this dissertation is to develop foundational techniques and prototype tools for verifying transformation programs, based on techniques from traditional programming language research. At first, I present a systematic analysis of declarative high-level transformation languages, like QVT and ATL, seeking to clarify their expressiveness. The analysis showed that virtually all examined languages were Turing-complete and thus could not be treated specially for verification.

I describe a joint effort in designing and validating an industrial modernization project using the high-level transformation language TXL. The transformation was validated using the translation validation and symbolic execution, catching 50 serious bug cases in an otherwise well-tested expert-written transformation, and thus showing the need for verification tools in this domain.

Using these experiences, I develop a formal symbolic execution algorithm for a small transformation language capturing key high-level constructs. A proto-type white-box test generation tool was implemented using the algorithm, and evaluated on a series of refactorings and model transformations comparing favorably against baseline black-box test generator in terms of test coverage.

I present the first formal operational semantics for a large subset of Rascal, called Rascal Light, which has been validated by formally proving a series of semantic properties regarding strong typing, safety and termination. This creates a solid basis for further foundational work on high-level transformation languages.

Finally, I develop a formal abstract interpreter for Rascal Light, with a modular abstract domain design and a Schmidt-style abstract operational semantics, adapting the idea of widening by *trace memoization* to work with inputs from infinite domains. The technique is evaluated on real Rascal transformations, showing that it is possible to effectively verify target type and shape properties.

Resumé

Højniveau transformationsprog såsom TXL og Rascal gør det nemmere at skrive og vedligeholde transformationer, f.eks. refaktoreringer og optimeringer, ved at understøtte specielle sprogs konstruktioner inklusiv kraftfulde mønstersamsvarsoperationer (orig. *pattern matching*), generisk gennemgang af datastrukturer, og fikspunktsiteration. Denne ekspresivitet gør det dog desværre svært at få en fuld forståelse for semantikken til disse sprog, og der eksisterer derfor kun få automatiske værktøjer der understøtter validering og verifikation af dem.

Hovedformålet med denne afhandling er at udvikle fundamentale teknikker og værktøjsprototyper som understøtter verifikation af transformationsprogrammer, baseret på teknikker fra traditionel programmeringssprogsforskning. I starten, præsenterer jeg en systematisk analyse af deklarativ højniveau transformationsprog, såsom QVT og ATL, for at tydeliggøre deres ekspresivitet. Analysen viste at stort set alle undersøgte sprog var Turing-komplette og kunne derfor ikke blive særbehandlet i verifikationsøjemed.

Jeg beskriver en fælles indsats hvis formål er at designe og validere et industrielt moderniseringsprojekt ved brug af højniveau transformationsproget TXL. Transformationen var valideret ved brug af oversættelsesvalidering (orig. *translation validation*) og symbolsk eksekvering, og opdagede 50 seriøse programfejl i en ellers veltestet transformation skrevet af en ekspert, hvilket fremviser behovet for verifikationsværktøjer i dette område.

Gørende brug af disse erfaringer, udvikler jeg en formel symbolsk eksekveringsalgoritme for et lille transformationsprog der modellerer de vigtigste højniveau konstruktioner. Jeg implementerede en white-box testgenereringsværktøjsprototype ved brug af algoritmen, og evaluerede denne vha. en række refaktoreringer og modeltransformationer, hvilket viste gode resultater i forhold til programgrendækning sammenlignet med black-box testgenereringsværktøjet som blev brugt som basislinje.

Jeg præsenterer den første formelle operationelle semantik for en stor delmængde af sproget Rascal, kaldet Rascal Light, hvilket er valideret ved at bevise en række semantiske egenskaber omkring stærkt typedisciplin, sikkerhed, og programafslutning. Dette lægger et solidt fundament for videre formelt arbejde omkring højniveau transformationsprog.

Afslutningsvis, udvikler jeg en formel abstrakt fortolker til Rascal Light, med et tilhørende modulært abstrakt domæne og en abstrakt operationel semantik i stil med Schmidt, og tilpasser konceptet om udvidelse (orig. *widening*) via memoisering af programmer til at virke med input fra uendeligt store domæner. Teknikken evalueres ved brug af rigtige Rascal transformationsprogrammer, hvilket viser at det er muligt at effektivt verificere de egenskaber omkring typer og form som vi havde til mål.

Foreword

How long a time is three years? It is too long, you get to know your project so well that you become an expert in its extremely narrow subject. It is too short, it ended at the moment you just started to know enough in the field to dare to tackle more complex, more interesting problems. Maybe, it is just enough time to earn a PhD.

The thing I recall most is my failing attempts at solving the problems I encountered, moreso than the successes that I present in this dissertation. Yet, these failures, remaining hidden, under the carpet, in my mind, are what makes research all the worthwhile. For how can we otherwise make worthwhile contributions, if there is nothing to dare; is it truly research if it can not fail, if it does not fail, if it will not fail *again*?

Acknowledgements

I would like to thank my family, who helped me with material and moral support, and especially during my PhD studies. I would like to thank the Danish society for funding my education throughout my life, all the way up to this point. This project in particular was supported mainly by The Danish Council for Independent Research under a Sapere Aude project grant no. 0602-02327B., VARIETE, and partially by ARTEMIS JU under grant agreement n°295397 together with Danish Agency for Science, Technology and Innovation.

I would like to thank my supervisors, Andrzej Wąsowski and Aleksandar Dimovski for comments, discussions and guidance during my studies, and for putting up with my occasional cloud-headedness. I would like thanks the rest of my collaborators, Claus Brabrand, Thomas Jensen, Alexandru Florin Iosif-Lazăr, Juha Erik Savolainen and Krzysztof Sierszecki, for their contributions to this project as well. I would like to

thank the committee, Riko Jacob, Xavier Rival and Tijs van der Storm, for evaluating this dissertation and providing valuable feedback.

I would like to thank Thomas Jensen, again, and his group, Celtique, at INRIA Rennes, for the great welcome I received during my stay-abroad, and for the interesting discussions during the reading groups. I would like to thank Lars Birkedal and his group at Aarhus University, and similarly Constantin Enea and Michaela Sighireanu and their group at Paris Diderot University (VII) for giving me the opportunity to present my work and discuss it during the two short visits I had during my PhD.

I thank Karl Portratz for introducing our team to the industrial modernization project at Danfoss, Rolf-Helge Pfeiffer for implementing the modernization transformation and Claus Brabrand for discussions on its design. I would like to thank Paul Klint and the Rascal team at CWI Amsterdam for answering my questions about the formal Rascal semantics in details and providing corrections on the corresponding report. I would like to thank Rasmus Ejlers Møgelberg for discussions about the correctness of the fixed-point domain abstraction theorem in the Rascal abstract interpreter.

Formal Notation

- For a binary relation $r \subseteq A \times A$, r^* is used to denote the reflexive-transitive closure, and similarly r^+ for the transitive (non-reflexive) closure.
- For a set A , $\wp(A)$ is used to denote the power set.
- For a function $f \in A \rightarrow B$, $\text{dom } f$ is used to represent all defined inputs i.e., $\{x \mid \exists y. y = f(x)\}$, $\text{img } f$ to represent all defined outputs $\{y \mid \exists x. y = f(x)\}$, and $\text{graph } f$ to represent the set $\{(x, f(x)) \mid x \in \text{dom } f\}$.
- For a function $f \in A \rightarrow B$, $\text{dom } f$, $f[a \mapsto b]$ is used to represent function updates, so that $f[a \mapsto b](a) = b$ and $f[a \mapsto b](a') = f(a')$ when $a' \neq a$.
- Approximate (symbolic or abstract) semantic components, sets and operations are marked with a wide hat \hat{a} .
- A sequence of es is represented by either marked by an underline \underline{e} or explicitly summarized e_1, \dots, e_n , the empty sequence is denoted by ε and the concatenation of sequences \underline{e}_1 and \underline{e}_2 is denoted $\underline{e}_1 \underline{e}_2$.
- Notation is overloaded in an intuitive manner for operations on sequences, so given a sequence \underline{v} , v_i denotes the i th element in the sequence, and $\underline{v} : t$ denotes the sequence $v_1 : t_1, \dots, v_n : t_n$.
- Operators from different sets are distinguished with a subscript, e.g., \top_A represents the top element of domain A ; the subscripts are left out when the intended use of operators is clear from context.

Contents

Foreword	i
Formal Notation	iii
Contents	v
List of Figures	ix
List of Tables	xi
1 Introduction	1
1.1 Context	1
1.2 Problem Definition	5
1.3 Aim and Scope	7
1.4 Contributions	9
1.5 Overview	12
2 High-Level Transformation Languages	13
2.1 Declarative Transformation Languages	13
2.2 Program Transformation Languages	17
2.3 Programs and Paradigms	19
3 Expressiveness of Declarative Model Transformation Languages	21
3.1 Surveys on Model Transformations	21
3.2 Computational expressiveness	22
3.3 Analysis Method	24
3.4 Analysis of computational expressiveness	25
3.5 Expressiveness and Verification	31

4	Validating a Software Modernization Transformation	33
4.1	Motivating Example	35
4.2	The Modernization Case Study	37
4.3	Designing the Transformation	41
4.4	Validating the Transformation	45
4.5	Bug Analysis	48
4.6	Formal Justification of the Procedure	52
4.7	Experiential Reflections	55
4.8	Related Modernization Efforts	56
4.9	Recap	57
5	Symbolic Execution of High-Level Transformation Languages	59
5.1	Overview	60
5.2	A Demonstration Language	63
5.3	Symbolic Execution	70
5.4	Evaluation	84
5.5	Related Tools and Approaches	91
5.6	Recap	92
6	The Formal Semantics of Rascal Light	95
6.1	Abstract Syntax	97
6.2	Semantics	101
6.3	Semantics Properties	130
6.4	Formal Semantics, Types and Intended Semantics	132
6.5	Recap	132
7	Verifying Rascal Transformations using Inductive Refinement Types	133
7.1	Motivating Example	134
7.2	Formal Language	139
7.3	Abstract Domains	141
7.4	Abstract Semantics	157
7.5	Evaluation	177
7.6	Related Work	184
7.7	Recap	185
8	Conclusion	187
	Bibliography	189
A	Modernization Transformation Proofs	217

Contents	vii
B SymexTRON Auxiliary Functions	219
C SymexTRON Proofs	223
D TRON Subject programs	225
E Rascal Light Semantics Proofs	237
F Rascal Light Verification Proofs	275
G Rascal Subject programs	279
H Rascal Verification Output	293

List of Figures

1.1	Constant Inlining Transformation	3
2.1	Meta-models for the Families to Persons [Allilaire and Jouault, 2007] transformation	14
2.2	The Families to Persons [Allilaire and Jouault, 2007] transformation in ETL	15
2.3	Two graph rewriting operations in GROOVE for the Families to Persons transformation	17
2.4	Negation Normal Form	18
2.5	Negation Normal Form Transformation in Stratego/XT	19
2.6	Negation Normal Form Transformation in Haskell using Uniplate	19
4.1	Anonymized example validation function from the Danfoss code base	35
4.2	Example TXL transformation rule from the modernization project: translating a conditional statement into a ternary expression.	44
4.3	A stub for an unknown Boolean function.	47
5.1	An example execution of the <i>Rename-Field</i> refactoring: rename the definition of <i>credit</i> to <i>balance</i> and update all references accordingly.	60
5.2	A simplified version of the <i>rename-field</i> refactoring example in TRON	61
5.3	High level architecture of the symbolic executor.	63

5.4	A heap instantiating one model of Figure 5.2a, inspired by the Account class in Figure 5.1. Dots (●) represent instances, diamond affixed lines represent containment links, dashed lines represent simple links.	66
5.5	Concrete semantics for TRON	68
5.6	Concrete semantics for TRON (Cont.)	69
5.7	An example heap for an initial state of an execution	74
5.8	One of the paths when symbolically executing the example Rename-Field refactoring, starting with the symbolic state presented in Figure 5.7. Double-stroked arrows represent deep containment constraints.	75
5.9	Symbolic execution rules for standard statements (σ is a symbolic variable store, h is a symbolic heap)	77
7.1	Expression Simplification	134
7.2	Abstract Operational Semantics Rules for Basic Expressions .	161
7.3	Abstract Operational Semantics Rules for Basic Expressions (Cont.)	162
7.4	Abstract Operational Semantics Rules for Basic Expressions (Cont. 2)	163
7.5	Inferred shape for Negation Normal Form transformation . .	182
7.6	Inferred shape for Rename Struct Field transformation	183
7.7	Flattening auxiliary begin statement	183

List of Tables

3.1	Paradigms of transformation languages	26
3.2	Expressiveness of transformation languages	27
4.1	Erroneous transformation cases caught by each step of the validation process.	48
5.1	Relating TRON features to existing high-level transformation languages	64
5.2	Characteristics and baseline meta-model coverage for the subject programs. Here LOC indicates lines of code, where the first component of the summation is the size of the data model and the second component is the size of the transformation.	87
5.3	Results of running the test generators on subject programs. The branch coverage indicates the percentage of the total reachable branches covered, and the time indicates the execution time of the tool from start to finish on the evaluation subjects.	88
7.1	Subject Transformations Statistics and Target Properties . . .	180
7.2	Results of abstractly interpreting subject programs. Runtime is median of three runs.	181

Chapter 1

Introduction

1.1 Context

Domain-specific languages or *DSLs* [Bentley, 1986; Hudak, 1996; van Deursen et al., 2000; Mernik et al., 2005; Fowler, 2011] allow expressing models and programs pertinent to a business domain (e.g., finance [Christerson et al., 2013], robotics [Nordmann et al., 2014], cryptography [Erkök and Matthews, 2009]) using a vocabulary familiar to domain users. These languages eschew generality in favor of providing a clear correspondence between domain requirements and models, being important for maintainability and evolution of modelled systems, and easing communication between domain users and technical experts.

DSL engineering relies heavily on the use of *transformations*, which allow converting between structured models or programs. Examples of such transformations include:

Model Translation Translating between structured models preserving consistency [Jouault et al., 2008], e.g., a class diagram to a relational diagram.

Desugaring Transformation of constructs from a language extended with *syntactic sugar* [Landin, 1964]—that primarily target easing programming—to constructs in a core language, e.g., translating **for**-loops to **while**-loops.

Refactoring Restructuring of source code [Fowler, 1999] to increase readability and maintainability while largely preserving intended semantics, e.g., renaming the field of a class.

Code Generation Converting constructs from one language to another for execution or further processing [Aho et al., 1986b], e.g., generating assembly code from an imperative language.

Optimization Transforming programs possibly given some static information in order to improve runtime and memory use [Aho et al., 1986c], e.g., software pipelining and partial evaluation [Jones et al., 1993].

Most practical DSLs have large abstract syntax—often spanning hundreds of syntactic elements—and correspondingly rich semantics; as a result, this makes writing transformations for these languages a tedious and error-prone process, more so using traditional programming languages like Java [Gosling et al., 2014], C[‡] [Hejlsberg et al., 2011] or Haskell [Marlow, 2010].

To simplify the task of writing and maintaining such transformations, specialized high-level transformation languages have been proposed across different communities in software language engineering (e.g. Rascal [Klint et al., 2011], Stratego/XT [Bravenboer et al., 2008]), programming languages (e.g. TXL [Cordy, 2006], Uniplate [Mitchell and Runciman, 2007]) and model-driven engineering (e.g. QVT [Object Management Group, 2016], ATL [Jouault et al., 2008], ETL [Kolovos et al., 2008]). Common for these languages is that they all provide rich features for traversing and manipulating abstract syntax, powerful pattern matching and querying, backtracking and generalized looping constructs.

Figure 1.1 shows an example Rascal programs which uses these rich features to inline constant declarations in a simple expression language. The `constsValid` function pattern matches on the declarations inside a module, checking whether there are duplicate declarations of constants using arbitrary elements match patterns `*_` to skip elements and non-linear patterns to check existence of two equal names (notice the two instances of `x` in the same pattern). The `inlineConsts` function iterates over all constant declarations using a pattern matching **for**-loop, removing the constant declaration from the list of declarations and traversing the body expression using **bottom-up visit** replacing all references to the constant with the provided value.

I will further illustrate how these features are concretely realized in Rascal:

```

1 data Decl = constdecl(str name, int val)
2           | vardecl(str name);
3 data Expr = mult(Expr l, Expr r) | add(Expr l, Expr r)
4           | var(str name) | const(int val);
5 data Mod = \mod(list[Decl] decls, Expr body);
6
7 bool constsValid(Mod m) {
8   switch (m.decls) {
9     case [*_ , constdecl(x, _), *_ ,
10          constdecl(x, _), *_ ]: return false;
11     case _: return true;
12   }
13 }
14 Mod inlineConsts(Mod m) {
15   assert constsValid(m);
16   for (cd: constdecl(x, v) <- m.decls) {
17     m.decls = m.decls - cd;
18     m.body = bottom-up visit(m.body) {
19       case var(x) => const(v)
20     };
21   };
22   return m;
23 }

```

Figure 1.1: Constant Inlining Transformation

- Traversals use **visit** which allows automated rewriting of abstract syntax trees using a sequence of pattern matching cases using either **bottom-up** ordering as in the example or **top-down** ordering.
- Pattern matching supports a rich set of operations, such as the constructor matching, arbitrary elements matching, and non-linear matching presented in the example. Other operations include type-based matching, negative matching, and deep matching that pattern matches arbitrarily nested values inside a structure.
- Backtracking happens both internally when constraints are enforced on non-deterministic matches—e.g., in our example when looking for two equal constant declarations—and externally, where a user can use **fail** in a target expression to rollback the state and trigger another possible match in a non-deterministic pattern or fall through.
- There are various looping constructs including the generalized **for**-loop used in our example that supports pattern matching iteration, a **solve**-loop that finds a fixedpoint of an operation over a set of variables, and fixedpoint-iterating traversal strategies (**innermost** and **outermost**).

This power of expressiveness unfortunately comes with a heavy price: the semantics of these languages becomes increasingly hard to reason about. There exist few automated tools that can handle validation and verification of transformation programs, and those that exist are limited in expressiveness of language constructs or supported properties. They can largely be divided in three categories:

1. *Black-box* techniques—e.g., Claessen and Hughes [2000], Yang et al. [2011], Finot et al. [2013] and Mongiovi et al. [2014]—rely solely on input and output specifications to validate the transformation programs. These techniques are independent of transformation program definitions which eases reuse for new languages, but can rarely provide a complete verification and can potentially miss a large amount of internal logic that is left unvalidated.
2. *White-box* techniques—e.g., Jackson et al. [2011], Büttner et al. [2012], and Clarisó et al. [2016]—that use the definitions of program definitions, but are limited in the the expressiveness of language constructs they handle (non-Turing complete) and thus do not support more interesting transformation programs such as refactorings or optimizations. Furthermore, many rely on ad-hoc encoding of the complete transformation programs to a logical formula to use in a solver (e.g. SMT or CLP) which is a process that is error prone, and hard to extend because there is a lack of explicit representation of the program state. Those that do use static analysis-based—e.g., Cuadrado et al. [2017]—techniques seem to scale better with regards to language expressiveness and properties, but currently only focus on simple type and structural properties and not on more complex properties such as abstraction of data (integers, strings, etc.), inductive properties or shapes of collections.
3. *Translation validation*-based techniques [Samet, 1976; Pnueli et al., 1998; Rival, 2004] can automatically prove sophisticated semantic properties about object programs, but only work for individual runs of a transformation. The strong guarantees they provide however are only for the provided input programs at each run, and no guarantees are provided on other inputs. This entails that errors are only detected late in the process, and fixing the errors at that stage can be very costly and time consuming.

The existing work therefore leaves space for improvement with regards to validating and verifying fully-expressive transformation programs effectively using their definitions.

1.2 Problem Definition

The void left by the existing work has consequences that go beyond mere technical curiosity. The availability of quality assurance techniques (including validation and verification) is a key characteristic for the trustworthiness of a programming language and paves way for adoption of the language in use of critical software systems. This is especially true of modern critical systems which are steadily increasing in complexity and it is here the aforementioned transformation languages would really shine in increasing readability and maintainability of such systems.

What particularly motivates this thesis is that despite the rich applied tradition of automated formal validation and verification methods in programming language research—deductive methods, abstract interpretation, symbolic execution, model checking—we paradoxically lack tools for quality assurance of transformation languages. If we are able to use these foundational techniques, we would benefit from the existing wealth of experience, be able to more effectively verify target transformations, and allow extending the techniques towards supporting properties more relevant to these languages—pertaining to typed syntactic structures and input/output relations—in ways more actively exploiting the information provided by the high-level features.

In summary, the problem statement addressed by this dissertation is:

How is it possible to extend automated formal methods from traditional programming language research to scale to handling the expressive features found in realistic high-level transformation languages?

Concretely, the problem statement will be addressed by addressing each of the following subproblems:

P1 Unclearly regarding expressiveness of available features in model and graph transformation languages. While some high-level transformation languages like TXL and Rascal are clearly programming languages with features targeting transformations, other languages that focus on transforming models or graphs like ATL and

ETL have a more declarative feel with limited use of explicit loops and recursion. As a result, it is unclear whether these languages in general are able to express general computation (Turing-complete), and whether it is possible to exploit the lack of expressiveness to provide better analysis methods. Clarifying this problem also clarifies whether the existing techniques that handle limited subsets are sufficient for handling more interesting transformations that have more complete use of available language features.

- P2 Available techniques for systematic testing of transformations do not effectively exploit definition of transformation programs** Traditionally, techniques that exploit program definitions (*white-box*) often generate better test cases than those that only use input specification (*black-box*), and I believe that this is also the case for transformations. This problem therefore primarily concerns how to extend existing effective white-box techniques to handle the language constructs and derived constraints present in fully-expressive transformation languages. This is further complicated by the fact that transformations manipulate rich typed data structures—containing a wealth of different types, that each are often both deep (many nested containments) and wide (many subcomponents)—and collections of these structures and so any naive algorithm will easily hit a combinatorial explosion.
- P3 Lack of formal techniques for automated verification of type and shape properties for fully-expressive transformations** Most interesting transformations like normalization procedures, refactorings and optimizations require the full expressive power of the high-level transformation languages to be encoded. Existing techniques which encode the complete transformation as a logical formula for an SMT solver (or similar systems) are unlikely to scale to this approach since their employed techniques are limited with regards to loops and complex recursion and bounded model finding approaches do not scale for large systems. Conversely, techniques like translation validation are useful for verifying complex properties but only work for individual runs of a transformation. This problem thus regards finding a technique for verifying transformations that i) runs automatically ii) is easily extensible, and iii) works for all executions of a transformation.

Importantly, the technique should be able to verify interesting properties about the transformations themselves, i.e., properties that relate to the typed abstract structures manipulated by the transformations such as types and shapes relating input and output values.

1.3 Aim and Scope

The primary aim of this thesis is to develop foundational theories and tools for high-level transformation languages based on traditional programming language techniques.

This aim will be fulfilled by satisfying the following objectives:

- O1 Identify key features of high-level transformation languages, classifying the languages according to their computational expressiveness** The purpose of this objective is two-fold i) it aims to set the context for objectives O3 and O4 by identifying the key features that should be handled by the respective analyses, and ii) it addresses problem P1 by clarifying the expressiveness of declarative transformation languages. I partly reach this objective by systematically investigating the documentation and relevant primary and survey papers about a representative selection of high-level transformation languages, particularly focusing on declarative model and graph transformation languages to clarify their expressiveness. Formal proofs of expressiveness are deemed out of scope, since most languages are not formally specified and it is both infeasible to exhaustively try to combine available features in a variety of distinct languages to check their expressiveness and may also be practically useless if relying on combinations of uncommon constructs. To get a first-hand grasp on these features, I use a selection of the languages to encode realistic transformations and specify the semantics for core high-level transformation language constructs in a way that makes it possible to formally reason about them.
- O2 Validate the semantic preservation for an industrial-scale transformation using translation validation** Having identified key features of high-level transformation languages in O1, this objective focuses on providing first-hand experience with understanding and validating industrial-scale transformation programs.

This serves two purposes: it pinpoints the language constructs used in realistic transformations, providing a clearer direction for objectives O3 and O4, and it provides a deeper understanding of the strength and limitations of translation validation-based approaches.

O3 Develop a symbolic executor that supports identified high-level language features for effective automatic testing of transformations This objective focuses on addressing problem P2 by developing a technique that is able to effectively generate test cases using the source code of input transformation programs.

The approach I take is to model identified key transformation features in a small formal language, and then use symbolic execution to run the transformation generating constraints representing program paths which can be used to generate test cases triggering each explored path. The main challenge here is to accurately track constraints on typed structured input during execution of the available expressive language constructs such as deep traversal in a sound way, while still can be provided to a solver to eliminate infeasible path during execution and generate relevant test cases.

O4 Develop a formal abstract interpreter that supports verifying type and shape properties for high-level transformation languages Using the experiences gathered with the symbolic executor in objective O3, I will go beyond testing of high-level transformations and focus on verification of type and shape properties. I use Schmidt-style abstract interpretation [Schmidt, 1998] to build a sophisticated static analysis that supports verifying type and shape properties for a large subset of a realistic high-level transformation language like Rascal, which both requires handling complex control flow present in these languages—including various control operators, backtracking and exceptions—and intelligently adapting the core idea of memoization—a form of widening on the traces of execution that the traces are finitely representable—to the various looping constructs and recursion while still preserving necessary precision.

1.4 Contributions

Core Contributions

This thesis realizes the stated objectives by the following series of contributions:

- C1** Systematic study fulfilling objective O1 showing that all surveyed declarative model and graph transformation languages, except bi-directional ones, are as expressive as ordinary programming languages (Turing-complete). This is significant because it shows that existing techniques that target less expressive subsets do not cover the complete target languages, and are unlikely to be able to handle more expressive transformations.

Technical Report Ahmad Salim Al-Sibahi. On the computational expressiveness of model transformation languages. *IT University. Technical Report Series*, January 2015. ISSN 1600-6100

- C2** An application of translation validation to a large industrial software modernization transformation written in the high-level language TXL converting an imperative configuration system in C++ to declarative constraints. This provided a deeper understanding of practical use of available high-level transformation constructs fulfilling objective O2.

Conference Paper Alexandru Florin Iosif-Lazăr, Ahmad Salim Al-Sibahi, Aleksandar S. Dimovski, Juha Erik Savolainen, Krzysztof Sierszecki, and Andrzej Wařowski. Experiences from designing and validating a software modernization transformation (E). In *30th IEEE/ACM International Conference on Automated Software Engineering, ASE 2015, Lincoln, NE, USA*, pages 597–607, 2015. doi: 10.1109/ASE.2015.84

- C3** Symbolic execution technique with fully formal semantics supporting high-level transformation language features such as deep type-directed querying, sets, and containment instantiated for a small IMP-like transformation language called TRON. This technique fulfills objective O3 showing that it is possible to extend symbolic execution to support high-level features by state of the art techniques like lazy initialization for pointers and symbolic set constraints with new innovation like deep containment constraints

that provide access to all elements of a particular type reachable for a given value abstracting over the intermediate structure. The technique is compared experimentally of a series of realistic transformations against a baseline black-box test generator optimizing towards literature-recommended metrics, showing that it is possible to achieve a significant improvement in coverage for most evaluation subjects.

Conference Paper Ahmad Salim Al-Sibahi, Aleksandar S. Dimovski, and Andrzej Wasowski. Symbolic execution of high-level transformations. In *Proceedings of the 2016 ACM SIGPLAN International Conference on Software Language Engineering, Amsterdam, The Netherlands*, pages 207–220, 2016. doi: 10.1145/2997364.2997382

Extended Technical Report Ahmad Salim Al-Sibahi, Aleksandar Dimovski, and Andrzej Wasowski. SymexTRON: Symbolic execution of high-level transformation languages: Symbolic execution of high-level transformations. *IT University. Technical Report Series*, September 2016. ISSN 1600-6100

- C4** A formal semantics for a large subset of the Rascal language, called Rascal Light, suitable for developing formal analyses of realistic high-level transformation programs. This represents the first formalization of a large subset of Rascal, which can be useful both for the Rascal community as a way to increase trust in the implementation, and high-level transformation language verification researchers who can use it as a basis for building their verification tools and techniques. The subset supports all relevant high-level features including the traversal construct **visit** with all available strategies, the generalized pattern matching **for**-loop, general looping constructs **while** and **solve**, the expressive pattern matching language with collection matches, deep matching and non-linear matching, control operators including backtracking and rollback, and exceptions. The semantics is based on the available documentation, testing of constructs, and expert contributions from the Rascal developers; I show the soundness of the semantics by proving a series of core theorems.

Technical Report Ahmad Salim Al-Sibahi. The Formal Semantics of Rascal Light. *CoRR*, abs/1703.02312, 2017a. URL <http://arxiv.org/abs/1703.02312>

C5 A Schmidt-style abstract interpreter derived from the formal operational semantics of a large relevant subset of the high-level transformation language Rascal supporting verification of inductive type and shape properties. This contribution fulfills objective O3 contributing theoretical insights including i) a modular domain construction that can be used to abstract individual types of values in languages operation on large domains like Rascal, ii) adaptations of trace memoization (core technique of Schmidt-style abstract interpretation) to various looping constructs, traversals and recursion using available information by each construct to maintain precision, and iii) a representation of state that has a first-class representation of control-flow operations to ensure a close correspondence between abstract and concrete semantic rules. The technique is evaluated on a series of real and realistic transformations and is able to show properties that include that a desugaring correctly translates sugared constructs, that a normalization procedure produces an output in normal form and a rename field refactoring produces programs that do not reference the old field name.

Manuscript Ahmad Salim Al-Sibahi, Aleksandar S. Dimovski, Thomas P. Jensen, and Andrzej Wasowski. Verifying Transformations using Inductive Shape Analysis. 2017¹

Other Scientific Activity

I have contributed to a series of papers on efficiently model checking program families by abstracting variability, which I reference below. The papers are indirectly related to my core objectives, since they use transformations to convert a process model with variability to a small set of process models where the variability has been abstracted away, making them useable with off-the-shelf model checkers.

- **Conference Paper** Aleksandar S. Dimovski, Ahmad Salim Al-Sibahi, Claus Brabrand, and Andrzej Wasowski. Family-based model checking without a family-based model checker. In *Model Checking Software - 22nd International Symposium, SPIN 2015, Stellenbosch, South Africa, August 24-26, 2015, Proceedings*, pages 282–299, 2015a. doi: 10.1007/978-3-319-23404-5_18

¹Author ordering to be determined, currently alphabetical.

- **Journal Paper** Aleksandar S. Dimovski, Ahmad Salim Al-Sibahi, Claus Brabrand, and Andrzej Wasowski. Efficient family-based model checking via variability abstractions. *International Journal on Software Tools for Technology Transfer*, pages 1–19, 2016. ISSN 1433-2787. doi: 10.1007/s10009-016-0425-2. URL <http://dx.doi.org/10.1007/s10009-016-0425-2>

Related to the papers above, I have held an invited talk at the 2015 Workshop on Software Product Line Analysis Tools (co-located with SPLC 2015) on our joint work:

- **Presentation** Aleksandar S. Dimovski, Ahmad Salim Al-Sibahi, Claus Brabrand, and Andrzej Wasowski. Family-based model checking using off-the-shelf model checkers: extended abstract. In *Proceedings of the 19th International Conference on Software Product Line, SPLC 2015, Nashville, TN, USA, July 20-24, 2015*, page 397, 2015b. doi: 10.1145/2791060.2791119. URL <http://doi.acm.org/10.1145/2791060.2791119>

In relation to my work on abstract interpretation of high-level transformations, I presented a poster at the Student Research Competition at POPL 2017. The poster presented work on shape abstractions for TRON, which was used as a stepping stone to C5.

- **Poster** Ahmad Salim Al-Sibahi. Abstract interpretation of high-level transformations. In *Student Research Competition at 44th ACM SIGPLAN Symposium on Principles of Programming Languages 2017, POPL SRC 2017, Paris, France, January 15-21, 2017*, 2017b. Poster

1.5 Overview

This thesis has seven further chapters explaining my stated contributions in detail. Chapter 2 provides background on high-level transformation languages and Chapter 3 summarizes our survey on expressiveness of a wide range of declarative model and graph transformation languages (C1). Chapter 4 summarizes the experiences behind validating a realistic modernization transformation (C2), and Chapter 5 summarizes a formal symbolic execution algorithm with support for high-level transformation language features (C3). The formal semantics of Rascal Light (C4) is presented in Chapter 6, and shape-based abstract interpreter (C5) is presented in Chapter 7. Finally, Chapter 8 concludes my results.

Chapter 2

High-Level Transformation Languages

Taking into consideration the ubiquity of transformations in all aspects of software engineering, it is no surprise that there exist various transformation languages across different communities each specialized to sets of relevant tasks. In this chapter, I summarize these approaches and discuss the available features and what makes them unique.

2.1 Declarative Transformation Languages

Model transformation [Mens and Gorp, 2006; Czarnecki and Helsen, 2006] is one of the most central techniques employed in the field of model-driven engineering. A *model transformation language* allows the programmer to translate elements from one or more source models to elements in one or more target models, given their formal descriptions or *meta-models*.

Declarative model transformation languages are largely rule-based—the programmer specifies how a subgroup of elements in the source model maps to the desired target elements—and can be categorised further based on whether they mainly use graph rewriting as a formalism or not. To give a more concrete grasp on how these languages work, I will use the traditional Families to Persons [Allilaire and Jouault, 2007] transformation as a running example.

The meta-models for both source and target models for the model transformation are shown in Figure 2.1. The source Families meta-model (Figure 2.1a) models a traditional family with a common last name, which contains a mother, father, sons and daughters, each having their own first name; note, that the links are modelled bidirec-

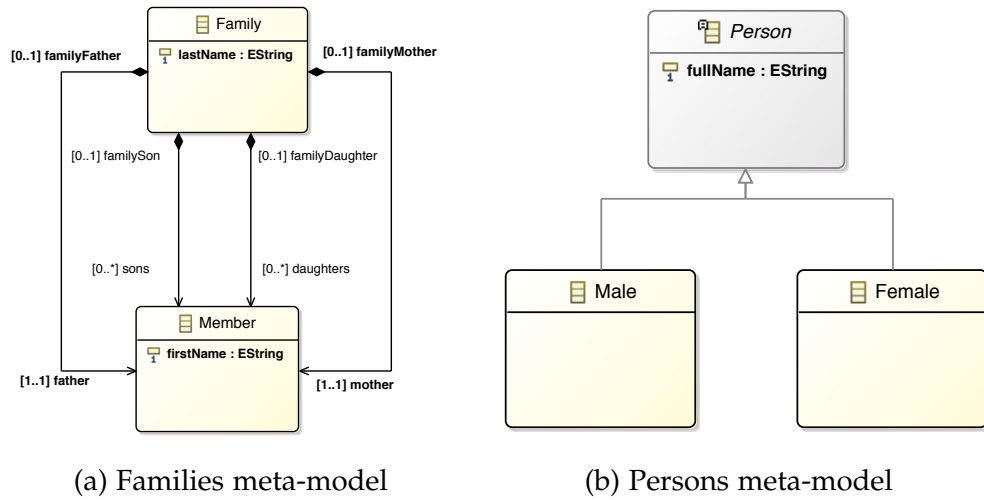


Figure 2.1: Meta-models for the Families to Persons [Allilaire and Jouault, 2007] transformation

tionally and so each member has a reference to the containing family (`familyMother`, `familyFather`, `familySon` or `familyDaughter`). The target Persons meta-model (Figure 2.1b) models individual persons, either male or female, including their full name.

Rule-based Model Transformation Languages

Purely rule-based languages focus on providing a set of declarative rules to describe a transformation between source and target models, given their meta-models. The rules in these languages usually support rich predicate languages like OCL [Warmer and Kleppe, 1998] for constraining the scope of rule application, may allow explicit dependency between rules, and may allow the transformation to be handled imperatively for more precision. Examples of these types of languages include ATL [Jouault et al., 2008], ETL [Kolovos et al., 2008], Tefkat [Lawley and Steel, 2006] and QVT [Object Management Group, 2016].

The example Families to Persons transformation is shown implemented in ETL¹ in Figure 2.2. The transformation consists of two *helper* operations and two transformation rules. The `familyName` operation retrieves the last name for a member using the reverse link to the family that is defined—recall that a member is contained in a family and so

¹Ported and simplified from the original ATL in Allilaire and Jouault [2007]

```

1 @cached
2 operation Families!Member familyName(): String {
3   if (self.familyFather.isDefined())
4     return self.familyFather.lastName;
5   else if (self.familyMother.isDefined())
6     return self.familyMother.lastName;
7   else if (self.familySon.isDefined())
8     return self.familySon.lastName;
9   else return self.familyDaughter.lastName;
10 }
11
12 @cached
13 operation Families!Member isFemale(): Boolean {
14   return self.familyMother.isDefined()
15   or self.familyDaughter.isDefined();
16 }
17
18 rule Member2Female
19 transform member : Families!Member
20 to person : Persons!Female {
21   guard: member.isFemale()
22   person.fullName =
23     member.firstName + " " + member.familyName();
24 }
25
26 rule Member2Male
27 transform member : Families!Member
28 to person : Persons!Male {
29   guard: not member.isFemale()
30   person.fullName =
31     member.firstName + " " + member.familyName();
32 }

```

Figure 2.2: The Families to Persons [Allilaire and Jouault, 2007] transformation in ETL

there must exist exactly one reverse link—and the `isFemale` operation determines whether a member is female (i.e., mother or daughter) or not (i.e., father or son). The `Member2Female` rule converts `Members` to `Female` persons, under the condition specified by **guard** that the input member is female, assigning the output person’s full name using the first and family names from the input member; the `Member2Male` works analogously but converts to `Male` persons assuming the input member is not female. In ETL (also ATL, QVT and similar languages) such rules are scheduled automatically, so `Member2Female` and `Member2Male` are applied on each matching element of type `Member` from the input model of type `Family` potentially producing a collection of `Persons`, either `Male` or `Female`.

Bidirectional Languages Bidirectional-languages [Hidaka et al., 2016] work similarly to other declarative transformation languages except that

transformations must allow conversion back from the target model to the source model. There are various approaches to this, for example requiring reversible rules as in BOTL [Braun and Marschall, 2003] and Triple Graph Grammars [Lauder et al., 2012], the requiring well-behaved operations to extract and recompose data as in Boomerang [Bohannon et al., 2008], or using an explicit synchronization mechanisms as for example in BeanBag [Xiong et al., 2009], SyncATL [Xiong et al., 2007] and X [Hu et al., 2008].

Graph Transformation Languages

Graph-based languages represent the source and target models using graphs—optionally having types or attributes—and describe transformations using sets of conditional graph rewriting rules. The graph rewriting rules specify how subgraphs in the source map to subgraphs in the target, and executing a transformation is done by applying the rewrite rules using some strategy for rule scheduling. Some graph-based languages like PROGRES [Schürr, 1994], GReAT [Balasubramanian et al., 2007], GrGen.NET [Jakumeit et al., 2010], Motif [Syriani and Vangheluwe, 2013], VIATRA2 [Varró and Balogh, 2007], VMTS [Levendovszky et al., 2005], SDM [Fischer et al., 2000] and GROOVE [Rensink, 2003] use a separate imperative control language to let users define rule scheduling strategies, while other languages like AGG [Taentzer, 2010] and Atom3 [de Lara et al., 2004] instead use a fixed strategy that applies rules as long as possible, but allow some form of *layering* that allows ordering of rules.

Two graph rewriting rules in GROOVE for the Families to Persons example transformation are presented in Figure 2.3. The rule in Figure 2.3a adds a label *male* (highlighted using + and is green colored) in case the target member has a `familyFather` link defined and is not already tagged by a *female* or *male* label (negative conditions are highlighted with ! and are red colored). The rule in Figure 2.3b corresponds to the `Member2Male` rule in ETL, and translates a family member tagged with a *male* label to a corresponding new `Male` person—if not already translated as indicated by the negative condition—with the full name being the concatenation of the first name and last name of the target member (operations are defined as subgraphs which rewrite to produce the correct value). This style of transformations is comparatively more visual than the purely rule-based style by ETL, but at the expense of be-

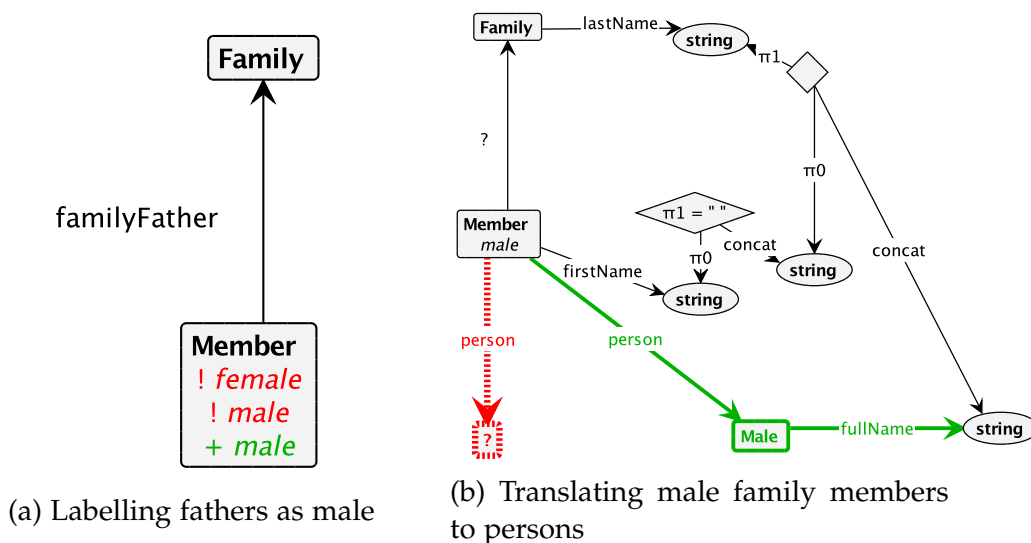


Figure 2.3: Two graph rewriting operations in GROOVE for the Families to Persons transformation

ing more low-level with regards to changes in the graph since additions and deletions of nodes and edges are done more explicitly.

2.2 Program Transformation Languages

The line between model and program transformation languages is fluid, but in general specialized program transformation languages focus more on transformation of syntax—including dealing with concrete syntax and parsing—and try to provide more expressive language features such as supporting higher-order transformations. I will use the Negation Normal Form transformation (NNF, see Figure 2.4) as a running example to present some of the different kinds of program transformation languages. The transformation converts propositional formula (Figure 2.4a)—containing atoms, negation, conjunction, disjunction and implication—to an equivalent formula where all negations are on atoms and without implications (Figure 2.4b).

Rewriting Languages

Rule-based program transformation languages like TXL [Cordy, 2006] are similar to rule-based transformation languages in that they represent a transformation as a set of rules that are scheduled automatically

$$\phi, \psi ::= x \in \text{Atom} \mid \neg\phi \mid \phi \wedge \psi \mid \phi \vee \psi \mid \phi \Rightarrow \psi$$

(a) Formula syntax

$$l ::= x \in \text{Atom} \mid \neg x \quad \phi_{\text{nnf}}, \psi_{\text{nnf}} ::= l \mid \phi_{\text{nnf}} \wedge \psi_{\text{nnf}} \mid \phi_{\text{nnf}} \vee \psi_{\text{nnf}}$$

(b) Normalized Formula syntax

Figure 2.4: Negation Normal Form

during execution; the key difference is that TXL transforms concrete syntax trees instead of structured models—supporting specification of and parsing using concrete grammars as first-class constructs—and supports more expressive features like run-time re-parsing, syntactic substitution and polymorphic rules.

Languages based more directly on term rewriting [Baader and Nipkow, 1998; Cirstea and Kirchner, 2001a,b] focus more on allowing on providing rich composable rewriting using different types of evaluation strategies that include support for backtracking, recursion and term traversal. Languages that primarily follow this paradigm, include programming languages like Stratego/XT [Bravenboer et al., 2008], ELAN [Borovanský et al., 1998] and Maude [Clavel et al., 2007], and frameworks like PLT Redex [Felleisen et al., 2009] for Racket [Felleisen et al., 2015] and Kiama [Sloane, 2009] for Scala [Odersky et al., 2004]. Figure 2.5 presents the NNF transformation in Stratego/XT: formulae are represented as terms using typed constructs, and the transformation `normalize` executes the sequence of rewriting rules tagged with NNF using the outermost strategy, which rewrites terms in a top-down fashion until a fixed-point is reached.

General Program Transformation Languages

General transformation languages like the aforementioned Rascal [Klint et al., 2011], are ordinary programming languages that provide high-level constructs suitable for transformation such as powerful pattern matching, traversal and backtracking. Other examples include languages like Xtend [The Eclipse Foundation, 2014] and frameworks like Uniplate [Mitchell and Runciman, 2007] for Haskell [Marlow, 2010]. Figure 2.6 presents the NNF transformation in Haskell using Uniplate; the transformation uses the higher-order function `rewrite` that traverses input formula exhaustively applying function `nnf'`, which is an ordinary

```

1 module nnf
2 signature
3 constructors
4   Atom      : ID -> Formula
5   Neg       : Formula -> Formula
6   And       : Formula * Formula -> Formula
7   Or        : Formula * Formula -> Formula
8   Imp       : Formula * Formula -> Formula
9 strategies
10  normalize = outermost(NNF)
11 rules
12  NNF : Neg(Or(e1, e2)) -> And(Neg(e1), Neg(e2))
13  NNF : Neg(And(e1, e2)) -> Or(Neg(e1), Neg(e2))
14  NNF : Neg(Neg(e)) -> e
15  NNF : Neg(Imp(e1, e2)) -> And(e1, Neg(e2))
16  NNF : Imp(e1, e2) -> Or(Neg(e1), e2)

```

Figure 2.5: Negation Normal Form Transformation in Stratego/XT

Haskell function defined by pattern matching that produces a possible result value for each applicable case.

```

1 data Formula = FAtom String
2             | (:) Formula
3             | Formula & Formula
4             | Formula | Formula
5             | Formula => Formula
6
7 nnf :: Formula -> Formula
8 nnf = rewrite nnf'
9   where nnf' ((:) phi) = Just phi
10         nnf' ((:) (phi & psi)) = Just ((:) phi | ((:) psi))
11         nnf' ((:) (phi | psi)) = Just ((:) phi & ((:) psi)
12         nnf' (phi => psi) = Just ((:) phi | psi)
13         nnf' _ = Nothing

```

Figure 2.6: Negation Normal Form Transformation in Haskell using Uniplate

2.3 Programs and Paradigms

In principle, any programming language is a transformation language, but I have in this chapter presented languages across different paradigms which have features designed primarily towards transformation of models and programs. In the chapters to come, I will focus on the programming-oriented high-level transformation languages,

because these allow expressing the more interesting transformations—refactorings, normalization procedures, optimizations, translations—and align the best with my goal of adapting techniques from programming language verification. Furthermore, as we will see in Chapter 3 there is little to gain of choosing the more declarative model and graph from a verification perspective.

In particular, later in this dissertation I describe an effort validating an industrial transformation in the mentioned TXL language (Chapter 4), which provided a particularly good match for the target code base because of its powerful capabilities for manipulating C++ syntax. To get a good grasp of the verification challenges introduced by the high-level transformation features, I develop a symbolic execution algorithm (Chapter 5) for a small purposely defined formal transformation language called TRON. Finally, I scale up the verification effort focusing on a large expressive subset of the Rascal transformation languages, for which I develop a verified formal semantics (Chapter 6) and an abstract interpretation based static analysis tool (Chapter 7). Rascal is particularly suitable as a target system because it provides a representative selection of high-level transformation language features—traversal, backtracking, fixed-point iteration—modelled explicitly, has an available up-to-date open source implementation,² and plenty of realistic programs that can be used for evaluation.

²<https://github.com/usethesource/Rascal>

Chapter 3

Expressiveness of Declarative Model Transformation Languages^{*}

The declarative model transformation languages presented in Chapter 2 provide a convenient rule-based interface for translating between various models. It is however unclear that the computational capabilities [Taylor, 1998] of such languages matches those of general programming languages, since modification of state can be limited and repetition can be implicit or bounded. The computational expressiveness of a language matters from a verification point of view, since a large variety of techniques only apply on languages with limited computational expressiveness. This chapter presents a systematic analysis of a wide selection of declarative model transformation languages, seeking to pinpoint the constructs that enables or limits the level of expressiveness.

3.1 Surveys on Model Transformations

Language Expressiveness To the author's knowledge there has been no overall study on computational expressiveness of model transformation languages, but there do exist several studies that compare specific aspects. Gomes et al. [2014] compare a variety of languages on their pattern matching capabilities in order to clarify the performance implications of each system¹. Different language features of ATL and SDM are

¹Additionally, a source of inspiration for selection of transformation languages.

^{*}This chapter is based on Al-Sibahi [2015], which is published as a technical report.

compared in a case study [Patzina and Patzina, 2012], presenting a table clarifying the capabilities of these languages, including support for recursion and in-place transformation. Finally, several model transformation language implementation papers [Varró and Balogh, 2007; Syriani and Vangheluwe, 2013] present small comparisons of language features; however, these do not provide an analysis of language expressiveness, and are neither comprehensive in their comparisons.

General Mens and Gorp [2006] provide definitions for kinds of model transformations, and classifies them according to several characteristics including whether they transform across different abstraction levels, between different meta-model definitions and are primarily syntactic or semantic. Czarnecki and Helsen [2006] provide a detailed classification of a large selection of model transformation languages according to the set of available features including constraint and transformation rule specifications, level of control over execution and modularity; Hidaka et al. [2016] provides a similar feature-oriented classification for bidirectional transformations. Several surveys [Amrani et al., 2012; Calegari and Szasz, 2013; Rahim and Whittle, 2015] focus on summarizing the state of the art in validation and verification of model transformation, classifying according to the techniques used, transformation features and properties supported. All of these surveys do not address my objective (O1, Section 1.3) which is to identify the computational expressiveness, except Hidaka et al. [2016] which does so only for bidirectional languages.

3.2 Computational expressiveness

Computability is usually defined in terms of the Church-Turing thesis [Kleene, 1967; Copeland, 2015] which states that any *effectively computable* mathematical function can be computed by a Turing machine [Hopcroft et al., 2007]. A language is said to be Turing-complete if it is as expressive as a Turing machine, in the sense that it can compute the same any function a Turing machine can compute. Known examples of Turing-equivalent systems include the λ -calculus [Barendt, 1985; Alama, 2014] and term-rewriting systems [Post, 1947].

While it can be hard to pinpoint the combination of features needed for a language to be Turing-complete, we can specify some general characteristics:

State change It should be possible to continue computation with a new state, whether by writing on the tape as the Turing machines, by reduction as in the λ -calculus or by subterm replacement as in term rewriting.

Branching It should be possible to make different decisions on input, similarly to how Turing-machines conditionally move the head depending on the current symbol or how term-rewriting systems express rules conditional on patterns.

Unbounded execution time The language must not bound the number of execution steps, e.g., by providing repetition constructs such as recursion or iteration.

Unbounded memory There must be no bounds on memory, analogously to the infinite extensibility of a tape in Turing-machines and the unbounded term size in λ -calculus.

Determining the computational expressiveness for a language has been long known to be important with regards to verification. In particular, it determines the class of interesting semantic properties we can effectively decide [Rice, 1953] and thus also affects the verification techniques applicable.

General results on graph rewriting

Graph rewriting is the base formalism for many declarative model transformation languages, as mentioned in Chapter 2. Therefore, it is important to understand the expressiveness of general graph rewriting systems, and clarify the conditions necessary to achieve expected expressiveness results.

Proposition 3.1. *Graph rewriting systems can express Turing-complete computation.*

Proof. Shown by Plump [1998]. □

Existing research shows that graph rewriting systems can express Turing-complete computation. Albeit similar, graph rewriting systems are richer than term rewriting systems—possibly having cycles and shared nodes—and do not share the same termination criteria; Plump [1998] shows it is possible to translate a non-terminating term rewriting

systems to a terminating graph rewriting systems and vice versa. Generally, a base requirement for graph rewriting systems to be terminating is that all the rules contained must either delete nodes or have application conditions [Plump, 1995; Bisztray and Heckel, 2010]. Any other graph rewriting system is trivially non-terminating since the same rule can be reapplied *ad infinitum*.

3.3 Analysis Method

The formal process to (dis)prove that a language is Turing-complete is by reduction, so to show that it is (im)possible to encode another Turing-equivalent system or an undecidable problem—e.g., Post correspondence problem [Post, 1946]—in the language. It is however practically infeasible to do such a formal proof for each of the various languages examined, since it requires manually setting up each system, learning the language sufficiently to perform a complete analysis, and verifying that such analysis is semantically correct.

Instead, I perform a literature-based analysis whereby I use published articles, language manuals and other related documentation to check whether there is a combinations of features that make a language Turing-complete. I will use the (semi-)formal semantics of target language as a primary source if provided, but in most cases I use a combination of sources to establish an overview of the target language capabilities. The results provided should be sufficiently accurate, and reflect better the features that ordinary users meet.

Collection process

The collection process proceeded as follows:

- If the language documentation explicitly mentions or proves that the language is Turing-complete, I note down the result and language features required to achieve such expressiveness.
- Otherwise, I search the documentation for whether there is a combination of constructs that make the target language Turing-complete.

The constructs required for Turing-complete computation differ across languages examined. For the subset of languages that directly

allow use of imperative constructs, it is straightforward to check the expressiveness. If the target language supports creation and processing of variably-sized data, conditionals and unbounded iteration or recursion, then it is definitely Turing-complete.

For primarily rule-based transformation languages with implicit control flow, the problem is more subtle. The analysis must consider possible limitations of rule definition and execution, and check whether non-terminating rules are rejected or iteration is only allowed on bounded collections.

3.4 Analysis of computational expressiveness

Paradigms The investigated languages are summarized according to supported paradigms in Table 3.1. All investigated declarative model transformation languages are unsurprisingly rule-based, except the operational subset of QVT and ATC which is a low-level implementation oriented language. The investigated languages are divided on the underlying formalism, with more than half the languages being primarily graph-oriented and the others being oriented towards meta-models. Furthermore, I have chosen to include two primarily bidirectional languages—BeanBag and BOTL—for comparison, although a more comprehensive analysis of expressiveness of bidirectional languages is available in Hidaka et al. [2016]. Finally, Table 3.1 shows a division regarding imperative constructs such as state updates and loops, spreading evenly between languages that do and do not support such constructs.

Expressiveness The expressiveness of investigated languages is presented in Table 3.2, which summarizes available repetition constructs and whether each language is Turing-complete. The repetition constructs supported are varied across different languages, but I divided them in three main categories according to whether the investigated languages support explicit loops like `for` and `while`, whether they support some form of recursion and whether they have implicit repetition in the execution engine, so that execution of the specified rule set continues until a fixed-point is reached.

All presented languages except the bidirectional ones are found to be Turing-complete based on the detailed analysis in Section 3.4; this is

consistent with the results by Hidaka et al. [2016] which argue that most bidirectional languages are by construction non-Turing complete, since they enforce functional relationships. Two of the presented languages (AGG and VMTS) notably support checking termination of transformations, but these checks are optional and not enforced.

Detailed analysis

This section discusses the details of the results presented in Table 3.1 and Table 3.2. It provides a summary for each language, whether the language is Turing-complete and what combination of features contribute to the expressiveness.

Language	Rule-based			Imperative
	General	Graph-oriented	Bidirectional	
PROGReS	✓	✓	✗	✓
AGG	✓	✓	✗	✗
GReAT	✓	✓	✗	✗
GrGen.NET	✓	✓	✗	✓
Motif	✓	✓	✗	✗
Atom3	✓	✓	✗	A01
Tefkat	✓	✗	✗	✗
QVT Rel.	✓	✗	✓	QR1
QVT Op.	✗	✗	✗	✓
BOTL	✓	✗	✓	✗
BeanBag	✓	✗	✓	✗
VIATRA2	✓	✓	✗	✓
VMTS	✓	✓	✗	VM1
ATL	✓	✗	✗	✓
ETL	✓	✗	✗	✓
SDM	✓	✓	✗	✓
ATC	✗	✗	✗	✓

A01 Constraints are specified by python statements

QR1 Can call QVT Operational mappings

VM1 Transformation on attributes using XSLT

Table 3.1: Paradigms of transformation languages

PROGReS The language supports specification of arbitrary graph rewriting rules, whose execution is controlled by an imperative language [Schürr, 1994; Patzina and Patzina, 2012; Rozenberg, 1997]. The language is Turing-complete by Proposition 3.1 since the imperative control subset has an unbounded looping construct `loop` [The PROGReS Developer Team, 1999], and it can therefore simulate a general graph rewriting system.

AGG Arbitrary graph rewriting rules are supported in AGG [Taentzer, 2004; Runge, 2006], and the rule interpretation engine only terminates

Language	Repetition			Turing-complete
	Looping	Recursion	Implicit	
PROGReS	✓	✓	✗	✓
AGG	✗	✗	✓	✓ !
GReAT	✗	✓	✗	✓
GrGen.NET	✓	✓	✗	✓
Motif	✓	✓	✗	✓
Atom3	✗	✗	✓	✓
Tefkat	✗	✓	✓	✓
QVT Rel.	✗	✓	✓	✓
QVT Op.	Ⓜ	✓	✗	✓
BOTL	✗	✗	Ⓜ	✗
BeanBag	Ⓜ	✓	✗	?
VIATRA2	✓	✓	✗	✓
VMTS	✗	✓	✗	✓ !
ATL	Ⓜ	AT1	Ⓜ	✓
ETL	ET1	ET1	Ⓜ	✓
SDM	SD1	SD1	✗	✓
ATC	✓	✓	✗	✓

Ⓜ Bounded number of steps

? Status generally unknown

! Optional termination check

AT1 Using global helpers or lazy rules

ET1 Using Epsilon Object Language (EOL)

SD1 By setting up the required control flow and path expressions

Table 3.2: Expressiveness of transformation languages

when no rule is further applicable. The language can check whether a given rule-set of a program is within the terminating fragment, but since this check is not enforced the language is considered Turing-complete by Proposition 3.1.

GReAT The language [Balasubramanian et al., 2007; Agrawal et al., 2003] supports defining transformation rules on graphs modularly as blocks that can be wired together, even recursively where the output of a block can be wired as input to a block further back in the transformation sequence. Since a rule sequence is only terminating when no further output is produced, the language is Turing-complete.

GrGen.NET The language user manual [Blumer et al., 2010] directly shows an implementation of a Turing-machine. Notably, the language works on graphs, supports recursive rules [Jakumeit, 2008], unbounded looping and general imperative constructs.

Motif The language [Syriani and Vangheluwe, 2013] transforms graphs and supports unbounded looping (FRule, SRule), backtracking and recursion (XRule), and it explicitly mentioned that it is possible to construct non-terminating programs. Therefore, the language can simulate a general graph rewriting system and is thus Turing-complete by Proposition 3.1.

Atom3 The framework uses graph-based meta-modelling [Vangheluwe et al., 2002], and transformations are specified using graph rewriting rules. The execution engine (Graph Rewriting Processor) [Levytskyy and Kerckhoffs, 2003; de Lara et al., 2004] is not limited and applies rules to exhaustion, making the language Turing-complete per Proposition 3.1. Remarkably, Motif is implemented in Atom3 which further strengthens the argument for Turing-completeness.

Tefkat This transformation language [Lawley and Steel, 2006] is based primarily on the second revised submission report on QVT [Duddy et al., 2004], which explicitly touts Turing-completeness. Transformation rules can specify expressive constraints on *both* source and target models, and it supports unbounded recursion.

QVT Relations The language [Object Management Group, 2016] is Turing-complete since it can simulate general graph rewriting systems by in-place transformations, where transformation rules (relations) are re-applied until all required constraints are satisfied. Additionally, it is possible further call imperative rules from the operational subset and arbitrary external code using the QVT Blackbox mechanism.

QVT Operational The operational subset of QVT [Object Management Group, 2016] is Turing-complete since it is possible to specify recursive rules [Kraas, 2014], in combination with creation of intermediate data, branching and looping. This makes it similar to many general-purpose imperative programming languages.

BOTL Model transformations in BOTL [Braun and Marschall, 2003] are bounded by the number of rules defined and only applied to finite sets of matches. Therefore, all transformations are known to be terminating and the language is not Turing-complete.

BeanBag The language [Xiong et al., 2009] supports synchronisation of two models by using expressions containing various equational constraints, including variable binding, bounded iteration (`forall`, `exists`) and recursion. All the repetition constructs must however satisfy a stability property—ensuring there always exists a valid synchronisation for given models—which limits the expressiveness of these constructs; the only way to construct non-terminating programs in the language is by using counter-intuitive circular equational constraints, and most program in practice always terminate. It is unclear whether it is possible to use circular constraints intuitively to perform arbitrary computation without breaking the stability property, and without further formal analysis it can not be decided whether the language is Turing-complete or not². This formal analysis is left as future work, but the paper [Xiong et al., 2009] presents a formal semantics which could be used as a starting point.

VIATRA2 The language [Varró and Balogh, 2007; Balogh and Varró, 2006] is based on two formalisms: graph transformation (GT) which

²Confirmed by personal correspondence with the main author, Dr. Yingfei Xiong.

allows the user to specify graph rewriting rules and abstract state machine (ASM) [Gurevich, 1995; Börger and Stärk, 2003]—a known Turing-complete formalism—that can be used to control the flow of execution. Notably, it is also possible to simulate a general graph rewriting system by using GT and the unbounded iterate construct from ASM, and so Proposition 3.1 applies.

VMTS This language [Levendovszky et al., 2005] is Turing-complete by Proposition 3.1 since it supports graph rewriting using a combination of visual model processors (VMP) and the visual control flow language (VCFL) [Lengyel et al., 2007]. Attribute transformations are performed using XSLT [Kay, 2017], which is known to be Turing-complete as well [Onder and Bayram, 2006].

ATL Ordinary ATL rules [Jouault et al., 2008; Troya and Vallecillo, 2011] cannot be recursive and are applied only a finite amount of times by the execution engine [Wagelaar et al., 2014]. Lazy rules [Tisi et al., 2011] can however be recursively called, which in combination with the complex constraints on source and target models makes the declarative part of ATL Turing-complete, similarly to Tefkat.

It is possible to call imperative OCL helpers from ATL rules, but all imperative constructs are bounded on the size of collections [Cengarle and Knapp, 2004]. OCL does however support recursion in helpers, and so Turing-complete computation could still be done that way.

ETL Epsilon Transformation Language [Kolovos et al., 2008, 2014] supports both ordinary and lazy rules, making the declarative part Turing-complete like ATL. Furthermore, it relies allows use of imperative constructs from the Epsilon Object Language (EOL), which includes unbounded loops (both for and while), branching and assignment.

SDM Story Driven Modelling (SDM) [Fischer et al., 2000] uses graph transformation as an underlying mechanism, where *story patterns* are used for describing modifications to models, similarly to graph rewriting rules. The patterns are executed via a control flow language that allows unbounded recursion [von Detten et al., 2012; Patzina and Patzina, 2012], and is therefore capable of expressing Turing-complete computation.

ATC As a low-level language [Estévez et al., 2006] for implementing model transformation engines, ATC supports many necessary primitives for model querying, matching, manipulation and transformation. In addition it contains common imperative language constructs, including unbounded while loops and is therefore Turing-complete.

3.5 Expressiveness and Verification

Eventhough the investigated languages focused more on providing a high-level declarative constructs for transforming models, the analysis presented in this chapter showed that virtually all the presented languages are very expressive and able to perform Turing-complete computation; the only notable exceptions were the bidirectional languages, BOTL and BeanBag, which were more limited by design.

This seems to indicate that transformation is a task that inherently requires a high level of expressiveness. Even a languages like ATL which typically bounds rule application and iteration, allows recursion via lazy rules and OCL helpers, and languages with termination checkers like AGG and VMTS chose to not strictly enforce termination.

These findings suggest that we can not treat declarative model transformation languages specially with regards to verification when we are interested in supporting the full feature set available. In particular, existing techniques that only support limited expressive subsets of transformation languages will only work in very specific cases, and which avoid features that are useful for more interesting transformations like refactoring.

General purpose high-level transformation languages like TXL and Rascal are known to also be Turing-complete, but provide several advantages from a verification point of view. Their data structures and basic control operations are similar to ordinary programming languages, so it is possible to more directly apply verification techniques such as symbolic execution and abstract interpretation from these communities. Furthermore, while they similarly to many of the declarative transformation languages, support high-level transformation features—backtracking, traversal and fixed-point iteration—they in contrast use an explicit execution model which is easier to reason about statically.

Chapter 4

Validating a Software Modernization Transformation^{*}

Joint work with: Alexandru F. Iosif-Lazăr, Aleksandar S. Dimovski, Juha Erik Savolainen, Krzysztof Sierszecki, Andrzej Wąsowski

Software modernization [Seacord et al., 2003] aims to tackle the challenge of growing size and complexity of legacy code by restructuring it to use improved technologies while preserving core functionality. To date, there is relatively little literature available on software modernization projects in practice, especially in the safety critical domain. We present experiences from an industrial software modernization project that is done in collaboration between IT University of Copenhagen and Danfoss Power Electronics,¹ a global producer of components and solutions for controlling electric motors. The modernization project is automated using a non-trivial code-to-model transformation written using the high-level rule-based transformation language TXL [Cordy, 2006], and has been particularly important for shedding light on the problems that occur in transformations, and the necessity of developing formal verification methods and tools.

¹<http://www.danfoss.com>

^{*}This chapter is based on Iosif-Lazăr et al. [2015] which is published a peer-reviewed conference publication. The primary experiments, which concern validating the transformation, were performed by the main author Alexandru F. Iosif-Lazăr based on an implementation that is done in collaboration between a separate team at ITU (Rolf-Helge Pfeiffer, Aleksandar Dimovski and Andrzej Wąsowski) and Danfoss. I participated in studying and classifying bugs in the transformation and contributed to the observations and formal justification in the paper. The gathered experience from this paper had an inspirational role for the rest of the dissertation.

Concretely, the modernization project involves a *configuration* tool used to adapt Danfoss frequency converters to a particular application. The configuration tool consists of thousands of C++ functions for accessing and validating the configuration parameters. The code base is being modernized so each configuration parameter function must be converted from the imperative C++ implementation to a declarative logical specification, so that the configuration software can be managed by an off-the-shelf constraint solver. Each existing function computes a result—either the value of a particular parameter or a Boolean value representing a check of accessibility and validity of parameters—and the expected declarative specification should be an expression or logical constraint that preserves the same result. The code modernization is done the automated transformation on each function individually; the transformation performs a sequence of syntactic replacements, gradually eliminating C++ pre-processor directives, local variables, C++ control statements and leaving behind a pure (side-effect free) expression.

TXL only ensures the syntactic correctness of transformations—enforcing grammar constraints through pattern matching—and does not guarantee that preservation of source program semantics. In the Danfoss case, determining the semantic equivalence is feasible due to the particular shape of programs: the employed subset of C++ is small with no recursion, small fixed bounds for loops, bounded input and no use of esoteric features like inline assembly.

We assessed that symbolic execution [King, 1976] is mature enough to handle this task efficiently, and implemented a *lightweight* wrapper that compiles both input and output of the transformation and employs the symbolic executor KLEE [Cadar et al., 2008] to assert their equivalence for all possible input. As the result of symbolically executing the programs, KLEE produces a set of *path conditions* that represent distinct execution paths of the two programs along with their return values as symbolic formulae; if the equivalence assertion is satisfied by all paths then the two programs are equivalent, and otherwise if there is a pair of compatible input-output path conditions where the assertion fails, then it outputs the path condition pair and a counter-example containing a set of concrete assignments to inputs that trigger the assertion failure.

Our contributions are:

- Synthesis of experiences from the design of a non-trivial modernization transformation for an industrial project.

- Designing and implementing a transformation validation technique for this case study that produces concrete counter-examples when semantic equivalence is not preserved between input and output programs.
- Lessons learnt from using the validation technique in the case study, including an analysis of the kind of errors that have been identified.

To the best of our knowledge this is the first analysis of transformation errors extracted from an industrial project of this scale and complexity (4119 functions or 14502 lines of input code of which: 105 error cases were found at transformation time, 104 error cases were found at verification time and 3910 were successfully verified).

4.1 Motivating Example

We illustrate the technique for modernizing a function and validating its correctness by using an example containing a simplified and anonymized version of an actual function from the Danfoss code base.

```

1 Configuration config = selectedConfParameter;
2 Option opt = selectedOptParameter;
3 bool result = false;
4 switch (config) {
5     case config1:
6         if (opt == option1) result = true;
7         break;
8     default:
9         result = true;
10        break;
11 }
12 return result;

```

(a) Example imperative parameter function

<pre> 1 (selectedConfParameter == config1 2 && selectedOptParameter == option1) 3 ? true 4 : false </pre>	<pre> 1 (selectedConfParameter == config1) 2 ? (selectedOptParameter == option1) 3 : true </pre>
--	--

(b) Actual output constraint from modernization transformation

(c) The expected output constraint

Figure 4.1: Anonymized example validation function from the Danfoss code base

The input program (Fig. 4.1a) is a function that validates the consistency of the values provided to two selected parameters (`selectedConfParameter` and `selectedOptParameter`) with respect to each other using constants `config1` and `option1`. The goal of our modernization transformation is to produce a declarative constraint that for the same inputs produces the same truth values.

Syntactically, the input program consists of a sequence of variable initializations, a switch statement with two cases—one containing an if-statement—and a return statement that returns the value stored in the variable `result`. The corresponding output program should be a pure C++ expression which produces the same result as the input given the same values. Running the TXL modernization transformation produces a declarative C++ expression (Fig. 4.1b) using the following steps:

1. Replacing the switch-statement with a nested if-else conditional.
2. Simplifying the conditional statement, so local variables are replaced by their assigned values and we end up with a straightforward control flow.
3. Reducing the conditional statement a ternary expression of the form `e1 ? e2 : e3`.

It is not immediately clear that the produced output (Fig. 4.1b) has equivalent semantics to the input (Fig. 4.1a), despite this being a relatively simple example. In order to check that this is indeed the case, we perform symbolic execution on the input and output programs using KLEE to gather path conditions for all possible execution paths of each program; when there exists discrepancy between the path conditions of the input and output programs, KLEE reports the discrepant path as a counter-example that acts as a witness of an execution that is possible in one program but not in the other.

For example, the path condition

$$\begin{aligned} & \textit{selectedConfigParameter} \neq \textit{config}_1 \\ & \wedge \textit{selectedOptParameter} = \textit{option}_1 \\ & \wedge \textit{result} = \textit{true} \end{aligned}$$

is generated for the input program in Fig. 4.1a but not for the transformed program in Fig. 4.1b proving that these programs are not semantically equivalent. By further investigation of the counter-example, we

are able to determine that the expected output program should instead be the one shown in Fig. 4.1c.

Furthermore, we were able to diagnose the trace of transformation rules applied to the input program successfully finding the rule that produced erroneous output. The erroneous rule was a rule which tried to simplify nested `if`-statements forgetting to take the `else`-branch into consideration. The validation technique thus proved itself incredibly useful in accurately tracking where the bugs occurred, allowing an experienced transformation specialist to fix this issue relatively swiftly.

4.2 The Modernization Case Study

The goal of the case study is to establish the feasibility of transforming a large imperative C++ code base for a configuration tool to a declarative model, in a manner that is *automatic*, *trustworthy* and *cost effective*. Cost effectiveness is understood here as being cheaper than reimplementing the code from scratch.

The modernization project combines transformation engineering and verification research in an exploratory case study, which acts as a pilot for larger modernization activities in the same organization. The researchers have access to the legacy C++ code base and to three Danfoss engineers/architects knowledgeable about the code, the context and the use case.

The study has two key propositions: i) To establish that freely available transformation and validation tools are sufficiently mature to execute this modernization process, and ii) to explain what kind of errors might appear in transformation projects involving complex code, even if implemented by experienced model transformation developers and language specialists.

Case Description

The configuration tool code base consists of 4119 C++ functions in total. These functions encode dependencies, visibility constraints, and default values of approximately one thousand configuration parameters of a frequency converter. There is limited use of C++ features: no object-oriented aspects are used, except for member access and limited encapsulation. Functions have straightforward control flow mostly consisting of conditionals and `switch`-statements, with occasional use of bounded

iteration via **for**-loops; there is no use of **goto** statements, unbounded loops or recursion. Other constructs include variable declarations and usage of local variables in arithmetic and comparison expressions, calls to pure functions, singleton members and static members (e.g., to convert physical units), and casts between different types (both C-style casts and `static_casts`).

There are 14502 source lines of code (SLOC) that need to be modernized in the pilot project, and more similarly-looking configuration tools for other products waiting for modernization afterwards. As many as 3348 of the 4119 functions are already in expression form; these do not need to be modernized, but should left intact by the modernization transformation. The remaining 771 functions have 14.47 SLOC of code on average.

The modernization to declarative logical expressions decreases the burden of maintenance by allowing configuration to be done using an off-the-shelf verifier [Felfernig et al., 2014]. The nature of the configuration project puts some strict constraints on the implementation and quality of the modernization transformation:

Automation It is infeasible to halt the normal development of the project for an extensive period of time . An automated transformation can be developed in parallel to the actively developed code base and can be efficiently run again on the updated code base in minutes, ensuring minimal impact.

Trustworthiness The configuration code contains crucial domain information and missing configuration constraints could lead to creation of unsafe configurations by customers.

Verification Background

In the general case, it is undecidable to show semantic preservation for transformations, but in practice many analysis problems appearing in engineering of real systems can be handled using incomplete and (partly) unsound [Livshits et al., 2015] methods. This is true for our particular case where most inputs have a fixed shape, and it is perfectly acceptable that there are few corner cases are manually handled during modernization.

Our approach is aligned with translation validation [Samet, 1976; Pnueli et al., 1998] in the sense that it validates concrete translations instead of the transformation tool or algorithm itself. The path conditions

produced by KLEE represent the semantic framework of the compared subset of C++ programs, and an SMT solver is used by KLEE to prove the equivalence of the path conditions, providing a counter example if they are not equivalent.

Research Questions and Methods

Our methodology generally follows the framework of action research [Wohlin et al., 2012]. Our research questions are as follows:

- RQ1** Is it feasible to design the aforementioned transformation using off-the-shelf technology in a limited time?
- RQ2** What are the main obstacles and challenges in designing and implementing the transformation?
- RQ3** Can high assurance methods be used at acceptable costs to validate the transformation?
- RQ4** What kind of errors are found in a transformation implemented by experts?

Questions RQ1–3 are interesting for companies and researchers looking into technology transfer in modernization projects, and the last research question is more relevant for researchers in the fields of model and program transformations. We are not aware of existing studies that provide a detailed account on the number and type of errors in realistic software transformations. This work hypothesizes the kinds of problems worth addressing, and concretely inspired the verification efforts presented in the later chapters of this thesis.

We decided to address RQ1 and RQ2 by systematically investigating the most effective way to implement the transformation, and evaluating a variety of approaches and technologies that fulfill the requirements elicited from the industry partner. A similar process was executed for investigating validation technologies to solve RQ3, where we recorded experiences during the process to report them in this paper. We have collected statistics about the effectiveness and efficiency of the transformation for RQ2, and measure the ration of false positives for RQ3, explaining why they appear when following our validation methods. Regarding RQ4, we collected counterexamples from the validation process, classified them, and qualitatively analyzed them to understand what kind of errors arise in the transformation.

Study Participants Three teams were involved in the project:

1. The industrial partner had a team of two engineers and an architect, which presented the software modernization problem, requirements and artifacts.
2. The transformation implementation team consisted of a transformation expert, a language semantics expert and a project leader; the transformation expert designed and implemented the transformation in dialog with the other members of the team.
3. The validation team consisted of an applied verification researcher, a junior PhD student in programming languages,² and the same language semantics expert and same project leader that were involved in designing the transformation.

Threats to Validity

Construct Validity The case was selected by the industrial partner according to the problem they wanted to solve, and the research team only had access to the previously mentioned parts of the configuration tool. It is entirely possible that the researchers could have misunderstood particular aspects of the software architecture, but given their expertise in the relevant areas, we believe that the impact of this would only be limited. Furthermore, the results are still interesting from a research point of view since transformation is of substantial complexity, and even if it was slightly misaligned with the requirements it still sheds a lot of light on pragmatics of such transformations. From the engineering point of view, the validation method provides actual counterexample for the detected errors and is easy to manually confirm any unwanted discrepancy.

Internal Validity The validation procedure has been developed in the same study in which it is evaluated, and there is certainly a risk that it had overlooked some important errors. However, the transformation and validation projects has been designed independently, and key developers of the two parts have not communicated in any significant matter; several months have passed between the end of the transformation implementation project and the beginning of the validation project. In

²Me

particular, the transformation had been designed with a focus on effectiveness with no thought of later verification in mind. The only minor risk of misinterpretation could be in the explanation of the programming errors that caused each type of bug, since the mapping was done by the validation team.

External Validity Limited external validity is in the very nature of an individual study, which is why we take care to describe the properties of the case in detail. In safety critical software, simple imperative code with bounded for-loops is common, and we thus believe that the findings can generalize to additional modernization projects.

Reliability The analyzed transformation errors can be biased towards the weaknesses of the particular development team. It is unlikely that the reported errors were trivial, since the involved teams were experts in model transformation and verification with more than 12 years of combined experience. Nevertheless, we report these findings only existentially—without generalizing—providing a single data point for the research space where very little evidence is available so far. There is a definite need for more studies to get a good understanding of the nature of design errors in transformative and generative programming.

4.3 Designing the Transformation

We found that implementing an automatic transformation is superior to a manual refactoring, as it allows minimizing down-time on the main development of the code base. The code to be modernized only needs to be frozen for the few minutes it takes to run the transformation and it can be evolved freely while the transformation is being implemented and tested. This is somewhat different from the standard use case for transformations which is automating repetitive tasks in conversions of data, code or models. This transformation was meant to be executed only once, but automation was key to minimize disruptions in the regular development process.

Observation 4.1. Automatic transformations allowed us to decouple and parallelize the regular development and the modernization activities.

Constructing a general transformation from imperative C++ code into a declarative form is impossible in general, but our objective was

much more modest: we only needed to handle programs with a fixed shape that covers the code at hand. We settled on saving resources whenever possible by sacrificing generality and even agreed to leave a small number of cases with complex fragments involving loops to be modernized manually instead of designing rules that would handle them correctly from first principles. The manual migrations can be hard-coded into a transformation as special cases, so that the transformation execution remains automatic.

Observation 4.2. Modernization is a one-time transformation, so it was beneficial to focus only at the code at hand without eagerly generalizing.

To control the lack of generality, the implementation follows a *fail-fast* programming style [Gray, 1986; Shore, 2004]: it succeeds on the inputs that it was designed for, but fails as early as possible on inputs that violate specific assumptions required (e.g., use of unsupported language elements). This was achieved by making preconditions for rules as precise as possible and writing explicit assertions when possible (where a rule is in principle applicable but does not cover all the cases). Without fail-fast programming we would have very limited trust that the transformation works correctly on hundreds of input code fragments that we were not able to inspect manually.

Observation 4.3. Fail-fast programming helped retaining quality on required inputs, safely avoiding supporting anything not necessary for the modernization project.

Initially, we considered using semantic analysis techniques to solve the task—e.g., static single assignment or type analysis—but it soon became clear that this would considerably raise the complexity of the implementation, and more importantly pollute the output with automatically generated identifiers that are unfamiliar to developers. The resulting models are not merely left as is, but are expected to be read and modified by humans afterwards. Full understanding of static semantics might only be required if one implements a general refactoring, but for known code base it seems much easier to work with syntactic transformations.

Observation 4.4. We found working with syntax directly significantly more cost effective for ad-hoc transformation tasks than relying on semantic information.

Tool Selection

Since the input language (C++) is rather complex, we understood early on that the transformation tooling should be driven by the availability of a good C++ grammar, rather than our personal experiences. Thus we considered several existing open source compiler front-ends (for instance GCC³), and language tools (such as Eclipse CDT⁴) for the task, but found them too challenging to use. We then considered transformation tools (see also Chapter 2) and found that both Spoofox [Kats and Visser, 2010] and TXL [Cordy, 2006] have C++ grammars, and that the latter had a better maintained grammar for the current version of the tool. We thus ended up using TXL since in addition to having a good C++ grammar, it handles grammatical ambiguities quite well and provides mechanisms for relaxing the existing grammar, making it easy to adapt to our needs. TXL is a standalone command-line tool, with few dependencies, and we got a simple proof of concept transformation working after only 4 hours of experiments with TXL.

Observation 4.5. Simplicity and integration with the languages to be transformed had a stronger influence on the selection of tools than the features or paradigm of the transformation language.

Transformation Implementation

TXL supports working with only one grammar at a time, and so the input and output syntax must conform to the same grammar. To overcome this we selected a subset of C++ expression language as our target language, since it sufficiently well captures propositional constraints over our finite set of domain variables. To handle this format we needed to relax the C++ grammar slightly to allow top-level expressions, and to preserve correctness we use a simple validation rule to that checks whether the output program is an expression in the subset of interest.

An example transformation rule for our modernization project is presented in Figure 4.2; fully capitalized identifiers refer to syntax trees. The rule matches a conditional statement without an else clause (line 3), translating it into a ternary expression (line 8) if the body statement is neither a compound statement nor another conditional (Lines 4–5); to fill in the missing branch in the ternary expression, a constant true

³<https://gcc.gnu.org>

⁴<http://eclipse.org/cdt/>

expression is used (Lines 6–7). The specified rule preconditions, in combination with the definition of the grammar and other rules, ensure that the rule is only applied when the body statement is largely consisting of a simple expression itself.

```

1 rule convert_simple_sel_stmt
2   replace [selection_statement]
3     'if '( EXP [expression] ')' STMT [statement]
4   where not STMT [contains_selection_stmt]
5   where not STMT [is_compound_stmt]
6   construct TRUE_STMT [true_case_statement]
7     'TRUE ';
8   by '( EXP ')' '?' '( STMT ')' ':' '( TRUE_STMT ')'
9 end rule

```

Figure 4.2: Example TXL transformation rule from the modernization project: translating a conditional statement into a ternary expression.

The overall algorithm applied by the transformation is:

1. The program fragment is checked for format assumptions: all branches return a value, there are no loops and goto jumps, no calls to non-pure methods, etc. The transformation is parametrized by a list of names of pure functions, since it does not do any semantic analysis to determine purity itself.
2. All preprocessor `#ifdef` directives in the program are cleaned up, and converted to ordinary `if` statements.
3. Local variable assignments are inlined so that variable accesses are replaced by assigned values—going from the last assignment to the first to maintain correct ordering—and corresponding declarations are removed.
4. All switch statements are converted to a series of `if` statements.
5. Sequenced `if` statements are simplified into nested `if` statements, such that the fragments are reduced to a single root statement.
6. `if` statements are converted to ternary expressions and `return` statements are replaced by the expression they return.

Observation 4.6. We succeeded in implementing a flow-aware syntactic transformation, including variable inlining, which enabled us to produce a result declarative expression that preserves identifiers and has reminiscent structure of the input programs.

Readability and recognizeability of the final output is important in modernization projects since developers are expected to further evolve the generate code.

Basic metrics

The entire transformation (including grammar definitions) spans 6515 SLOC. The core C++ grammar—provided from the TXL website—has 137 nonterminal rules (595 lines of code), and we additionally (re)defined 98 grammar rules to adapt the grammar to our needs. The transformation has 171 function definitions and 297 rule definitions, making 468 definitions of operations in total.

It took 3 months full-time work of experienced software developer to implement this transformation, including learning TXL, domain understanding, unit testing and meetings with the industrial partner. The cost is deemed acceptable, especially given that the company has several similar products to modernize for which the transformation and accompanying experience would be largely reusable.

The transformation execution lasts 30 minutes on the 4119 functions out of which 105 functions are not handled; these non-handled cases were caught by our fail-fast approach, and marked by the transformation to be manually migrated.

4.4 Validating the Transformation

Our transformation uses the full expressiveness of TXL including the expressive pattern matching capabilities, implicit tree traversals, and shared global state. Furthermore, there is use of reflective features such as *dynamic reparsing* which allows serialization of abstract syntax trees to textual syntax and then reinterpretation as other syntactic structures. Individual rules and functions are written non-modularly and might intermediately break syntactic and semantic correctness properties; this makes it hard to verify individual rules compositionally, and the transformation must be validated as a whole.

Observation 4.7. Our transformation was written non-modularly, which made it impossible to do compositional verification of the transformation rules.

An concrete example of a non-modular rule in our transformation is

$$\text{if } (E) \text{ return } true \rightarrow \text{return } E$$

which is only correct under assumption that the `if` statement occurs last in the main function. Since our input size and shape is fixed, we decided it would be more feasible to treat the transformation as a black-box than try to tackle the complexity of verifying rules using white-box techniques.

Observation 4.8. We were able to treat a complex transformation as a black-box, reducing validation to checking whether the provided input and corresponding transformed output are semantically equivalent.

Approach

Ordinary TXL rules and functions are constrained to replacing well-formed syntactic trees with newly built ones, making it hard for TXL programs to produce syntactically incorrect output. The correctness criterion of the transformation was to produce *semantically* equivalent C++ programs, and therefore our validation technique must be able to reason about the semantics of C++ programs.

We considered several analysis techniques, finding that symbolic execution [King, 1976] is able to build a very precise semantic model for the programs we are considering. We decided on using the precise symbolic executor KLEE [Cadar et al., 2008] that handles the majority of features used in the code base, and integrating it in the automation process proved to be cost-effective.

One of the challenges of using KLEE is that it requires compiling the input code to LLVM [Lattner and Adve, 2004] intermediate representation (LLVM-IR), including all external libraries, because otherwise the symbolic execution may provide incorrect results. Our code base however contained calls external functions with unknown implementations, and therefore we had to instrument the code to get a closed program:

1. We created stubs for unknown functions—ordinary, static member and singleton member functions alike—such that a set of arguments is matched to the same symbolic result on every call.
2. We created stubs for the data structures with constructors that straightforwardly initialize all members to symbols and overloaded equality to do structural comparison.

Function stubs are created in the *symbolic function* style [Corin and Manzano, 2011]. For each stub, a memoization table which matches input arguments with symbolic return values is allocated. When the function is called, it looks up the arguments in the memoization table: if they are found, it returns the same symbolic result as before; otherwise, a new record that stores the arguments along with a fresh symbolic result variable is created in the table.

```

1      bool defined(int p) {
2          static node<int, bool> *results;
3          static int* counter = new int(0);
4          bool* val = new bool;
5          if(!getResult(&results, &results, p, counter, val)) {
6              char symbolicname[40];
7              sprintf(symbolicname, "defined%d", *counter);
8              *val = klee_range(0, 2, symbolicname);
9          }
10         return *val;
11     }

```

Figure 4.3: A stub for an unknown Boolean function.

The stub for the Boolean function defined with unknown concrete implementation is presented in Figure 4.3. The static memoization table results and call counter are reused for all calls to the function. The function getResult looks up the input argument p in the memoization table and updating the value of the pointer val if p was already memoized. In case this is the first time p is encountered, then the call counter is incremented and val points to a new memory address, that is marked as symbolic by klee_range.

Observation 4.9. We were able to use off-the-shelf tools to perform semantic verification of programs, with a moderate amount of effort required to instrument the input to be useable with the tool.

Validation Experiment

Our validation experiment answers the following three questions refining RQ3:

RQ3.1 How large a part of the transformed code base can be verified automatically?

RQ3.2 How much additional effort would it require to verify the rest of the code base?

RQ3.3 How can our verification effort be generalized to other similar modernization projects?

We address RQ3.1 in Section 4.5 by running the verification procedure on the input and transformed output, reporting and classifying the results. Section 4.7 discusses the challenges (and solutions to these challenges) that appeared during verification (RQ3.2) and what parts of our verification procedure is generalizable to other transformation tools and projects (RQ3.3). Executing the validation procedure on all transformed functions lasts 7 minutes, where 3348 of the functions are evaluated trivially to pass verification—the output is identical to the input.

4.5 Bug Analysis

There were 771 functions out of 4119 in the code base that could not be trivially validated. We present the statistics in Table 4.1.

Table 4.1: Erroneous transformation cases caught by each step of the validation process.

	Step	#Cases
1	Failing transformation precondition (not handled, requiring manual inspection)	105
2	Failing silently due to unhandled syntactic structures (caught statically by TXL during preliminary steps of verification)	3
3	Caught by C++compiler	3
4	Checked for equivalence using KLEE	640
4a	Validated being equivalent	562
4b	Concrete bug cases with provided counter-examples	50
4c	False positives with spurious counter-examples (due to over-approximation of functions, and representation mismatch)	28
5	Unhandled cases containing assertions (intentional, due to design limitations of the validation technique)	20

Observation 4.10. It was possible to analyze a substantial amount of the modernized code automatically with only 20 corner cases left to be handled manually.

We analyze the bugs found by verification (3 & 4b) below. In Section 4.7 we will discuss the types of spurious counter-examples (4c), the unhandled cases due to design limitations (5), and propose solutions for these issues.

Analysis of Bugs Found by Verification

Our technique had identified seven bugs present in the transformation. While these bugs varied in nature, they had one important thing in common: they were all related to execution semantics and would have been hard to find with a syntactic check or a static semantics check using e.g., type analysis.

Observation 4.11. When the code base of our modernization project reached a certain complexity it became infeasible to find all bugs through expertise and unit testing. Validation of semantics was essential to ensure that the output code worked correctly.

Bug 4.1: *Function call is dropped in some paths.* The most widespread errors in the output expressions were missing function calls, that were present in the original code. This was not contained to calls of a particular function or particular place in the output expression, and other calls to the same function might still be present in the rest of the output expression.

The bug was caused by an *incomplete rewrite rule*: the rule matched a functional call in a return statement, but did not use the function call when constructing the replacement.

Bug 4.2: *Structure replaced by a constant integer.* Another widespread but simple bug is where the input function declares a variable with a class type, calls its object initializer with multiple arguments and then returns the variable. Here, the transformation seemed to instead return the first argument given to the object initializer which often had an incompatible type like `int`.

This bug happened due to *misuse of deep pattern search* and a *broken rule assumption*. The variable inlining procedure expected variables with simple types and so used a deep pattern match to extract the target

value; for structure types, the deep pattern will however simply match the first argument value of the initializer—ignoring the rest—and use that for the inlining.

Bug 4.3: *Conditional branches are dropped.* This bug was also caused by *incomplete rewrite rules*, and made the transformation ignore all branches following a nested `if`-statement (also referred to in Section 4.1). The transformation rule matches a nested conditional possibly having an `else`-branch, correctly handling the nesting but not the `else`-branch.

Bug 4.4: *The unexpected exceptions.* This bug was surprising since it mostly happened in huge functions with heavy nesting of control flow. While the input function seems total and returns a correct result on all paths, the transformation produces an output which contains a branch that throws an exception stating that the branch should be invalid.

This bug happens due to *overconstrained pattern matching* and *broken rule assumption*. When a sequence of nested conditionals followed by a return statement is matched by the transformation, it tries to put the final return statement inside the branches of the previous conditional. However, the pattern was overconstrained—not matching the inputs it should have handled—and so the containing rule was never applied; later, when the transformation tries to convert the statement to an expression it finds a branch with no return statements and replaces it with a throw statement—as part of the fail-fast approach—since it did not expect this case to be possible.

Bug 4.5: *Use of undeclared variables.* This bug was caught during compilation, and happened when the original input contained declarations to local variables that were not inlined correctly by the transformation. Recall, that the transformation removes all local declarations after inlining, but in this case there were still some variable accesses that were not inlined which resulted in undeclared variable errors when running the result expression by the compiler.

This bug occurs due to a combination of *dynamic reparsing capabilities* and *wrong target type* in expression. To control the number of iterations of inlining substitution, the transformation replaces variable nodes with the string representation of their assigned expressions, using the textual output capabilities of TXL. This ensures that the substitution terminates,

but might also be incorrect if the transformation has not finished migrating the serialised subtrees. In general, it is caused by challenges in implementing a semantics-sensitive operation (i.e., inlining) syntactically.

Bug 4.6: *Negation dropped in result.* The simplest bug found by the KLEE-based verifier is where the transformation had transformed the whole input correctly except a negation operation which was missing in the output. This bug occurs due to *misuse of deep pattern search* of TXL, because a transformation rule searched for a more specific syntax type than necessary and thus ignored more complex types of terms. This bug is particularly interesting since it would not be caught by a simple type-based static analysis, but would be caught by a more sophisticated analysis such as the inductive shape analysis in Chapter 7.

Bug 4.7: *Conditional with error code assignment dropped.* A particularly interesting bug is where the input has a function that contains a conditional statement assigning a value to an error code pointer variable, in addition to returning a separate value. In this case, the transformation will produce output that will completely remove the conditional branch and only keep the final return value, making the function produce an incorrect result.

This bug happens due to the *dynamic reparsing capabilities* and *eager removal of source data*. It originates in the inlining phase where some abstract syntax is broken by wrongly inserted textual syntax, and subsequently a rule that removed empty conditional branches was applied.

Bug 4.8: *Variable declarations without assignment not handled.* This bug was caught statically in the cases when the transformation finished, but the output was empty. Similar to Bug 4.2, a combination of a *broken rule assumption* and *misuse of deep pattern search* was the cause of this bug. Inlining assumes that declaration and initialization of local variables happen in a single statement, but the cases triggering this bug had separate declaration and initialization. The declaration removal rule used deep search to identify statements containing local variable assignments and since program consisted of a single large if-statement with the assignments, it was completely removed.

Classification summary Most cases were affected by Bug 4.1 where there were 23 cases in total, and followed by Bug 4.2 which had 15 cases in total; both of which were simple in nature; this is unsurprising since

function calls and object initialisations are common constructs in C++, and a simple mistake in the transformation of these features will therefore affect a large number of analyzed functions. The more interesting Bugs 4.3 and 4.4 had 5 cases in total each; this type of bugs often appeared in larger files with a complex nesting of conditionals, and would therefore have been hard to immediately spot manually or with simpler unit tests. Finally, the remaining bugs (4.5, 4.6, 4.7, 4.8) had respectively 3, 1, 1 and 3 cases in total; these errors represent corner cases that were either not caught by the preconditions of the transformation, or occurred where an intermediate assumption of the transformation was wrong.

Observation 4.12. Simple bugs hit wide, complex bugs hit deep. Simple semantic errors affected a large number of functions while complex errors were found in a few but bigger functions.

4.6 Formal Justification of the Procedure

The transformation translates many functions individually, each of which needs to be translated in a semantics preserving manner. We view these functions (and programs in general) as black-boxes relating input and output.

Definition 4.1. A *program* P is a set of imperative statements that state how to calculate the designated output variable ret from a set of input variables $\mathbb{X}_{\text{in}} = \{x_1, \dots, x_n\}$.

Definition 4.2. A *concrete store* σ is a *partial function* mapping program variables Var into values Val , i.e. $\sigma : \text{Var} \rightarrow \text{Val}$. The set of values Val range over constants from C++ base types: Boolean, bounded integer, float, etc.

A concrete execution starts with an *initial state*, σ_{in} , where all input variables are assigned some initial values. The effect of executing each statement s of program P in a state σ is a successor state σ' , written as $\sigma \xrightarrow{s} \sigma'$. When all statements are executed, the program reaches a final state σ_{out} , in which the value of the output variable ret is well-defined (i.e., $\text{ret} \in \text{dom } \sigma$).

Definition 4.3. A *concrete execution path* $\pi = \sigma_{\text{in}}, \sigma_1, \dots, \sigma_{\text{out}}$ of the program P is a sequence of states, such that σ_{in} is an initial state, every next state in the sequence is obtained by sequentially executing statements from P one by one, and σ_{out} is the final state.

Definition 4.4 (Concrete program path semantics). The *path semantics* of program P —called $\llbracket P \rrbracket^{\text{trace}}$ —is defined to be the set of all valid concrete execution paths π of P .

Definition 4.5 (Denotational program semantics). The *denotational semantics* of program P is a partial function $\llbracket P \rrbracket : \text{Val}^k \rightarrow \text{Val}$ defined by: $\llbracket P \rrbracket(\sigma_{\text{in}}(i_1), \dots, \sigma_{\text{in}}(i_k)) = \sigma_{\text{out}}(\text{ret})$, for any concrete execution path $\pi \in \llbracket P \rrbracket^{\text{trace}} = (\sigma_{\text{in}}, \dots, \sigma_{\text{out}})$.

Definition 4.6 (Semantic equivalence). Two programs P and P' are *semantically equivalent*, written $P \sim P'$, if for any collection of values v_1, \dots, v_k it holds: $\llbracket P \rrbracket(v_1, \dots, v_k) = \llbracket P' \rrbracket(v_1, \dots, v_k)$.

Determining the semantic equivalence of programs using concrete path semantics is intractable due to the immense range of input values. Instead, we use a symbolic execution to cluster the input values using a set of constraints called path conditions.

Symbolic execution

In *symbolic execution* the program does not assign concrete values to its variables; instead, it assigns symbolic expressions containing uninterpreted symbols abstractly representing user-assignable values in a concrete execution of the program. We let Sym represent the set of uninterpreted symbols $x^?, y^? \in \text{Sym}$ which we write as variables with a lifted question mark. In the initial execution state, each possible input variable x_i is usually assigned a corresponding unique symbol $x_i^?$.

For example, let x and y be input integer variables, then the concrete semantics of $\text{ret} = x + y$ is the set $\{([x \mapsto 0, y \mapsto 0], [x \mapsto 0, y \mapsto 0, \text{ret} \mapsto 0]), ([x \mapsto 0, y \mapsto 1], [x \mapsto 0, y \mapsto 1, \text{ret} \mapsto 1]), \dots\}$, where initial states are all possible assignments of integer values to x and y . However, the symbolic path semantics of $\text{ret} = x + y$ will contain only one symbolic execution path $([x \mapsto x^?, y \mapsto y^?], [x \mapsto x^?, y \mapsto y^?, \text{ret} \mapsto x^? + y^?])$, where $x^?$ and $y^?$ are fresh symbols.

Symbolic execution approximates a set of different concrete paths into a single symbolic one by following all branches whenever a branching or looping statement is encountered. In the same time, for each branch it maintains a set of constraints called the *path condition*, which must hold on the execution of that path.

Definition 4.7. A *symbolic expression* $\hat{e} \in \widehat{\text{Exp}}$ can be built out of constant values from Val , symbols from Sym , and arithmetic-logic operations.

Definition 4.8. A *symbolic store* $\widehat{\sigma}$ is a function mapping program variables Var into symbolic expressions $\widehat{\text{Exp}}$, i.e. $\widehat{\sigma} : \text{Var} \rightarrow \widehat{\text{Exp}}$. The initial symbolic state $\widehat{\sigma}_{\text{in}}$ maps each input variable $x_i \in \mathbb{X}_{\text{in}}$ to a fresh symbol $x_i^? \in \text{Sym}$.

Definition 4.9 (Constrained symbolic state). A *constraint* is any Boolean symbolic expression $\widehat{b} \in \widehat{\text{BExp}}$. A *constrained symbolic state* is a pair $(\widehat{\sigma}, \widehat{b})$, which constraints the symbolic expressions in $\widehat{\sigma}$ with a Boolean symbolic expression \widehat{b} .

Definition 4.10 (Symbolic execution path). A *symbolic execution path* of the program P is a sequence of constrained symbolic states $((\widehat{\sigma}_{\text{in}}, \text{true}), (\widehat{\sigma}_1, \widehat{b}_1), \dots, (\widehat{\sigma}_{\text{out}}, \widehat{b}_{\text{out}}))$, where the initial state $\widehat{\sigma}_{\text{in}}$ is unconstrained, and the constraint produced for the final state, \widehat{b}_{out} , represents the path condition.

Note how the resulting set of symbolic execution paths partitions the set of concrete execution paths. For the program that computes the absolute value of an integer variable x , there are two different paths returned by symbolic execution: $(([x \mapsto x^?], \text{true}), ([x \mapsto x^?, \text{ret} \mapsto x^?], x^? \geq 0))$ and $(([x \mapsto x^?], \text{true}), ([x \mapsto x^?, \text{ret} \mapsto -x^?], x^? < 0))$. If the initial value of x is non-negative, then the return value is the symbolic expression $x^?$; otherwise, the return value is $-x^?$. Hence, the set of all concrete execution paths—determined by the input values of x —has been partitioned in two sets: those for which $x \geq 0$ holds and those for which $x < 0$ holds.

Proposition 4.1. *For each concrete execution path $\pi = (\sigma_{\text{in}}, \sigma_1, \dots, \sigma_{\text{out}})$ of the program P , there exists the corresponding symbolic execution path $\widehat{\pi} = ((\widehat{\sigma}_{\text{in}}, \text{true}), \dots, (\widehat{\sigma}_{\text{out}}, \widehat{b}_{\text{out}}))$, such that the following equations are satisfied:*

$$\begin{aligned} \widehat{\sigma}_{\text{in}} &= [x_i \mapsto x_i^?, \dots, x_n \mapsto x_n^?] \\ \theta &= [x_0^? \mapsto \sigma_{\text{in}}(x_0), \dots, x_n^? \mapsto \sigma_{\text{in}}(x_n)] \\ ((\widehat{\sigma}_{\text{out}})\theta)(\text{ret}) &= \sigma_{\text{out}}(\text{ret}) \\ (\widehat{b})\theta &= \text{true} \end{aligned}$$

Theorem 4.1. *Two programs P and P' are semantically equivalent $P \sim P'$ iff for each value $v \in \text{Val}$ it holds:*

$$\bigvee_{(\widehat{\sigma}_{\text{out}}, \widehat{b}_{\text{out}}) \in \Sigma} (\widehat{\sigma}_{\text{out}}(\text{ret}) = v \wedge \widehat{b}_{\text{out}}) \iff \bigvee_{(\widehat{\sigma}'_{\text{out}}, \widehat{b}'_{\text{out}}) \in \Sigma'} (\widehat{\sigma}'_{\text{out}}(\text{ret}) = v \wedge \widehat{b}'_{\text{out}})$$

where

$$\Sigma = \left\{ (\hat{\sigma}_{\text{out}}, \hat{b}_{\text{out}}) \mid ((\hat{\sigma}_{\text{in}}, \text{true}) \dots (\hat{\sigma}_{\text{out}}, \hat{b}_{\text{out}})) \text{ is a symbolic path of } P \right\}$$

$$\Sigma' = \left\{ (\hat{\sigma}'_{\text{out}}, \hat{b}'_{\text{out}}) \mid ((\hat{\sigma}_{\text{in}}, \text{true}) \dots (\hat{\sigma}'_{\text{out}}, \hat{b}'_{\text{out}})) \text{ is a symbolic path of } P' \right\}$$

4.7 Experiential Reflections

Unreliability of academic tools is a known challenge in software engineering action research; due to lack of support services, these tools are rarely adopted by companies, except in rare cases where the company is able to further develop and maintain the tool in-house. Our experiences with this project tells us that this issue is less detrimental for modernization projects, since the tools are only required for use in a short transitional period and there is more freedom for customization.

The applied design and validation principles translate easily to other program and model transformation languages. Even though we verify expressive semantic properties between input and output, the transformation itself is treated as black-box and the method is oblivious to the choice of transformation language. Uncertainty remains regarding the generalizability of the technique with regards to transformations that should support a broader variety of input, and whether the identified bugs are specific to this case, this input and output languages, and TXL.

Validation Challenges

In practical projects, many of the idealized assumptions made in research are not fulfilled, e.g., that data is represented in a particular way or that we have access to all library code at validation time. We will discuss three concrete issues we had when validating the transformed code, and how we effectively mitigated them.

Representation of Boolean expressions

In C++ integer valued expression can be used as logical tests (inside if-statements etc.), where any non-zero value counts as true and zero counts as false. If an integer variable a is used only as a logical condition both in the input and output programs it would be pragmatically fine. However, our transformation contains simplification rules which convert statements of form `if (a) return true; else return false;` to

return a; which clearly has different semantics for non-Boolean values. In cases where we are certain that specific integer variables are only used as conditionals we instruct KLEE to assume that these variables have values lying in range $[0, 2)$.

Over-approximation of Function Semantics

We use stubs to model external functions whose implementation we do not know, which is over-approximating since we only know that equal arguments provide the same result values. If any of these functions however had equivalent implementations and we used a different stub for each one, calling the two stubs with the same parameters would result in distinct return values. This led to a number of false positives where KLEE decided that the input and output programs were not equivalent, and was solved by using the same stub for functions which were known to be identical *a priori*.

Assertions

When KLEE meets a C++ assertion that fails on a possible path, it will immediately halt execution for that particular path. This concretely means that it will never check whether the input and output functions have the same results, or in this case rather both fail. Instead of using the default assertion function, one could instead use a stub that throws a recoverable error on failing conditions; thereafter, one could check whether both the input and output programs failed on the same paths and if they did one could consider the paths to be equal. Our code base contains 20 cases affected by this limitation, but implementing the suggested solution was not feasible in the allocated time.

4.8 Related Modernization Efforts

Semantic Designs (SD)⁵ presents a similar automation-based modernization effort is presented in a joint project with Boeing⁶ [Akers et al., 2007] which uses Semantic Designs's commercial transformation and analysis tool DMS [Baxter et al., 2004] to convert an old component-based C++ code base to a standardized CORBA [Siegel, 2000] architecture. This is in contrast to our case study which relied solely on freely

⁵<https://www.semanticdesigns.com>

⁶<http://www.boeing.com>

available tools like TXL and KLEE, and we more importantly had a significantly more thorough validation effort with precise descriptions of the challenges we encountered and the lessons we learned.

A series of papers [Selim et al., 2013, 2014, 2015] discuss a case study that aims to design and verify a model transformation for modernizing an existing collection of proprietary models to conform to the standardised AUTOSAR [Bunzel, 2011] format. The transformation [Selim et al., 2015] was initially encoded in a limited expressive subset of the ATL model transformation language [Jouault et al., 2008] and then verified for structural properties [Selim et al., 2013]. The same verification effort was then repeated [Selim et al., 2014] more efficiently by symbolically executing a version of the transformation re-encoded in DSLTrans [Barroca et al., 2010]. While these verification tools and the presented case study have significant contributions to the model transformations community, they were not applicable in our study due to the difference in expressiveness between TXL and the verified non-Turing-complete subset of ATL/DSLTrans, and the complexity of the property we wanted to check (behavioral equivalence versus structural properties).

4.9 Recap

This chapter presented a collaborative effort focusing on designing and verifying an industrial modernization project implemented using the high-level transformation language TXL. It showed the usefulness of the high-level transformation language features—such as traversals and fixed-point iteration in the rule-based execution—in providing a practical and cost-effective development cycle for migrating a large code base. More importantly for this dissertation, the modernization project highlighted the practical need for formal verification methods such as symbolic execution: even though our the modernization transformation was written by an expert transformation programmer and had been well-tested, the validation effort found seven serious bugs spread over 50 concrete instances. We furthermore provided a detailed analysis of the bugs and their causes, which was a particularly useful inspiration for directing the verification work in the rest of the thesis.

Chapter 5

Symbolic Execution of High-Level Transformation Languages^{*}

The *rename field* refactoring in Figure 5.1 is an example of a transformation that is best written in a high-level transformation language. The transformation changes the name of a given field in the definition of a class and ensures that all relevant field accesses use the new field name [Fowler, 1999].

This transformation, while simple in nature, requires a lot of boilerplate to define in traditional programming languages—like C or Java—since one needs to construct visitors that recursively traverse each of the many cases in the target program syntax tree. In contrast, high-level transformation languages provides these features directly in the form of *type-directed querying and manipulation*. These operations allow *deep* matching and rewriting of structures by following the types of objects and references between them. In our example in Figure 5.1, a type-directed query makes it possible to retrieve all accesses to the *credit* field where the target expression has type *Account*, using only a couple of lines.

Transformations are complex programs and as such prone to bugs; for our *rename field* example, a bug could be that we modified accesses to fields with similar names but an unrelated class. Due to the complexity of transformations these bugs are hard to find and expensive to fix; this creates a need for automated formal techniques that support verifying the correctness of transformations, which is an important step to-

^{*}This chapter is based on Al-Sibahi et al. [2016] which is the technical report version of the conference paper, Al-Sibahi et al. [2016].

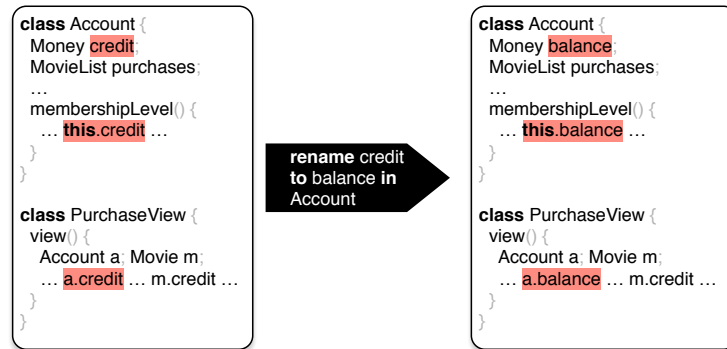


Figure 5.1: An example execution of the *Rename-Field* refactoring: rename the definition of `credit` to `balance` and update all references accordingly.

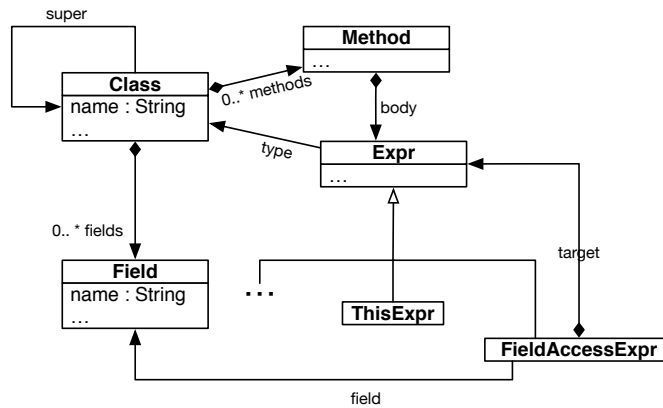
wards increasing the trustworthiness of our language implementations and tools [Cadaru and Donaldson, 2016; Schäfer et al., 2009; Hoare, 2005].

In this chapter, we present a foundational symbolic execution technique that handles high-level transformation features as first-class. Concretely, our contributions are:

- A formal symbolic execution technique supporting complex concepts such as symbolic sets, ownership constraints, type-directed querying and manipulation, and fixed-point iteration.
- An evaluation of the symbolic execution technique when used for white-box test generation using realistic model transformations and refactorings, showing that our symbolic executor makes white-box test generation for transformation feasible.
- A comparison of our symbolic execution technique to object-oriented symbolic executors, highlighting the difficulties of dealing with target high-level features as second-class.

5.1 Overview

Let us start by discussing our motivating *rename field* example (Figure 5.1) in more detail. Observe that the refactoring program needs two key components: *type definitions* (sometimes called *meta-model*) for the data and the actual transformation *code*. We will use (minimalistic)



(a) Abstract syntax for simple object-oriented programs

```

1 input:
2   target_class: Class,
3   old_field: Field,
4   new_field: Field
5 precondition:
6   old_field ∈ target_class.fields
7   ∧ new_field ∉ target_class.fields
8
9 target_class.fields :=
10  (target_class.fields \ old_field) ∪ new_field
11 foreach faexpr ∈ target_class match* FieldAccessExpr do
12   if faexpr.field = old_field ∧
13     faexpr.target.type = target_class then
14     faexpr.field := new_field
15   else skip
  
```

(b) The *Rename-Field* Refactoring implemented in the small formal transformation language, TRONFigure 5.2: A simplified version of the *rename-field* refactoring example in TRON

class diagrams to show the former (e.g., Figure 5.2a), and a compact formally defined transformation language TRON for the latter (e.g., Figure 5.2b). These two notational choices incorporate some key common characteristics of transformations, which we discuss below.

Figure 5.2a shows the types for the abstract syntax of a hypothetical object-oriented language, which we will use to model the programs we are refactoring. We show the classes¹ and properties relevant for our

¹The abstract syntax types describe a hypothetical object-oriented programming language and not of the formal core transformation language (TRON).

refactoring, while omitting irrelevant details. In our example, each class has a name (an attribute), and *contains* two collections, one for fields and one for methods. Recall that in class diagrams, a black diamond is used to represent *containment* references, which are traditionally found both in object-oriented modeling languages and grammar-based languages like TXL.

Additionally, each class may simply *refer* to a possible super-class, which is denoted using a line without the diamond symbol. All the represented features—classes, containment references, simple references, and simply-typed attributes—are typical of transformation languages we want to handle by our technique.

For simplicity of presentation, we let the body of each method in our example be an expression. Expressions themselves can come in many different kinds (thus the use of *inheritance*), but we only show expressions representing **this** and field access expressions since they are the ones relevant for the example.

A simplified implementation of the *rename field* refactoring, is presented in Figure 5.2b using the core transformation language (TRON). We will discuss the example based on general intuitions, presenting the formal semantics of the core language in Section 5.2–Section 5.3. In the start (Lines 1–4), we list the input parameters—references to a class, the field with the old name, and the replacing field with the new name—and the application precondition (Lines 5–7) which specifies that the old field has to be contained in the fields of the input class, whereas the new must not.

We begin the refactoring by removing the old field definition from the fields of the class and adding the new field definition (Lines 9–10). Then, in Line 11, all field access expressions in the class are matched and gathered into a single set using a *deep type-directed query*, which collects instances of `FieldAccessExpr` contained transitively in the input class. After the deep type-directed query, Line 11 binds each element of the matched objects to *faexpr* executing Lines 12–15 for each of these objects. If the expression accesses the refactored field (Lines 12–13) then the field reference is updated to point at the new field (Line 14). It is typical for transformation languages that references are redirected or attributes are changed.

The example demonstrates how TRON, despite being minimalistic, supports the key high-level transformation features such as type-

directed querying and iteration via **foreach**, collection operations, and imperative updates for manipulation (like ATL or Rascal).

Symbolic Execution of Transformations

Symbolic execution is an effective way to check the presence of bugs in transformations, since it systematically explores the various transformation program paths. Figure 5.3 presents an overview of our symbolic execution technique for transformations; here, the symbolic executor expects as input the transformation program, along with the required type definitions. The initial step is to run the symbolic executor (see Section 5.3) on the input transformation and generate a finite set of path conditions. These path conditions are logical formulae constraining the shape, types and range of input data, achieved by refining input constraints according to the semantics of each statement in the given transformation.

We use the model finder to prune those paths which produce unsatisfiable formulae so that only valid paths are considered. In our implementation, the model finder uses the relational constraint solver KodKod [Torlak and Jackson, 2007] to check the existence of a suitable model satisfying a target formula within a bounded scope, possibly failing when either the formula is unsatisfiable or the scope is too small.




Figure 5.3: High level architecture of the symbolic executor.

5.2 A Demonstration Language

We developed the small formal transformation language TRON as a methodological device, to keep the formal work, discussions, and the

presentation focused, and to allow agile experimentation; TRON is a decoy language *not* meant to be used by programmers, but is designed so that our ideas remain applicable to real-world transformation languages.

Table 5.1 shows how TRON incorporates characteristic features of high-level transformation languages. We present TRON in two parts: i) the meta-model that captures structures of the manipulated data, and ii) the operational part of the language that describes computations.

Language TRON Feature	ATL	Haskell	Maude
Containment	Containment references	Algebraic Data Types	Many-Sorted Terms
Set expressions	OCL collections and collection operations	Standard library	Standard library
Shallow matching	Type testing via <code>oclIsKindOf</code>	Pattern matching	Rewrite rules
Deep matching	Transformation rule definition	Generic traversal via Uniplate	Rewrite rules and strategies
Fixedpoint iteration	Lazy rules, recursive helpers	Recursive functions	Rewrite strategies

Table 5.1: Relating TRON features to existing high-level transformation languages

Data Model. The data in TRON is described by types that capture the common features of high-level transformation languages: constructors, containment, references and generalization. It is essentially a formal model for the kind of structures like the one represented in Figure 5.2a.

A data model is a tuple: $(\text{Class}, \text{Field}, \text{gen}, \text{ref})$, where *Class* is the set of *classes*, *Field* is the set of *fields*, partitioned into *contained* fields $\text{Field}_\blacklozenge$ and *referenced* fields $\text{Field}_\rightsquigarrow$. Later, we use c to range over class names (*Class*), and f to range over field names (*Field*). A class has at

most one superclass, described by the generalization relation: $\text{gen} \subseteq \text{Class} \times \text{Class}$, where $c \text{ gen } c'$ means that c is a subtype of c' . Each field has a corresponding type, a class. This is represented by the references relation $\text{ref} \subseteq \text{Class} \times \text{Field} \times \text{Class}$, where $\text{ref}(c, f, c')$ means that the class c has a field f of type c' . We generally expect that gen has the expected properties of a generalization relation, namely that there is a strict ordering of generalization (no cycles); similarly, we expect that reference definitions in ref are not overridden by subtypes, i.e. if for any class c a supertype has defined a typing for a field f , then c must have the same typing for f .

To get all the fields defined for a class c or any of its supertypes, we define the following function:

$$\text{fields}(c) = \{(f, c'') \mid c \text{ gen}^* c' \wedge \text{ref}(c', f, c'')\}$$

We do not explicitly handle simple types in our formal model, instead they can be modeled theoretically as classes. For instance, we assume a class `Integer` with instances representing integer numbers; integer attributes can then be modeled as references to this class. We do handle simple types in our symbolic execution tool by using symbolic variables of corresponding simple types.

Heap Representation. Concrete TRON programs are executed over finite concrete heaps ($h \in \text{Instance} \times \text{Field} \rightarrow \wp(\text{Instance})$) that contain instances organized into structures using containment links and simple links (a link is a concrete instantiation of a field). In particular each link f of an instance o , can point to a set of instances os . Instances are typed at runtime using a type environment ($\Gamma \in \text{Instance} \rightarrow \text{Class}$). An example heap is presented in Figure 5.4, which describes a possible definition of the `Account` class used in Figure 5.1.

For the remainder of this chapter we will only consider well-formed heaps where all instances are typed and their structure conforms to the static typing provided by the data model. Furthermore, we assume that in well-formed heaps each instance can at most be pointed to by a single containment link (no sharing) and that there are no cycles in containment (acyclicity). Note that these restrictions do not apply for simple links, which still allow cycles and sharing.

Abstract Syntax. The core TRON constructs include access to variables and fields, constants, object construction, assignment, sequencing and

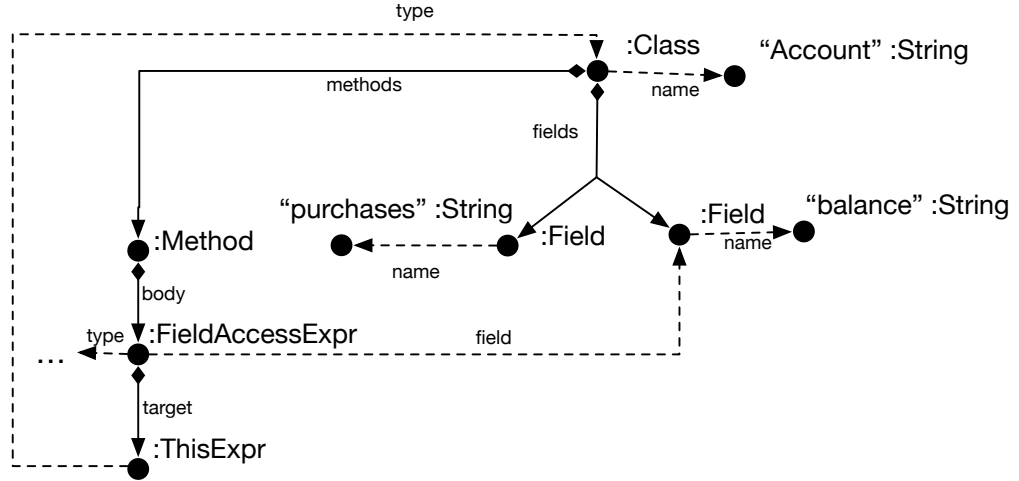


Figure 5.4: A heap instantiating one model of Figure 5.2a, inspired by the Account class in Figure 5.1. Dots (●) represent instances, diamond affixed lines represent containment links, dashed lines represent simple links.

branching. The syntax is summarized in the following grammar:

$$\begin{aligned}
 \text{SetExpr } \ni e &::= x \mid \emptyset \mid e_1 \cup e_2 \mid e_1 \cap e_2 \mid e_1 \setminus e_2 \\
 \text{BoolExpr } \ni b &::= e_1 \subseteq e_2 \mid e_1 = e_2 \mid \neg b \mid b_1 \wedge b_2 \\
 \text{MatchExpr } \ni me &::= e \mid e \mathbf{match} c \mid e \mathbf{match}^* c \\
 \text{Statement } \ni s &::= \mathbf{skip} \mid s_1; s_2 \mid x := e \mid x := e.f \\
 &\quad \mid x := \mathbf{new} c \mid e_1.f := e_2 \mid \mathbf{if} b \mathbf{then} s_1 \mathbf{else} s_2 \\
 &\quad \mid \mathbf{foreach} x \mathbf{in} me \mathbf{do} s \mid \mathbf{fix} e \mathbf{do} s
 \end{aligned}$$

where x is a variable, f is a field name, and c is a class name. The set expressions (e) and Boolean expressions (b) are standard. Match expressions (me) include “ e **match** c ” which allows finding all objects computed by e that are instances of class c . For example, consider a set of objects representing program expressions to be referenced $exprs = \{te_1, te_2, fae_1, fae_2\}$ where te_i is of type ThisExpr (Figure 5.2a) and fae_i is of type FieldAccessExpr. The TRON expression “ $exprs$ **match** ThisExpr” returns the set $\{te_1, te_2\}$, similarly “ $exprs$ **match** FieldAccessExpr” returns $\{fae_1, fae_2\}$ and “ $exprs$ **match** Expr” return the complete set $exprs$.

A deep variant of the pattern matching, e **match**^{*} c , is also provided. It matches objects nested at an arbitrary depth inside other ob-

jects, following the containment references ($\text{ref}\blacklozenge$). This is similar to the matching capabilities in many of the model transformation, term and graph rewriting languages. A classical example here would be to get all objects representing program variables in a term, i.e., the expression $\text{expr match}^* \text{Var}$ —for a class `Var` representing variables—would return a set that has all variables transitively contained in expr .

Most of the statements, s , are standard formulations from Java or IMP; from left to right, the statements are: skip, sequencing, branching, variable assignment, assignment of a field value, object creation (**new**) and assignment to a field.

There are two looping constructs in TRON. The “**foreach** x **in** me **do** s ” iterates over the set of elements matched by me , binding each element to x , and executes then statement s for each of them. The “**fix** e **do** s ” loop executes the body s , and continues to do so as long as the values of e after and before iteration differ; therefore expression e defines the part of the heap which is relevant for this fixed point iteration (a control condition). By allowing the statement to explicitly depend on a local control condition, it is possible to create temporary helper values on the heap (outside e) without influencing the loop termination. This allows explicit modeling of the implicit fix point iteration that is also supported by many high-level transformation languages where rewrite rules are repeatedly applied until no rule is further applicable.

Concrete Semantics

The formal concrete semantics of the TRON constructs is presented in Figure 5.5 and Figure 5.6. Most constructs are handled as expected from the informal description: the only constructs that require more elaboration are match expressions and field updates. Match expressions evaluation requires the auxiliary functions `match`, which gathers the subset of instances which are of a subtype of the target type, and the auxiliary function `dcs`, which collects all are transitively contained instances for a given set of instances in the heap. Field updates require that assigned values are correctly typed via the typed relation, and ensure to remove old containment references to the assigned instances in case of assignment to a containment reference (change of ownership).

Set Expressions

$$\mathcal{E}[[e]]\sigma \in \wp(\text{Instance})$$

$$\mathcal{E}[[x]]\sigma = \sigma(x) \quad \mathcal{E}[[\emptyset]]\sigma = \emptyset$$

$$\mathcal{E}[[e_1 \cup e_2]]\sigma = \mathcal{E}[[e_1]]\sigma \cup \mathcal{E}[[e_2]]\sigma$$

$$\mathcal{E}[[e_1 \cap e_2]]\sigma = \mathcal{E}[[e_1]]\sigma \cap \mathcal{E}[[e_2]]\sigma$$

$$\mathcal{E}[[e_1 \setminus e_2]]\sigma = \mathcal{E}[[e_1]]\sigma \setminus \mathcal{E}[[e_2]]\sigma$$

Boolean Expressions

$$\mathcal{B}[[b]]\sigma \in \{\text{tt}, \text{ff}\}$$

$$\mathcal{B}[[\neg b]]\sigma = \neg \mathcal{B}[[b]]\sigma$$

$$\mathcal{B}[[e_1 \subseteq e_2]]\sigma = \mathcal{E}[[e_1]]\sigma \subseteq \mathcal{E}[[e_2]]\sigma$$

$$\mathcal{B}[[e_1 = e_2]]\sigma = \mathcal{E}[[e_1]]\sigma = \mathcal{E}[[e_2]]\sigma$$

$$\mathcal{B}[[b_1 \wedge b_2]]\sigma = \mathcal{B}[[b_1]]\sigma \wedge \mathcal{B}[[b_2]]\sigma$$

Match Expressions

$$\mathcal{M}[[me]](\sigma, \Gamma, h) \in \wp(\text{Instance})$$

$$\mathcal{M}[[e]](\sigma, \Gamma, h) = \mathcal{E}[[e]]\sigma$$

$$\mathcal{M}[[e \text{ match } c]](\sigma, \Gamma, h) = \text{match}(\mathcal{E}[[e]]\sigma, c, \Gamma)$$

$$\mathcal{M}[[e \text{ match}^* c]](\sigma, \Gamma, h) = \text{match}(\text{dcs}(\mathcal{E}[[e]]\sigma, h), c, \Gamma)$$

$$\text{match}(os, c, \Gamma) = \{o \mid o \in os \wedge \Gamma(o) \text{ gen}^* c\}$$

$$\text{dcs}(os, h) = \{o' \mid o \in os \wedge o \diamond_h^* o'\}$$

$$o \diamond_h o' \text{ iff } \exists f \in \text{Field}_\diamond. o' \in h(o, f)$$

Figure 5.5: Concrete semantics for TRON

Statements

$$\text{update}(o, f, os, h) = \begin{cases} h[(o, f) \mapsto os] & \text{if } f \in \text{Field}_{\rightsquigarrow} \\ h'[(o, f) \mapsto os] & \text{if } \begin{matrix} f \in \text{Field}_{\blacklozenge} \wedge \\ \nexists o' \in os. o' \blacklozenge_h^* o \end{matrix} \end{cases}$$

$$\text{where } h' = [(o, f) \mapsto \text{do-f}(h(o, f), f, os) \mid (o, f) \in \text{dom } h]$$

$$\text{do-f}(os', f, os) = \begin{cases} os' & \text{if } f \in \text{Field}_{\rightsquigarrow} \\ os' \setminus os & \text{if } f \in \text{Field}_{\blacklozenge} \end{cases}$$

$$\text{typed}_{\Gamma}(o, f, os) \text{ iff } \exists c. (f, c) \in \Gamma(o) \wedge \forall o' \in os. \Gamma(o') \text{ gen}^* c$$

$$\frac{}{\text{skip}, \sigma, \Gamma, h \Longrightarrow \sigma, \Gamma, h}$$

$$\frac{s_1, \sigma, \Gamma, h \Longrightarrow \sigma'', \Gamma'', h'' \quad s_2, \sigma'', \Gamma'', h'' \Longrightarrow \sigma', \Gamma', h'}{s_1; s_2, \sigma, \Gamma, h \Longrightarrow \sigma', \Gamma', h'}$$

$$\frac{\sigma' = \sigma[x \mapsto \mathcal{E}[[e]]\sigma]}{x := e, \sigma, \Gamma, h \Longrightarrow \sigma', \Gamma, h}$$

$$\frac{\begin{matrix} o' \text{ fresh} \\ \sigma' = \sigma[x \mapsto \{o'\}] \quad \Gamma' = \Gamma[o' \mapsto c] \\ h' = h[(o', f) \mapsto \emptyset \mid (f, -) \in \text{fields}(c)] \end{matrix}}{x := \text{new } c, \sigma, \Gamma, h \Longrightarrow \sigma', \Gamma', h'}$$

$$\frac{\mathcal{E}[[e]]\sigma = \{o\} \quad \sigma' = \sigma[x \mapsto h(o, f)]}{x := e.f, \sigma, \Gamma, h \Longrightarrow \sigma', \Gamma, h}$$

$$\frac{\begin{matrix} \mathcal{E}[[e_1]]\sigma = \{o\} & \mathcal{E}[[e_2]]\sigma = os \\ \text{typed}_{\Gamma}(o, f, os) & \text{update}(o, f, os, h) = h' \end{matrix}}{e_1.f := e_2, \sigma, \Gamma, h \Longrightarrow \sigma, \Gamma, h'}$$

$$\frac{\mathcal{B}[[b]]\sigma = \text{tt} \quad s_1, \sigma, \Gamma, h \Longrightarrow \sigma', \Gamma', h'}{\text{if } b \text{ then } s_1 \text{ else } s_2, \sigma, \Gamma, h \Longrightarrow \sigma', \Gamma', h'}$$

$$\frac{\mathcal{B}[[b]]\sigma = \text{ff} \quad s_2, \sigma, \Gamma, h \Longrightarrow \sigma', \Gamma', h'}{\text{if } b \text{ then } s_1 \text{ else } s_2, \sigma, \Gamma, h \Longrightarrow \sigma', \Gamma', h'}$$

$$\frac{\begin{matrix} s, \sigma, \Gamma, h \Longrightarrow \sigma', \Gamma', h' \\ \mathcal{E}[[e]]\sigma = \mathcal{E}[[e]]\sigma' \end{matrix}}{\text{fix } e \text{ do } s, \sigma, \Gamma, h \Longrightarrow \sigma', \Gamma', h'}$$

$$\frac{\begin{matrix} s, \sigma, \Gamma, h \Longrightarrow \sigma'', \Gamma'', h'' \\ \mathcal{E}[[e]]\sigma \neq \mathcal{E}[[e]]\sigma'' \\ \text{fix } e \text{ do } s, \sigma'', \Gamma'', h'' \Longrightarrow \sigma', \Gamma', h' \end{matrix}}{\text{fix } e \text{ do } s, \sigma, \Gamma, h \Longrightarrow \sigma', \Gamma', h'}$$

$$\frac{\mathcal{M}[[me]](\sigma, \Gamma, h) = os \quad x \leftarrow os \vdash s, \sigma, \Gamma, h \xrightarrow{\text{each}} \sigma', \Gamma', h'}{\text{foreach } x \in me \text{ do } s, \sigma, \Gamma, h \Longrightarrow \sigma', \Gamma', h'}$$

$$\frac{}{x \leftarrow \emptyset \vdash s, \sigma, \Gamma, h \xrightarrow{\text{each}} \sigma, \Gamma, h}$$

$$\frac{\begin{matrix} s, \sigma[x \mapsto \{o\}], \Gamma, h \Longrightarrow \sigma'', \Gamma'', h'' \\ x \leftarrow os \vdash s, \sigma'', \Gamma'', h'' \xrightarrow{\text{each}} \sigma', \Gamma', h' \end{matrix}}{x \leftarrow \{o\} \uplus os \vdash s, \sigma, \Gamma, h \xrightarrow{\text{each}} \sigma', \Gamma', h'}$$

Figure 5.6: Concrete semantics for TRON (Cont.)

5.3 Symbolic Execution

We discuss the main design principles of our symbolic executor. Although, the technique has been developed for TRON, the design decisions were driven by the desire to handle the modeled high-level transformation language in general.

Representing Rich States Symbolically

Spatial Constraints. Transformations manipulate structured data, not just simple values, and so the symbolic states of our executor describe primarily the possible shapes of the memory heap. Following other symbolic executors for object-oriented languages [Khurshid et al., 2003], we use *spatial constraints* to restrict the shapes admitted by an execution path. These constraints are first order formulae restricting values that are pointed to by links. In the style of the Lazier# algorithm [Deng et al., 2012], we distinguish between two kinds of symbolic objects: *symbolic instances* and *symbolic references*.

A *symbolic instance* ($\widehat{o} \in \widehat{\text{Instance}}$) abstracts over a unique instance. Instances cannot alias, so two different symbolic instances always point to two different class instances in memory, even if they have the same type. A *symbolic reference* ($x, y \in \text{Symbol}$) points to a class instance that may be aliased by other reference symbols, and, indeed, by some symbolic instances. The separation of symbolic instances and symbolic references allows separating reasoning about the structural representation of data from aliasing by references. We can lazily reason about aliasing without committing pre-maturely to a particular concretization of the heap structure. This is particularly important for our symbolic executor, as it handles deep containment constraints, which are hard to reason about and are heavily affected by aliasing (more about deep containment constraints below).

In traditional symbolic execution [Khurshid et al., 2003], whenever a field is accessed, the executor branches to initialize it to a new symbolic instance, or to alias an existing symbolic instance. In contrast, the Lazier# algorithm, simply assigns a distinct symbolic reference to each fresh field access and aliasing is only explicitly treated if the substructure of that symbolic reference is further explored.

For objects created using `new`, we eagerly generate a new concrete instance and exclude it from aliasing with pre-existing symbolic references, as new objects cannot alias previously existing ones (assuming

correctness of the memory manager). To emphasize this in the rules below, we mark the explicitly created instances with a dagger (\hat{o}^\dagger).

Set Symbols. In addition to ordinary symbolic references, we introduce *set symbols*, in the style of Cox et al. [2015] ($X^?, Y^? \in \text{SetSymbol}$), which abstract over finite sets of instances with unknown cardinality. This addition may seem very simple at first, but is key for our symbolic executor: it allows us to range over sets without prematurely concretizing their cardinality, contained objects, or their structure and aliasing.

Set Expressions and Set Constraints. Set symbols can be combined using symbolic set expressions:

$$\widehat{\text{SetExpr}} \ni \hat{e} ::= X^? \mid \emptyset \mid \{x_0^?, \dots, x_n^?\} \mid \hat{e}_1 \cup \hat{e}_2 \mid \hat{e}_1 \cap \hat{e}_2 \mid \hat{e}_1 \setminus \hat{e}_2$$

The symbolic set expressions mimic the set expressions of TRON, presented in Section 5.2, but without match expressions and with support for literal set constructors over simple symbolic references $\{x_1^?, \dots, x_n^?\}$. The meaning of the latter is a set of a fixed cardinality n , whose all elements are distinct (so, as a side effect, it also precludes aliasing between symbols listed). We use it to concretize the cardinality and content of sets during iteration.

Set expressions are embedded into constraints in a standard manner, using subset and equality constraints:

$$\widehat{\text{BoolExpr}} \ni \hat{b} ::= \hat{e}_1 \subseteq \hat{e}_2 \mid \hat{e}_1 = \hat{e}_2 \mid \neg \hat{b} \mid \hat{b}_1 \wedge \hat{b}_2$$

During symbolic execution the set comprehensions and reference symbols interplay to our benefit, allowing to describe assumptions about sets more lazily. For example, consider the constraint that equates two sets of cardinality 3 of unknown references: $\{x_1^?, x_2^?, x_3^?\} = \{y_1^?, y_2^?, y_3^?\}$. Generating this constraint allows to avoid deciding prematurely, which of the six possible aliasing configurations between x_i s and y_j s we are seeing, something which would not scale if done repeatedly.

Containment Constraints. A special feature of our symbolic executor is its ability to reason about the deep containment constraints of the manipulated data structures, which are extremely common in language processing (abstract syntax trees) and in data modeling. Besides eliminating many false positives, reasoning about containment also allows implementing deep matching.

To model deep containment constraints, we define a containment relation as the union of all links typed by containment fields, and insist that, for any two objects, their containments sets (the transitive closure of the containment relation) are disjoint. Furthermore, we enforce the acyclicity of the containment relation, ensuring that the irreflexive transitive closure of containment does not contain the identity pair for any object. We are using a solver (KodKod) that allows reasoning about transitive closures.

In order to perform symbolic deep matching, we on-demand bind a symbolic *containment set* for each instance \hat{o} used in a deep-matching query against a type c , to keep track of descendants. For example, this allows us to track all instances representing program variables for a some symbolic instance \hat{o} representing a program expression. This set is further constrained and unfolded during execution, in order to maintain sound access to instances that are contained by \hat{o} .

Type Constraints. We introduce type approximation in our symbolic executor, in order to not concretize types of instances (objects) prematurely. Many transformation rules operate on data constrained by types with inheritance, so the actual type of parameters might be unknown during symbolic execution. We maintain a bounding constraint on types, and refine it during execution by-need during branching and concretization cycles.

A type constraint environment ($\hat{\Gamma}$) maps each symbolic reference, set symbol, and each, symbolic instance \hat{o} to a type bound (cs_{in}, cs_{ex}) . The bound restricts the types of concrete values assignable to a symbol in question. A type bound is a tuple (cs_{in}, cs_{ex}) where the first component cs_{in} is a set of classes that specify the possible supertypes of a symbol and the second component cs_{ex} is a set of classes that specify excluded supertypes of a symbol, e.g. $(\{Expr\}, \{ThisExpr, FieldAccessExpr\})$ represents all subtypes of Expr (expressions) that are not subtypes of ThisExpr or FieldAccessExpr ('this' and field access expressions).

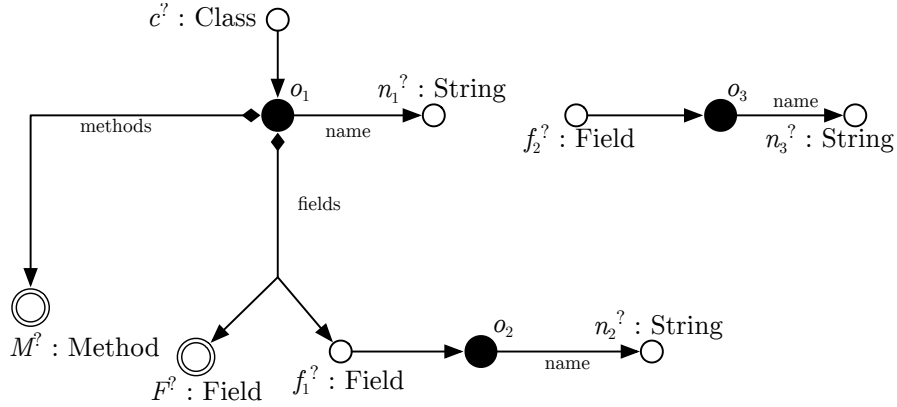
We only consider type bounds (cs_{in}, cs_{ex}) that are *well-formed*. That is: (i) the set of possible super-types cs_{in} cannot be empty, and (ii) none of the super-types is excluded: there is no class $c \in cs_{in}$ which is a subtype of an excluded super-type $c' \in cs_{ex}$. We also maintain an invariant that the set of possible super-types cs_{in} given by Γ is a singleton for symbolic instances. We will simply write c as a shorthand for $(\{c\}, \{c' \mid c' \text{ gen } c\})$ when the type of an element is precisely known.

Symbolic Heaps. A *symbolic heap* combines all types of constraints discussed above to describe possible concrete heaps and typings that could have been created during the execution. We define a symbolic heap to be a tuple $(\widehat{z}, \widehat{\ell}, \widehat{d}, \widehat{\Gamma}, \widehat{b})$, where $\widehat{z} \in \widehat{\text{Symbol}} \rightarrow \widehat{\text{Instance}}$ is a symbolic reference environment—partial mapping of symbolic references to symbolic instances that they are constrained to point to; $\widehat{\ell} \in \widehat{\text{Instance}} \times \text{Field} \rightarrow \widehat{\text{SetExpr}}$ collects the symbolic instances, by mapping fields of symbolic instances to symbolic set expressions; \widehat{d} is an environment storing deep containment constraints $\widehat{d} \in \widehat{\text{Instance}} \times \text{Class} \rightarrow \widehat{\text{SetExpr}}$ for all symbolic instances, $\widehat{\Gamma}$ is the type constraint environment, and \widehat{b} is the path constraint so far, in the execution leading to this symbolic heap.

An example symbolic heap is presented in Figure 5.7 using both the above syntax and a diagram². Dot vertices (\bullet) denote symbolic instances, white circles (\circ) denote symbolic references and large double-stroked white circles (\odot) denote symbolic reference sets.

Satisfiability of Symbolic Heaps We say that \widehat{h} is *satisfiable* if there is at least a pair of a concrete heap h and type environment Γ that are consistent with the constraints present in \widehat{h} . To check the consistency of the concrete heap h and type environment Γ against the constraints in \widehat{h} ($h, \Gamma \models^m \widehat{h}$) we need a model $m \in (\text{Symbol} \rightarrow \text{Instance}) \cup (\text{SetSymbol} \rightarrow \wp(\text{Instance}))$ which assigns to each symbolic reference a concrete instance, and to each symbolic reference set a concrete set of instances; we assume that symbolic instances are mapped directly one-to-one to concrete instances. The symbolic reference environment \widehat{z} is satisfied by model m , if m is an extension of \widehat{z} , i.e., $\forall x \in \text{dom } \widehat{z}. \widehat{z}(x) = m(x)$. The symbolic shape environment $\widehat{\ell}$ is consistent with heap h , if they agree on the structure of all defined links given the model m , i.e. $\forall (o, f) \in \text{dom } \widehat{\ell}. m(\widehat{\ell}(o, f)) = h(o, f)$; here the application of m to set expressions is extended to work by replacing all sub symbolic references and symbolic reference sets in the set expression with their value in m . The heap h and type environment Γ are consistent with the deep containment constraints \widehat{d} if \widehat{d} capture all necessary descendants for each class c for a particular instance

²We use our own diagram notation for objects instead of the UML one because it is more compact and allows us to neatly represent non-standard concepts like symbolic values and containment.



Symbolic references $\hat{z} = [c^? \mapsto \hat{o}_1, f_1^? \mapsto \hat{o}_2, f_2^? \mapsto \hat{o}_3]$

Symbolic instances $\hat{\ell} = [(\hat{o}_1, \text{name}) \mapsto n_1^?, (\hat{o}_1, \text{methods}) \mapsto M^?, (\hat{o}_1, \text{fields}) \mapsto F^? \uplus \{f_1^?\}, (\hat{o}_2, \text{name}) \mapsto n_2^?, (\hat{o}_3, \text{name}) \mapsto n_3^?]$

Containment constraints $\hat{d} = []$ (not accumulated yet)

Type constraints $\hat{\Gamma} = [\hat{o}_1 \mapsto \text{Class}, \hat{o}_2 \mapsto \text{Field}, \hat{o}_3 \mapsto \text{Field}, c^? \mapsto \text{Class}, n_1^? \mapsto \text{String}, f_1^? \mapsto \text{Field}, n_2^? \mapsto \text{String}, f_2^? \mapsto \text{Field}, n_3^? \mapsto \text{String}, M^? \mapsto \text{Method}, F^? \mapsto \text{Field}]$

Path condition $\hat{b} = \text{true}$ (not accumulated yet)

Figure 5.7: An example heap for an initial state of an execution

o , i.e. $m(\hat{d}(o, c)) = \{o' \mid o \text{ owns}_h^+ o' \wedge \Gamma(o') \text{ gen}^* c\}$, where for two instances o, o' then $o \text{ owns}_h o'$ iff there exists exactly one containment field $\hat{f} \in \text{Field}_\blacklozenge$ such that $h(o, \hat{f}) = o'$.

The symbolic type environment $\hat{\Gamma}$ is consistent with the concrete type environment Γ if each symbolic expression (symbolic reference, symbolic reference set or symbolic instance) has a type bound that is consistent with the types assigned in Γ given mapping m ; a type bound $(cs_{\text{in}}, cs_{\text{ex}})$ is consistent with a type c if there exist a $c' \in cs_{\text{in}}$ such that c is a subtype of c' ($c \text{ gen}^* c'$) and there doesn't exist a $c'' \in cs_{\text{ex}}$ which c is a subtype of ($\neg(c \text{ gen}^* c'')$). Finally, logical constraints in \hat{b} are consistent with the model m if the expression $m(\hat{b})$ we get by substituting all symbols in \hat{b} with m is true.

One symbolic heap is *stronger* than the other if all models (here all satisfying concrete heaps) of the former are also models of the latter. For conciseness, we let $(\hat{z}, \hat{\ell}, \hat{d}, \hat{\Gamma}, \hat{b}_1) \wedge \hat{b}_2$ mean $(\hat{z}, \hat{\ell}, \hat{d}, \hat{\Gamma}, \hat{b}_1 \wedge \hat{b}_2)$.




Figure 5.8: One of the paths when symbolically executing the example Rename-Field refactoring, starting with the symbolic state presented in Figure 5.7. Double-stroked arrows represent deep containment constraints.

Manipulating Symbolic State During Execution

Figure 5.8 shows an example path of the symbolic executor when executing the Rename-Field refactoring from Figure 5.2b starting with the symbolic state presented in Figure 5.7. The execution proceeds in the following steps:

- The initial statement on lines 9-10 replaces the old field (represented by symbol $f_1^?$) with the new field ($f_2^?$), such that the ‘fields’ reference of the target class ($c^?$) now points at $f_2^?$ instead of $f_1^?$. Remark, that changes are highlighted in red to ease understanding.
- Then we perform a deep matching on line 11, prompting the symbolic executor to create a deep containment constraint with type

FieldAccessExpr—represented by a double-stroked arrow (\Rightarrow)—for the location assigned to $c^?$; the containment constraint points to a symbolic set reference $FE_0^?$ (not shown in Figure) abstracting the set of all concrete instances of type FieldAccessExpr reachable from the concrete instance corresponding to $c^?$.

- In order to iterate over the elements of $FE_0^?$, we non-deterministically chose to partition it into disjoint symbol $fe^?$ and symbolic set $FE_1^?$, executing the body of the **foreach** with *faexpr* assigned to $fe^?$.
- To check the condition of the **if**-statement at lines 7-8, we perform a couple of field accesses, which triggers lazy initialization to creates two new symbolic instances: one which is assigned to the symbolic reference $fe^?$ and one which is assigned to its *target* field.
- We non-deterministically chose to execute the **then** branch—further constraining the values of the fields of $fe^?$ —executing the field update statement at line 9, which updates the field access expression to point at the new renamed field instances.
- Finally, we are ready for another iteration at line 6, and this time non-deterministically chose to stop, further constraining $FE_1^?$ to be \emptyset (thus disappearing in the figure).

We shall now define how these and other execution steps are realized. We start discussing the *basics* of the presentation format and the simple rules. Then we proceed to the four major ideas in our symbolic executor: *lazy initialization* during heap access and modification, *containment handling* when updating containment links, *lazy iteration* in **foreach**-loops, and *deep containment constraints* for handling matching expressions.

Basics. During the execution we maintain a store σ mapping variable names to symbolic set expressions. We use two symbolic evaluation functions for TRON's set ($\widehat{\mathcal{E}}[[e]]\widehat{\sigma} = \widehat{e}$) and Boolean expressions ($\widehat{\mathcal{B}}[[b]]\widehat{\sigma} = \widehat{b}$). They take concrete expressions with a store, and return resulting symbolic expressions by syntactically substituting all variables with their symbolic values as defined by σ . For example, we have $\widehat{\mathcal{B}}[[x \subseteq y]] [x \mapsto \{x^?\} \cup \{z^?\}, y \mapsto Y^?] = \{x^?\} \cup \{z^?\} \subseteq Y^?$.

The main judgement has the following format: $s, \widehat{\sigma}, \widehat{h} \longrightarrow \widehat{\sigma}', \widehat{h}'$, denoting that the statement s evaluated in the symbolic store $\widehat{\sigma}$ and heap

$$\begin{array}{c}
\text{SKIP} \frac{}{\text{skip}, \hat{\sigma}, \hat{h} \longrightarrow \hat{\sigma}, \hat{h}} \qquad \text{SEQ} \frac{s_1, \hat{\sigma}, \hat{h} \longrightarrow \hat{\sigma}'', \hat{h}'' \quad s_2, \hat{\sigma}'', \hat{h}'' \longrightarrow \hat{\sigma}', \hat{h}'}{s_1; s_2, \hat{\sigma}, \hat{h} \longrightarrow \hat{\sigma}', \hat{h}'} \\
\\
\text{AGN} \frac{}{x := e, \hat{\sigma}, \hat{h} \longrightarrow \hat{\sigma}[x \mapsto \hat{\mathcal{E}}[e]\hat{\sigma}], \hat{h}} \\
\\
\text{IFT} \frac{s_1, \hat{\sigma}, \hat{h} \wedge \hat{\mathcal{B}}[b]\hat{\sigma} \longrightarrow \hat{\sigma}', \hat{h}'}{\text{if } b \text{ then } s_1 \text{ else } s_2, \hat{\sigma}, \hat{h} \longrightarrow \hat{\sigma}', \hat{h}'} \qquad \text{IFF} \frac{s_2, \hat{\sigma}, \hat{h} \wedge \neg \hat{\mathcal{B}}[b]\hat{\sigma} \longrightarrow \hat{\sigma}', \hat{h}'}{\text{if } b \text{ then } s_1 \text{ else } s_2, \hat{\sigma}, \hat{h} \longrightarrow \hat{\sigma}', \hat{h}'} \\
\\
\text{FIXS} \frac{s, \hat{\sigma}, \hat{h} \longrightarrow \hat{\sigma}', \hat{h}'' \quad \hat{h}' = \hat{h}'' \wedge (\hat{\mathcal{E}}[e]\hat{\sigma} = \hat{\mathcal{E}}[e]\hat{\sigma}')}{\text{fix e do } s, \hat{\sigma}, \hat{h} \longrightarrow \hat{\sigma}', \hat{h}'} \\
\\
\text{FIXD} \frac{\hat{h}'' = \hat{h}''' \wedge (\hat{\mathcal{E}}[e]\hat{\sigma} \neq \hat{\mathcal{E}}[e]\hat{\sigma}'') \quad s, \hat{\sigma}, \hat{h} \longrightarrow \hat{\sigma}'', \hat{h}''' \quad \text{fix e do } s, \hat{\sigma}'', \hat{h}'' \longrightarrow \hat{\sigma}', \hat{h}'}{\text{fix e do } s, \hat{\sigma}, \hat{h} \longrightarrow \hat{\sigma}', \hat{h}'} \\
\\
\text{NEW} \frac{x^?, \hat{\sigma}^+ \text{ fresh} \quad \hat{\sigma}' = \hat{\sigma}[x \mapsto \{x^?\}] \quad \hat{z}' = \hat{z}[x^? \mapsto \hat{\sigma}^+] \quad \hat{\Gamma}' = \hat{\Gamma}[x^? \mapsto c, \hat{\sigma}^+ \mapsto c] \quad \hat{\ell}' = \hat{\ell}[(\hat{\sigma}^+, f) \mapsto \emptyset \mid f \in \text{fields}(c)]}{x := \text{new } c, \hat{\sigma}, (\hat{z}, \hat{\ell}, \hat{d}, \hat{\Gamma}, \hat{b}) \longrightarrow \hat{\sigma}', (\hat{z}', \hat{\ell}', \hat{d}, \hat{\Gamma}', \hat{b})}
\end{array}$$

Figure 5.9: Symbolic execution rules for standard statements (σ is a symbolic variable store, h is a symbolic heap)

\hat{h} , produces a new symbolic store $\hat{\sigma}'$ and heap \hat{h}' . The basic symbolic execution rules for statements are shown in Figure 5.9, including assignments, sequencing, branching, object creation and the **fix**-loop. These steps are essentially the same as in other existing symbolic executors. For straight-line statements the branch condition b remains unchanged during execution (SKIP, AGN, SEQ, NEW). For branching statements it is amended (cf. IFT, IFF, FIXS and FIXD). Any execution continues as long as the heap is satisfiable, so satisfiability of \hat{h} is an implicit premise in all the rules. The branching rules are also non-deterministic; the non-determinism corresponds to branching (back-tracking) in the symbolic

executor. Finally, the **NEW** rule, creates a new object of type c by allocating a symbolic instance and a symbolic reference $x^?$ pointing to it. All fields of the new instance are initialized to be empty sets.

Loops with unbounded iteration count give rise to infinite paths in symbolic execution (for instance by considering larger and larger inputs). In order for the symbolic execution algorithm to terminate, we bound the number of paths to be explored: **FixD** can only be applied a bounded number of times in a given execution of a given loop.

Lazy Initialization and Field Access. The rule for symbolic execution of field access is as follows:

$$\text{Acc} \frac{\widehat{\text{singleton}}(\widehat{\mathcal{E}}[[e]]\widehat{\sigma}, \widehat{h}) \ni (x^?, \widehat{h}'') \quad \widehat{\text{inst}}(x^?, \widehat{h}'') \ni (\widehat{o}, \widehat{h}') \quad \widehat{h}' = (\widehat{z}', \widehat{\ell}', \widehat{d}', \widehat{\Gamma}', \widehat{b}')}{x := e.f, \widehat{\sigma}, \widehat{h} \longrightarrow \widehat{\sigma}[x \mapsto \widehat{\ell}'(o, f)], \widehat{h}'}$$

Symbolically executing a field access $x := e.f$ requires the following steps:

1. Symbolically evaluating e to a symbolic set expression \widehat{e} .
2. Using $\widehat{\text{singleton}}$ function³ to get a single symbol $x^?$ representing the value of \widehat{e} .
 - a) If \widehat{e} is not already a single symbol, then the $\widehat{\text{singleton}}$ function will generate a fresh symbol $x^?$ with the correct type and add the constraint $\widehat{e} = x^?$ to the heap, returning a new heap if satisfiable.
3. Lazily assigning a symbolic instance \widehat{o} to $x^?$ —if not already assigned—using the $\widehat{\text{inst}}$ function, which non-deterministically either creates a new symbolic instance \widehat{o} with the right type and shape, or picks an existing symbolic instance \widehat{o} with compatible type bounds to treat aliasing.
4. Looking up the value of f of the assigned symbolic instance \widehat{o} in the spatial part of the heap $\widehat{\ell}'$ assigning the resulting value to variable x .

³All our auxiliary functions are formally defined in Appendix B.

Containment and Field Updates. The symbolic execution of a field update statement $e_1.f := e_2$ follows a similar pattern to field access:

$$\text{UPD} \frac{\widehat{\text{singleton}}(\widehat{\mathcal{E}}[\![e_1]\!] \widehat{\sigma}, \widehat{h}) \ni (x^?, \widehat{h}''') \quad \widehat{\text{inst}}(x^?, \widehat{h}''') \ni (\widehat{\sigma}, \widehat{h}'') \quad \widehat{\text{update}}(\widehat{\sigma}, f, \widehat{\mathcal{E}}[\![e_2]\!] \widehat{\sigma}, \widehat{h}'') = \widehat{h}'}{e_1.f := e_2, \widehat{\sigma}, \widehat{h} \longrightarrow \widehat{\sigma}, \widehat{h}'}}$$

After evaluating and resolving e to a symbolic instance $\widehat{\sigma}$, the $\widehat{\text{update}}$ function is used to update field f of $\widehat{\sigma}$ to point to the evaluated value of e_2 in the spatial constraints:

$$\widehat{\text{update}}(\widehat{\sigma}, f, \widehat{e}, (\widehat{z}, \widehat{\ell}, \widehat{d}, \widehat{\Gamma}, \widehat{b})) = \begin{cases} (\widehat{z}, \widehat{\ell}[(\widehat{\sigma}, f) \mapsto \widehat{e}], \widehat{d}, \widehat{\Gamma}, \widehat{b}) & \text{if } f \in \text{Field}_{\dots} \\ (\widehat{z}, \widehat{\ell}'[(\widehat{\sigma}, f) \mapsto \widehat{e}], \widehat{d}', \widehat{\Gamma}', \widehat{b} \wedge \widehat{b}') & \text{if } f \in \text{Field}_{\blacklozenge} \end{cases}$$

where $\widehat{\ell}' = \widehat{\text{disown}}(\widehat{e}, \widehat{\ell})$

$$(\widehat{d}', \widehat{\Gamma}', \widehat{b}') = \widehat{\text{dc-containment}}(\widehat{e}, c, \widehat{z}, \widehat{d}, \widehat{\Gamma})$$

If f is a containment field we must further ensure that $\widehat{\sigma}$ is the unique owner of \widehat{e} which $\widehat{\text{update}}$ does by calling $\widehat{\text{disown}}$ and $\widehat{\text{dc-containment}}$.

$$\widehat{\text{disown}}(\widehat{e}, \widehat{\ell}) = [(\widehat{\sigma}, f) \mapsto \widehat{\text{do-f}}(\widehat{e}', \widehat{e}) \mid (\widehat{\sigma}, f, \widehat{e}') \in \text{graph } \widehat{\ell}]$$

$$\text{where } \widehat{\text{do-f}}(\widehat{e}', \widehat{e}) = \begin{cases} \widehat{e}' & \text{if } f \in \text{Fields}_{\dots} \\ \widehat{e}' \setminus \widehat{e} & \text{if } f \in \text{Fields}_{\blacklozenge} \end{cases}$$

The $\widehat{\text{disown}}$ function presented above modifies each containment link in the spatial constraints $\widehat{\ell}$ to exclude the target symbolic expression \widehat{e} . The $\widehat{\text{dc-containment}}$ function analogously first excludes \widehat{e} from all deep containment constraints to ensure that there are no stale references to the values of \widehat{e} , and then tries to correctly propagate the effects of the assignment of \widehat{e} back to the containment constraints:

- For every containment constraint $\widehat{d}(\widehat{\sigma}, c) = \widehat{e}'$ with the same type as \widehat{e} or a super-type of it, we generate a new set symbol $X^?$ with the target type, replace \widehat{e}' with it and then add the constraint $X^? = \widehat{e}' \vee X^? = \widehat{e} \cup \widehat{e}'$ to the heap, which signifies that \widehat{e} might have been added to the deep containment constraints of $\widehat{\sigma}$; this highlights an interesting interaction between containment links and deep containment constraints which is not immediately obvious, but is necessary to maintain consistency while still keeping a high-level of symbolic abstraction.

- For subtypes of c , we do almost the same but use the constraint $(X^? = \hat{e}' \vee X^? = \hat{e}' \cup Y^?) \wedge \hat{e} = Y^? \uplus Z^?$ instead, where $Y^?$ and $Z^?$ are fresh set symbols with $Y^?$ having type c and $Z^?$ having the type of \hat{e} excluding \hat{c} (in a type bound); this ensures that we refer to all elements in \hat{e} of type c and only those.

Lazy Iteration with First-class Set Expressions. We introduce *lazy iteration* over first-class symbolic set expressions. In particular, consider the operational rule for the **foreach**-loop below:

$$\text{FOR} \frac{\widehat{\text{init}}(me, \hat{h}) = (e, \hat{\zeta}) \quad x \leftarrow \hat{\mathcal{E}}[[e]]\hat{\sigma} \vdash s, \hat{\sigma}, \hat{h}, \hat{\zeta} \xrightarrow{\text{each}} \hat{\sigma}', \hat{h}', \hat{\zeta}'}{\text{foreach } x \in me \text{ do } s, \hat{\sigma}, \hat{h} \longrightarrow \hat{\sigma}', \hat{h}'}$$

The first step is to use the $\widehat{\text{init}}$ function to get the expression e to be iterated over, and initialize a control state $\hat{\zeta}$ used during iteration depending on the kind of matching expression me provided; we will treat $\hat{\zeta}$ abstractly for now, and define it precisely later in this section. The other step is to use the $\xrightarrow{\text{each}}$ -judgement to iterate over the values of \hat{e} depending on $\hat{\zeta}$, executing the **foreach**-body s at each iteration. The two rules for the $\xrightarrow{\text{each}}$ -judgement, are provided below:

$$\text{FORB} \frac{\widehat{\text{next}}(\hat{e}, \hat{h}, \hat{\zeta}) \ni (\text{break}, \hat{h}')}{x \leftarrow \hat{e} \vdash s, \hat{\sigma}, \hat{h}, \hat{\zeta} \xrightarrow{\text{each}} \hat{\sigma}, \hat{h}', \hat{\zeta}}$$

$$\text{FORC} \frac{\widehat{\text{next}}(\hat{e}, \hat{h}, \hat{\zeta}) \ni (\text{cont}(x^?, \hat{e}', \hat{\zeta}''), \hat{h}''') \quad s, \hat{\sigma}[x \mapsto x^?], \hat{h}''' \longrightarrow \hat{\sigma}'', \hat{h}'' \quad x \leftarrow \hat{e}' \vdash s, \hat{\sigma}'', \hat{h}'', \hat{\zeta}'' \xrightarrow{\text{each}} \hat{\sigma}', \hat{h}', \hat{\zeta}'}{x \leftarrow \hat{e} \vdash s, \hat{\sigma}, \hat{h}, \hat{\zeta} \xrightarrow{\text{each}} \hat{\sigma}', \hat{h}', \hat{\zeta}'}$$

Both of the above rules depend on the $\widehat{\text{next}}$ function which given target expression \hat{e} , current control state $\hat{\zeta}$ and heap \hat{h} , provides a set of possible next actions; a possible action \hat{a} is either **break** which signals that iteration should stop, or is $\text{cont}(x^?, \hat{e}', \hat{\zeta}')$ which signals that another iteration should happen with symbol $x^?$ bound to the range variable, \hat{e}' (disjoint from $x^?$) representing the rest of the values to be iterated over, and new control state $\hat{\zeta}'$. The first rule of the $\xrightarrow{\text{each}}$ -judgement check whether **break** is a possible next state and if so it will stop iteration with new heap \hat{h}' . The second rule checks whether **cont** is a possible next state, executes the loop body s with $x^?$ bound to the range variable x , and finally continues iteration over \hat{e}' in the updated states.

Now, observe how laziness is achieved with two key ideas: we never explicitly concretize me , leaving the level of concretization required to be decided by the $\widehat{\text{next}}$ function according to the control state $\widehat{\zeta}$, and we iterate using a symbolic reference $x^?$ without requiring an assignment of a symbolic instance (to treat possible aliasing) as this point. Furthermore, by parameterizing the rules for **foreach** over functions $\widehat{\text{init}}$ and $\widehat{\text{next}}$, it would be easy to add new kinds of expressions without affecting the rules.

Type-Directed Matching with Containment Constraints. We will now discuss how the control state $\widehat{\zeta}$ and functions $\widehat{\text{init}}$ and $\widehat{\text{next}}$ interact with matching expressions. We define the control state as follows:

$$\zeta^? ::= \text{ns} \mid \text{ms}(c) \mid \text{ms}^*(c, \widehat{e}, \widehat{d})$$

Each alternative control state contains the required components to execute a given matching expression: ns is used for ordinary iteration, $\text{ms}(c)$ is used for shallow matching of elements against c and $\text{ms}^*(c, \widehat{e}, \widehat{d})$ is used for deep matching of elements against type c , a set of additional elements that must be iterated \widehat{e} and a copy of the deep containment constraints \widehat{d} ; the copy of containment constraints is kept in order to retrieve the deep containment constraint values that were available before iteration, which would represent the concrete objects that would have been matched by a concrete deep match operation at the time of evaluation. The $\widehat{\text{init}}$ function is therefore defined to map each expression to its initial control state:

$$\begin{aligned} \widehat{\text{init}}(e, \widehat{h}) &= (e, \text{ns}) & \widehat{\text{init}}(e \text{ \textbf{match}} c, \widehat{h}) &= (e, \text{ms}(c)) \\ \widehat{\text{init}}(e \text{ \textbf{match}}^* c, (\widehat{z}, \widehat{\ell}, \widehat{d}, \widehat{\Gamma}, \widehat{b})) &= (e, \text{ms}^*(c, \widehat{\mathcal{O}}, \widehat{d})) \end{aligned}$$

The $\widehat{\text{next}}$ function is more interesting since it calculates the possible next actions for iteration. For sake of readability, we will define the $\widehat{\text{next}}$ function using inference rules, which describe the possible values in the result set. We define two rules for ordinary iteration:

$$\begin{aligned} \text{ITER-EMP} & \frac{}{\widehat{\text{next}}(\widehat{e}, \widehat{h}, \text{ns}) \ni (\text{break}, (\widehat{h} \wedge \widehat{e} = \emptyset))} \\ \text{ITER-CONT} & \frac{\widehat{\text{partition}}(\widehat{e}, \widehat{h}) = (x^?, X^?, \widehat{h}')}{\widehat{\text{next}}(\widehat{e}, \widehat{h}, \text{ns}) \ni (\text{cont}(x^?, X^?, \widehat{\text{ns}}), \widehat{h}')} \end{aligned}$$

There are two possible actions: we can stop iterating if it is possible to constraint \widehat{e} to be \emptyset (ITER-EMP), and we can try to use $\widehat{\text{partition}}$ to split \widehat{e} into a symbol $x^?$, and a disjoint set symbol $X^?$ which continue iteration with (ITER-CONT). The $\widehat{\text{partition}}$ function essentially generates fresh symbol $x^?$ and set symbol $X^?$ with the right types adding the constraint $\widehat{e} = x^? \uplus X^?$ to \widehat{h} .

For matching iterations, the rules for $\widehat{\text{next}}$ are defined as follows:

$$\begin{array}{c} \text{ITERM-EMP} \frac{}{\widehat{\text{next}}(\widehat{e}, \widehat{h}, \text{ms}(c)) \ni (\text{break}, (\widehat{h} \wedge \widehat{e} = \emptyset))} \\ \text{ITERM-SUCS} \frac{\widehat{\text{partition}}(\widehat{e}, \widehat{h}) = (x^?, X^?, \widehat{h}'') \quad \widehat{\text{match}}(x^?, X^?, c, \widehat{h}'') \ni (\text{tt}, \widehat{h}')}{\widehat{\text{next}}(\widehat{e}, \widehat{h}, \text{ms}(c)) \ni (\text{cont}(x^?, X^?, \text{ms}(c)), \widehat{h}')} \\ \text{ITERM-FAIL} \frac{\widehat{\text{partition}}(\widehat{e}, \widehat{h}) = (x^?, X^?, \widehat{h}'') \quad \widehat{\text{match}}(x^?, X^?, c, \widehat{h}'') \ni (\text{ff}, \widehat{h}')}{\widehat{\text{next}}(\widehat{e}, \widehat{h}, \text{ms}(c)) \ni (\text{break}, \widehat{h}')} \end{array}$$

The terminating rule (ITERM-EMP) is essentially unchanged, and the other two continuation rules use $\widehat{\text{partition}}$ to divide the target symbolic expression \widehat{e} into a symbol $x^?$ and a symbolic reference set $X^?$. The symbol $x^?$ is then matched against $c^?$ using $\widehat{\text{match}}$ which returns a set of states each indicating whether matching $x^?$ against c was successful or not; if the match is successful (ITERM-SUCS)— $(\text{tt}, \widehat{h}')$ is included in the result—then \widehat{h}' constraints the type of $x^?$ to be a subtype of c , and otherwise if the match failed (ITERM-FAIL)—so $(\text{ff}, \widehat{h}')$ is included in the result—then \widehat{h}' constraints the type bounds of $x^?$ and $X^?$ to exclude c as a possible supertype. A match is always successful if the type of $x^?$ is a subtype of c , always fails when the c is unrelated to or excluded from type bounds of $x^?$, and allows both when the type of $x^?$ is a supertype of c .

Finally, the rule for $\widehat{\text{next}}$ for deep matching delegates to an auxiliary definition $\widehat{\text{next-ms}}$ (ITERMS):

$$\text{ITERMS} \frac{\widehat{\text{next-ms}}(\widehat{e}, c, \widehat{e}', \widehat{d}, \widehat{h}) \ni (\widehat{a}, \widehat{h}')}{\widehat{\text{next}}(\widehat{e}, \widehat{h}, \text{ms}^*(c, \widehat{e}', \widehat{d})) \ni (\widehat{a}, \widehat{h}')}$$

There are three possible rules for $\widehat{\text{next-ms}}$, which are defined as follows:

$$\text{ITERMSA-EMP} \frac{\widehat{\text{next}}(\widehat{e}', \widehat{h} \wedge \widehat{e} = \emptyset, \text{ns}) \ni (\widehat{a}, \widehat{h}')}{\widehat{\text{next-ms}}(\widehat{e}, c, \widehat{e}', \widehat{d}, \widehat{h}) \ni (\widehat{a}, \widehat{h}')}$$

$$\text{ITERMSA-FAIL} \frac{\widehat{\text{partition}}(\widehat{e}, \widehat{h}) = (x^?, X^?, \widehat{h}''') \quad \widehat{\text{match}}(x^?, X^?, c, \widehat{h}''') \ni (\text{ff}, \widehat{h}''')}{\widehat{\text{dcs}}(x^?, c, \widehat{d}, \widehat{h}''') \ni (\widehat{e}'', \widehat{h}'') \quad \widehat{\text{next-ms}}(X^?, \widehat{e}' \cup \widehat{e}'', \widehat{h}'') \ni (\widehat{a}, \widehat{h}')} \widehat{\text{next-ms}}(\widehat{e}, c, \widehat{e}', \widehat{d}, \widehat{h}) \ni (\widehat{a}, \widehat{h}')$$

$$\text{ITERMSA-SUCS} \frac{\widehat{\text{partition}}(\widehat{e}, \widehat{h}) = (x^?, X^?, \widehat{h}''')}{\widehat{\text{match}}(x^?, X^?, c, \widehat{h}''') \ni (\text{tt}, \widehat{h}'') \quad \widehat{\text{dcs}}(x^?, c, \widehat{d}, \widehat{h}'') \ni (\widehat{e}'', \widehat{h}'')} \widehat{\text{next-ms}}(\widehat{e}, c, \widehat{e}', \widehat{d}, \widehat{h}) \ni (\text{cont}(x^?, X^?, \text{ms}^*(c, \widehat{e}' \cup \widehat{e}'', \widehat{d})), \widehat{h}')$$

The first rule constraints $\widehat{e} \in \emptyset$, and continues iterating over the symbolic set \widehat{e}' which is used by the other rules to collect deep containment constraint values during iteration (ITERMSA-EMP).

The second rule uses $\widehat{\text{partition}}$ and $\widehat{\text{match}}$ on c getting an unsuccessful match (ITERMSA-FAIL); we use the $\widehat{\text{dcs}}$ function to assign a location \widehat{d} to $x^?$, lookup the deep containment constraint value $\widehat{d}(\widehat{d}, c) = \widehat{e}''$ —creating it in the provided heap if non-existing—and adding \widehat{e}'' to the control state. The rule then continues iterating over the rest $X^?$, since $x^?$ did not match the target type c and so must be skipped.

The third and final rule gets a successful match (ITERMSA-SUCS) and so we use the $x^?$ for the next iteration; we still need to consider possible descendants of $x^?$ of type c and so use $\widehat{\text{dcs}}$ to get the deep containment constraint value and add it to the control state. Observe how the use of deep containment constraints allows us to provide a higher-level abstraction over structures focusing only on instances of the target type, and without explicitly considering all intermediate shapes of data.

Relating Concrete and Symbolic Semantics

To recover a deterministic semantics for programs from our provided non-deterministic operational symbolic semantics, one could define the symbolic semantics of a program as the set of all output pairs of symbolic stores and heaps obtained from non-deterministically executing each different feasible path in the program.

$$\widehat{S}[\![s]\!](\widehat{\sigma}, \widehat{h}) = \left\{ (\widehat{\sigma}', \widehat{h}') \mid s, \widehat{\sigma}, \widehat{h} \longrightarrow \widehat{\sigma}', \widehat{h}' \right\}$$

A useful property to show is then that the deterministic symbolic semantics is sound with regards to the concrete semantics, i.e.:

Theorem 5.1 (Soundness). *If $\exists m.\sigma = m(\hat{\sigma}) \wedge \Gamma, h \stackrel{m}{\models} \hat{h}$ and $\widehat{S}[\![s]\!](\hat{\sigma}, \hat{h}) = \widehat{\mathbf{M}}$ then for all $(\hat{\sigma}, \hat{h}) \in \widehat{\mathbf{M}}$ there exists σ', Γ' , and h' such that we have a concrete execution $s, \sigma, \Gamma, h \Longrightarrow \sigma', \Gamma', h'$ and exists a model m' such that $\sigma' = m'(\hat{\sigma}')$ and $\Gamma', h' \stackrel{m'}{\models} \hat{h}'$.*

5.4 Evaluation

We implemented our technique in a prototype tool,⁴ which we evaluate to show the concrete benefits of our technique.

Test Generation

White-box test generation is a classical application of symbolic execution, and we aim to use it as an example for evaluating our symbolic execution algorithm. We have built a white-box test generator and compare its effectiveness against a baseline black-box test generator.

We aim to compare the test generators according to their effectiveness, which is how well a test suite exercises the transformation-under-test (TUT). We use branch coverage as the target metric, which we define for TRON constructs as follows: for `if`-statements both branches must be taken, for `fix`-loops we check whether it is run one or more times⁵ and for `foreach`-loops we check whether it is run zero, one or more times.

White-box Test Generator. The white-box test generator is a simple extension of the symbolic execution algorithm presented in Section 5.3, requiring only two new additions:

1. Memoising a copy of the spatial constraints whose values are not modified by field updates, thus keeping track of the initial structure of input.
2. A translator between the output model given by the model finder to concrete data usable by the target TRON program.

⁴<https://github.com/itu-square/SymexTRON>

⁵`fix`-loops must run the body at least once, according to the semantics

Black-box Test Generator. The black-box test generator optimizes towards *meta-model coverage*, which is the literature-recommended metric [Wang et al., 2006; Finot et al., 2013]. A test suite is said to have full meta-model coverage if each subtype of relevant classes is present in at least in one test case, and each relevant field is instantiated with each valid multiplicity (i.e. zero, one or many).

Subject Programs. The subject programs were selected according to three criteria: be an interesting representative variety of realistic transformations, be independently specified to avoid bias, and be feasible to implement in TRON. To fulfill the first criterion, we chose transformations from two categories: model transformations and refactorings. For the second criterion, we ported the model transformations from the ATL transformation zoo⁶ and chose the refactorings from Fowler’s classic collection [Fowler, 1999]. The third criterion is achieved by picking suitably sized transformations that satisfy our resource and design constraints, since it takes time to manually port complex transformations correctly (despite language similarity) and TRON lacks abstractions for modularity that full languages have. We ended up with 3 model transformations and 4 refactorings, all of which we describe below.

Refactorings.

- An extended version of the *Rename field* refactoring used as the running example (see Figure 5.2).
- *Rename method*: renames a target method in a class, and ensures that all calls to the correct overloading of this method (with the right types) must refer to the updated name.
- *Extract superclass*: creates a common superclass of two classes with similar structure ensuring that all common fields are pulled up and that both classes inherit from it.
- *Replace delegation with inheritance*: Makes a class inherit directly from a type instead of using a field for delegation, updating all method calls targeting that field to use **this** as a target instead.

⁶<https://www.eclipse.org/atl/atlTransformations/>

Model Transformations.

- The *Families to Persons* model transformation (Fam2Pers), converts a model of a traditional family with a mother, father and possibly children to a collection of individuals with explicit gender (male or female).
- The classical *Class to Relational* model transformation (Class2Rel), which converts an object-oriented class model to a relational database schema.
- The *Path expression to Petri net* model transformation (Path2Petri), which converts a path expression with states and transitions to a full Petri net with named places, different types of arcs and weighted transitions.

The characteristics and baseline meta-model coverage is presented in Table 5.2, and the source code of all the programs is available in Appendix D.

Porting Transformations to TRON. By design TRON is a minimal language, and so there are non-core transformation languages features that must be handled when porting transformations from fully-featured languages to TRON. In particular, three features had to be handled for the considered subject programs: functions, implicit tracing links and circular data dependencies (the latter two present in model transformation languages like ATL). When a transformation is ported to TRON one must take care to correctly inline function calls, which is done by replacing the calls with the function body, substituting the parameters with the provided arguments, renaming local variables to avoid clashes, and converting any explicit recursion to use the ‘**foreach**’-statement or ‘**fix**’-statement. To handle tracing links one must take care to augment the meta-model to include them explicitly, and to explicate assignment to the tracing links on object creation in the transformation; for circular data dependencies one must ensure to initialize all relevant objects before the transformation.

Set-up. The experiment was set-up to automatically run both test generators automatically on all the described subject programs. The white-box test generator was bounded in the number of iterations (2, except

Program	LOC	Meta-model coverage (%)	
		Black-box	White-box
RenameField	53 + 12 = 65	96.55	70.69
RenameMethod	53 + 28 = 81	96.55	93.10
ExtractSuper	53 + 29 = 82	98.28	31.03
ReplaceDelegation	53 + 30 = 83	100.00	85.19
Fam2Pers	21 + 56 = 77	100.00	88.00
Path2Petri	42 + 58 = 100	100.00	37.50
Class2Rel	34 + 100 = 134	100.00	100.00

Table 5.2: Characteristics and baseline meta-model coverage for the subject programs. Here LOC indicates lines of code, where the first component of the summation is the size of the data model and the second component is the size of the transformation.

for Fam2Pers which uses 3 due to meta-model constraints) and instances considered by the model finder (6 for model transformations, 10 for refactorings), and a time-out of 1 hour was put in place. We also added light-weight support for bidirectional fields in the model finder to better support the model transformations. The prototype symbolic executor, was implemented in Scala 2.11.7 [Odersky and Rompf, 2014] and the evaluation was run on a 2.3 GHz Core i7 MacBook Pro (OS X 10.11). The external model finder KodKod was configured to use the parallel SAT solver Plingeling [Biere, 2014].

Test Generation Results. We ran a series of toy programs exercising the various constructs of TRON as a warm up for our test generators; the white-box test generator achieved 100% code coverage for all programs beating the black-box test generator, and all under 30 seconds of execution time.

Table 5.3 shows the results of running the test generators on the selected subject programs (model transformations and refactorings).

Refactorings. The white-box test generator achieves better code coverage than the baseline black-box test generator for all the refactorings, reaching 100% coverage for two. We hypothesise that refactorings do many targeted modifications of complex models, making it hard to

Program	Branch coverage (%)		Time (s)	
	Black-box	White-box	Black-box	White-box
RenameField	60.00	100.00	141.5	336.6
RenameMethod	26.67	93.33	141.7	3600.0
ExtractSuper	75.00	100.00	124.8	386.4
ReplaceDelegation	47.06	76.47	115.5	3600.0
Fam2Pers	100.00	100.00	4.8	135.4
Path2Petri	88.89	33.33	1.8	3600.0
Class2Rel	70.83	75.00	3.8	3600.0

Table 5.3: Results of running the test generators on subject programs. The branch coverage indicates the percentage of the total reachable branches covered, and the time indicates the execution time of the tool from start to finish on the evaluation subjects.

generate tests that cover the required paths without access to the transformation code; The test generated by black-box had high meta-model coverage (see Table 5.2) as expected, but achieved low branch coverage since it did not fully exercise the transformation (Table 5.3); this is in contrast to the white-box test generator which generated more focused test cases.

Due to the nature of symbolic execution, the white-box test generator was slower than the black-box test generator. We believe that the higher precision in bug finding offsets the runtime cost; test generation is an occasional offline task and 1 computer hour is not unreasonable to use compared to the many hours a programmer would otherwise have spent on the same task.

Model Transformations. The white-box test generator achieved good results for model transformations, performing better than the black-box test generator for the Fam2Pers and Class2Rel transformations, and worse for the Path2Petri transformation. We suspect that the black-box test generator performs well on model transformations because they primarily translate structured data according to their meta-model types, without relying on complex constraints and cases. Therefore, having a test suite with high meta-model coverage will generate the necessary different types to trigger the right execution paths resulting in acceptable branch coverage. The white-box test generator performed

less impressively on model transformations than refactorings because the symbolic executor did not reach the right paths before the time-out triggered. The sequential composition of complex for-loops results in an explosion of paths to be explored, and so it takes significantly more time to explore the whole program and generate interesting tests.

Comparing Coverage Criteria. The black-box test generator optimizes towards achieving maximal meta-model coverage (see Table 5.3), but seems to achieve mixed branch coverage results, and although it performed better for model transformations than for refactorings, it never reached full branch coverage. This indicates that there is little guarantee that high meta-model coverage ensures high branch coverage, which is an interesting experience for building both white-box and black-box tools.

The symbolic executor achieved high branch coverage for most subjects without achieving the same high meta-model coverage—the lowest being 31.03% meta-model coverage for the ExtractSuper refactoring that we achieved full branch coverage for—which indicates that there is no correlation the other way as well. It would of course require a more extensive empirical study to conclusively affirm our hypothesis.

Comparison with Symbolic Executors for Object-Oriented Languages

Two of the best known and well-supported symbolic executors for object-oriented programming languages are Symbolic PathFinder [Pasareanu et al., 2013] and Microsoft IntelliTest/Pex [Tillmann and de Halleux, 2008]. We will describe our experiences trying to encode various high-level transformation features—sets, containment and deep matching—and the difficulties faced when these features are not handled first-class.

Symbolic Support for Sets. The first challenge we faced was how to symbolically encode sets in traditional symbolic executors. Symbolic PathFinder does not directly support standard Java collections like HashSet or TreeSet, and using those collections during symbolic execution leads to errors: a symbol does not have a hash value, and it is not possible to symbolically compare two objects. One could try a more cumbersome encoding by using lists and handling inequality constraints explicitly, but it is unclear how to generate interesting instances of such

sets automatically. Pex tries to dynamically construct instances of sets using arrays, but it is hard in practice to make it generate an array of distinct elements that is usable to construct an interesting set.

We treat set values first-class which exploits the support of set theories in model finders like KodKod and SMT solvers [Kröning et al., 2009]. We believe that implementing this could be beneficial for these traditional symbolic executors as well.

Enforcing Deep Containment Constraints. Another challenge is enforcing containment-like constraints with acyclicity and non-sharing for abstract syntax trees. Traditional symbolic executors support deep containment constraints neither directly nor indirectly. Hypothetically, acyclicity constraints could be enforced statically by giving all objects unique identifiers and fixing an ordering between contained objects and parents; however, this requires both the management of a complex system on top of existing dynamic structures, and it is unclear how to efficiently handle dynamic to such links. In contrast, first class support of these constraints in TRON makes it easy to handle dynamic updates and allows specialized techniques to be used to handle such constraints.

Deep Matching and Visitors. Deep matching is straightforward to encode using reflection, but reflection is not handled well by symbolic executors, and so is to be avoided.

Traditionally traversal of abstract syntax is done using recursive visitors, which uses plain classes and thus is better supported. However, this approach is non-optimal from a symbolic execution point of view, since to reach the relevant part of an abstract syntax tree—like a field access expression—one has to consider and enumerate all intermediate shapes—i.e., classes, methods, different kinds of statements and containing expressions in our example—which hits a combinatorial explosion, even with reasonably small bounds.

In contrast, we abstract away intermediate shapes with deep containment constraints, which allows reasoning about only the parts of the data structure we are interested in. Transformations like refactorings often perform local changes on the abstract syntax trees, and so this approach seems especially beneficial in those cases.

Threats and Limitations

The main threat to validity of the experiment is that we implemented the subject programs ourselves, introducing the possibility of bias and errors in the implementation. For the model transformations, we mitigated this by choosing existing ones from ATL, and for the refactorings we chose a number of standard ones from Fowler’s authoritative book [Fowler, 1999]. Furthermore, minor implementation mistakes are of lesser importance since the number of found errors is not an evaluation criterion.

Inozemtseva and Holmes [2014] show that test coverage is not a strongly correlative measure for effectiveness. However, arguably a test suite which has low code coverage is going to miss bugs because it simply does not visit code present in some of the branches. The black-box test generator has been implemented by us optimizing for the standard meta-model coverage metric [Finot et al., 2013; Wang et al., 2006] to avoid bias, since we could not find an existing third-party tool that was suitable for our purposes.

We are not experts on Symbolic PathFinder and Pex, and could have missed better ways to encode high-level features. We mitigated this by systematically reading the available documentation, and searching on forums and mailing lists for answers to similar challenges.

5.5 Related Tools and Approaches

Symbolic Execution of High-level Transformation Languages. Simple symbolic execution algorithms [Lucio and Vangheluwe, 2013] exist for significantly less expressive transformation languages like DSLTrans, which bounds loops and does not permit dependent state and loop iterations. This lack of expressiveness allows the symbolic executor to be heavily specialized and quick, but the algorithm is hard to generalize for more expressive Turing-complete languages like TRON.

More complex whitebox-based algorithms are presented by ATL-Test [González and Cabot, 2012] and TETRA Box [Schönböck et al., 2013]. These tools only support a class of transformations that can not modify input state. Therefore, it is not possible to easily express complex transformations like the refactorings considered in this paper. Furthermore, the method presented in this paper, is fully-formalized and

evaluated, showing applicability of our framework for a broader range of transformations.

Test Generation for Transformations. The latest survey on verification of model transformations [Rahim and Whittle, 2015] shows that most test generation techniques for model transformations focus on black-box testing, which do not account for concrete transformation semantics and thus may fail to cover program statements as shown in our evaluation. There is a test generation tool for Maude [Riesco, 2010, 2012] based on bounded narrowing (practically, symbolic execution for rewriting languages). However, complex transformations are hard to write in the style of term-rewriting systems—since they modify object graphs—and the object-oriented extension is not as far as we understand supported by the test generator. A black-box test generation tool called Dolly [Mongiovi et al., 2014] is used to test C and Java based refactoring engines with promising results. As our evaluation results indicate that white-box based techniques have better effectiveness than black-box based ones, it could be interesting to see whether we could adapt some of our novel ideas for a language like Java and increase the number of bugs found.

Dynamic Symbolic Execution Our evaluation shows that it is relatively cheap to generate random initial test input with acceptable coverage. We believe that it might be fruitful to look at hybrid white-box approaches like Dynamic Symbolic Execution [Korel, 1990; Godefroid et al., 2005] to increase performance and scalability for larger transformations.

5.6 Recap

This chapter presented a formal symbolic execution algorithm for a small formal transformation language TRON capturing key high-level transformation features from Chapter 2. In order to feasibly model the features symbolically without hitting a combinatorial explosion, we had to combine ideas from existing work, i.e., lazy initialization and symbolic reference sets, with novel ideas that were targeted towards the modelled high-level features, i.e. containment handling, lazy iteration and deep containment constraints. The symbolic technique was implemented in a prototype white-box test generation tool, which compared

favorably against the baseline black-box test generator, achieving high test coverage for virtually all the selected subject transformation. This shows great promise in adapting programming language techniques to transformation languages, and made it possible for us in the later chapters to tackle a large subset of Rascal by using the techniques and experiences we gathered.

Chapter 6

The Formal Semantics of Rascal Light

Rascal [Klint et al., 2009, 2011] is a high-level transformation language that aims to simplify software language engineering tasks like defining program syntax, analyzing and transforming programs, and performing code generation. The language provides several features including built-in collections (lists, sets, maps), algebraic data-types, powerful pattern matching operations with backtracking, and high-level traversals supporting multiple strategies.

Interaction between different language features can be difficult to comprehend, since most features are semantically rich. My goal is to provide a well-defined formal semantics for a large subset of Rascal called Rascal Light—in the spirit of (core) Caml Light [Leroy, 1997], Clight [Blazy and Leroy, 2009], and Middleweight Java [Bierman et al., 2003]—suitable for developing formal techniques, e.g., type systems and static analyses.

Scope

Rascal Light aims to model a realistic set of high-level transformation language features, by capturing a large subset of the Rascal operational language. Unlike the small formal transformation language TRON (Chapter 5), Rascal Light targets being practically usable, i.e., it should be possible to translate many realistic pure Rascal programs to this subset by an expert programmer without losing the high level of abstraction. The following Rascal features are captured in Rascal Light:

- Fundamental definitions including algebraic data-types, functions with type parameters, and global variables.

- Basic expressions, including variable assignment, definition and access, exceptions, collections (including sets, lists and maps), **if**-expressions, **switch**-statements, **for**-loops and **while**-loops including control flow operators.
- Powerful pattern matching operations including type patterns, non-linear pattern matching, (sub)collection patterns, and descendant pattern matching.¹
- Backtracking using the **fail** operator, including roll-back of state.
- Traversals using generic **visit**-expressions, supporting the different kinds of available strategies: **bottom-up(-break)**, **top-down(-break)**, **innermost**, and **outermost**.
- Fixed-point iteration using the **solve**-loop.

The following Rascal features are considered out of scope:

- Concrete syntax declaration and literals, string interpolation, regular expression matching, date literals and path literals
- The standard library including Input/Output and the Java foreign function interface
- The module system, include modular extension of language elements such as datatypes and functions
- Advanced type system features, like parametric polymorphism, the numerical hierarchy and type inference

In Rascal, Boolean expressions can contain pattern matching subexpressions and backtracking is affected by the various Boolean operators (conjunction, disjunction, implication). In Rascal Light, backtracking is more restricted, which I believe heavily simplifies the semantics while losing little expressiveness in practice; most programs could be rewritten to support the required backtracking at the cost of increased verbosity.

Method

I have used the language documentation,² and the open source implementation of Rascal,³ to derive the formalism. The syntax is primarily

¹Similar to type-directed querying in TRON

²<http://tutor.rascal-impl.org/Rascal/Rascal.html>

³<https://github.com/usethesource/rascal>

based on the μ Rascal syntax description [CWI Amsterdam, 2017], but altered to focus on the high-level features. The semantics is based on the description of individual language features in the documentation. In case of under-specification or ambiguity, I used small test programs to check what the expected behavior is. I thank the Rascal developers for our personal correspondence which further clarified ambiguity and limitations of the semantics compared to Rascal.

To discover possible issues, the semantics has been implemented as a prototype and tested against a series of Rascal programs. To strengthen the correctness claims, I have proven a series of theorems of interest in Section 6.3, since as Milner et al. [1990] states:

The robustness of the semantics depends upon theorems

— The Definition of Standard ML

6.1 Abstract Syntax

Rascal Light programs are organized into modules that consist of definitions of global variables, functions and algebraic data types. I will in the rest of this chapter assume that modules are well-formed: top-level definitions and function parameters have unique names, all function calls, constructors and datatypes used have corresponding definitions, and all variables are well-scoped with no shadowing. To maintain a clean presentation, I will not write the definition environment explicitly, but mention the necessary definitions as premises when required.

Rascal Light has three kinds of definitions: global variables, functions and algebraic data-types. Global variables are typed and are initialized with a target expression at module loading time. Functions have a return type, a unique name, a list of typed uniquely named parameters, and have an expression as a body. Algebraic data-types have unique names and declare a set of possible alternative constructors, each taking a list of typed fields as arguments.

$$\begin{array}{ll}
 d ::= \mathbf{global} \ t \ y = e & \text{(Global variables)} \\
 | \mathbf{fun} \ t \ f(\underline{t \ x}) = e & \text{(Function Definition)} \\
 | \mathbf{data} \ at = k_1(\underline{t \ x}) | \dots | k_n(\underline{t \ x}) & \text{(Algebraic Data Type Definition)}
 \end{array}$$

The operational part of Rascal consists of syntactically distinct categories of statements and expressions; I have chosen to collapse the two categories in Rascal Light since statements also have result values that can be used inside expressions. Most constructs are standard imperative ones, such as blocks, assignment, branching and loops (including control operations); I have chosen to use **local-in** for representing blocks in Rascal Light instead of curly braces like in Rascal, to distinguish them from set literal expressions; blocks contain locally declared variables and a sequence of expressions to evaluate.

Notable Rascal-specific constructs include a generalized version of the **switch**-expression that supports rich pattern matching over both basic values and algebraic data types, the **visit**-expression which allows traversing data types using various strategies (Example 6.1), and the **solve**-expression which continuously evaluates its body expression until target variables reach a fixed-point in assigned values (Example 6.2). The **fail** control operator allows backtracking inside **switch** and **visit** statements to try other possible matches (Example 6.3).

Example 6.1 (Expression simplifier). An expression simplifier can use **visit** to completely simplify all sub-expressions no matter where they are in the input expression:

```

1  data Expr = intlit(int v) | plus(Expr lop, Expr rop) | ...;
2
3  Expr simplify(Expr e) =
4    bottom-up visit(e) {
5      case plus(intlit(0), y) => y
6      case plus(x, intlit(0)) => x
7    };

```

▲

Example 6.2 (Lattice fixed-point). The Kleene fixed-point of a continuous function in a lattice for a domain *Val* with functions for the bottom element and least-upper bound can be computed using **solve**:

```

1  Val fix(Fun f) = {
2    Val v = bottom();
3    solve(v) {
4      v = lub(v, apply(f, v))
5    };
6  }

```

▲

Example 6.3 (Knapsack problem). The knapsack problem concerns finding a subset of items with greatest value under a specified maximum

weight. The function below uses backtracking to (slowly) find the optimal solution.

```

1  set[Item] slowknapsack(set[Item] items, int maxWeight) = {
2      set[Item] res = {};
3      solve(res) {
4          switch(items) {
5              case {*xs, *ys}:
6                  if (sumweights(xs) > maxWeight
7                      || sumvalues(xs) < sumvalues(res)) fail;
8                  else res = xs;
9              };
10         res
11     };
12 }

```



$e ::= vb$	(Basic Values, $vb \in \{1, \text{"foo"}, 3.14, \dots\}$)
x	(Variables, $x, y \in \text{Var}$)
$\ominus e$	(Unary Operations, $\ominus \in \{-, \dots\}$)
$e_1 \oplus e_2$	(Binary Operations, $\oplus \in \{+, -, \times, \dots\}$)
$k(\underline{e})$	(Constructor Applications, $k \in \text{Constructor}$)
$[\underline{e}]$	(List Expressions)
$\{ \underline{e} \}$	(Set Expressions)
$(\underline{e} : \underline{e}')$	(Map Expressions)
$e_1[e_2]$	(Map Lookup)
$e_1[e_2 = e_3]$	(Map Update)
$f(\underline{e})$	(Function Calls, $f \in \text{Function}$)
return e	(Return Expressions)
$x = e$	(Variable Update Assignments)
if e then e_1 else e_2	(Conditionals)
switch e do \underline{cs}	(Case Matching)
st visit e do \underline{cs}	(Deep Traversal)
break continue fail	(Control Operations)
local \underline{t} x in \underline{e} end	(Blocks)
for g e	(Iteration)
while e e	(While Expressions)
solve \underline{x} e	(Fixedpoint Expressions)
throw e	(Exception Throwing)

try e_1 catch $x \Rightarrow e_2$	<i>(Exception Catching)</i>
try e_1 finally e_2	<i>(Exception Finalization)</i>
$cs ::= \mathbf{case} \ p \Rightarrow e$	<i>(Case)</i>

There are two kinds of *generator expressions*: *enumerating assignment* where the range variable is assigned to each element in the result of a collection-producing expression, and *matching assignment* that produces all possible assignments that match target patterns (defined later).

$g ::= x \leftarrow e$	<i>(Enumerating Assignment)</i>
$p := e$	<i>(Matching Assignment)</i>

The **visit**-expression supports various strategies that determine the order a particular value is traversed w.r.t. its contained values. The **top-down** strategy traverses the value itself before contained values, and conversely **bottom-up** traverses contained values before itself. The **break** versions stop the traversal at first successful match, and the **outermost** and **innermost** respectively apply the **top-down** and **bottom-up** until a fixed-point is reached.

$st ::= \mathbf{top-down}$	<i>(Preorder Traversal)</i>
bottom-up	<i>(Postorder Traversal)</i>
top-down-break	<i>(First-match Preorder Traversal)</i>
bottom-up-break	<i>(First-match Postorder Traversal)</i>
outermost	<i>(Fixedpoint Preorder Traversal)</i>
innermost	<i>(Fixedpoint Postorder Traversal)</i>

Like Rascal, Rascal Light has a rich pattern language that not only includes matching on basic values and constructors, but also powerful matching inside collections and descendant patterns that allow matching arbitrarily deeply contained values.

$$\begin{array}{ll}
p ::= vb & \text{(Basic Value Patterns)} \\
| x & \text{(Variable Patterns)} \\
| k(\underline{p}) & \text{(Deconstructor Patterns)} \\
| t x : p & \text{(Typed Labelled Patterns)} \\
| [\underline{\star}p] & \text{(List Patterns)} \\
| \{\underline{\star}p\} & \text{(Set Patterns)} \\
| !p & \text{(Negation pattern)} \\
| /p & \text{(Descendant Patterns)}
\end{array}$$

Patterns inside collections can be either ordinary patterns or *star patterns* that match a subcollection with arbitrary number of elements.

$$\begin{array}{ll}
\star p ::= p & \text{(Ordinary Pattern)} \\
| \star x & \text{(Star Pattern)}
\end{array}$$

Rascal Light programs expressions evaluate to values, which either are basic values, constructor values, collections or the undefined value (■).

$$\begin{array}{ll}
v ::= vb & \text{(Basic Values)} \\
| k(\underline{v}) & \text{(Constructor Values, } k \in \text{Constructor)} \\
| [\underline{v}] & \text{(List Values)} \\
| \{\underline{v}\} & \text{(Set Values)} \\
| (\underline{v} : \underline{v}') & \text{(Map Values)} \\
| \blacksquare & \text{(Undefined Value)}
\end{array}$$

6.2 Semantics

I present a formal development of the dynamic aspects of Rascal Light, using a natural semantics specification. Natural (big-step) semantics [Kahn, 1987] is particularly suitable for Rascal, because it closely mimics semantics of an interpreter and the high-level features—exceptions, backtracking, traversals—introduce a rich control-flow that depends not only on the structure of the program but also on the provided input. There is no concurrency or interleaving in Rascal, and so there is no need for a more fine-grained operational semantics like (small-step) structural operational semantics (SOS) [Plotkin, 2004].

Value Typing

Rascal (Light) is strongly typed and so all values are typed. The types are fairly straight-forward in the sense that most values have a canonical type and there is a sub-typing hierarchy (explained shortly) with a bottom type **void** and a top type **value**.

$t ::= tb$	(Base Types, $tb \in \{\text{Int, Rational, String, } \dots\}$)
at	(Algebraic Data Types, $at \in \text{DataType}$)
set $\langle t \rangle$	(Sets)
list $\langle t \rangle$	(Lists)
map $\langle t_1, t_2 \rangle$	(Maps)
void	(Bottom Type)
value	(Top Type)

I provide a typing judgment for values of form $\boxed{v : t}$, which states that value v has type t . Basic values are typed by their defining basic type, the undefined value (\blacksquare) has the bottom type **void**, and constructor values are typed by their corresponding data type definition assuming the contained values are well-typed. The type of collections is determined by the contained values, and so a least upper bound operator is defined in types—following the sub-type ordering—which is used to infer a precise type for the value parameters.

$$\begin{array}{c}
 \text{T-BASIC} \frac{vb \in \llbracket tb \rrbracket}{vb : tb} \quad \text{T-BOT} \frac{}{\blacksquare : \mathbf{void}} \\
 \\
 \text{T-CONSTRUCTOR} \frac{\mathbf{data} \ at = \dots \mid k(\underline{t}) \mid \dots \quad \underline{v} : \underline{t}' \quad \underline{t}' <: \underline{t}}{k(v_1, \dots, v_n) : at} \\
 \\
 \text{T-SET} \frac{\underline{v} : \underline{t}}{\{\underline{v}\} : \mathbf{set}\langle \underline{t} \rangle} \quad \text{T-LIST} \frac{\underline{v} : \underline{t}}{[\underline{v}] : \mathbf{list}\langle \underline{t} \rangle} \quad \text{T-MAP} \frac{\underline{v} : \underline{t} \quad \underline{v}' : \underline{t}'}{(\underline{v} : \underline{v}') : \mathbf{map}\langle \underline{t}, \underline{t}' \rangle}
 \end{array}$$

The subtyping relation has form $\boxed{t <: t'}$, stating t is a subtype of t' . We let $\boxed{t \not<: t'}$ denote the negated form where none of the given cases below matches. Sub-typing is reflexive, so every type is a sub-type of itself; **void** and **value** act as bottom type and top type respectively.

$$\text{ST-REFL} \frac{}{t <: t} \quad \text{ST-VOID} \frac{}{\mathbf{void} <: t} \quad \text{ST-VALUE} \frac{}{t <: \mathbf{value}}$$

Collections are covariant in their type parameters, which is safe since all values are immutable.

$$\text{ST-LIST} \frac{t <: t'}{\mathbf{list}\langle t \rangle <: \mathbf{list}\langle t' \rangle} \quad \text{ST-SET} \frac{t <: t'}{\mathbf{set}\langle t \rangle <: \mathbf{set}\langle t' \rangle}$$

$$\text{ST-MAP} \frac{t_{\text{key}} <: t'_{\text{key}} \quad t_{\text{val}} <: t'_{\text{val}}}{\mathbf{map}\langle t_{\text{key}}, t_{\text{val}} \rangle <: \mathbf{map}\langle t'_{\text{key}}, t'_{\text{val}} \rangle}$$

The least upper bound on types is defined as follows:

$$\begin{aligned} \sqcup \varepsilon &= \mathbf{void} & \sqcup t, t' &= t \sqcup \sqcup t' \\ t_1 \sqcup t_2 &= \begin{cases} t_1 & \text{if } t_2 = \mathbf{void} \vee t_1 = t_2 \\ t_2 & \text{if } t_1 = \mathbf{void} \\ \mathbf{list}\langle t'_1 \sqcup t'_2 \rangle & \text{if } t_1 = \mathbf{list}\langle t'_1 \rangle \wedge t_2 = \mathbf{list}\langle t'_2 \rangle \\ \mathbf{set}\langle t'_1 \sqcup t'_2 \rangle & \text{if } t_1 = \mathbf{set}\langle t'_1 \rangle \wedge t_2 = \mathbf{set}\langle t'_2 \rangle \\ \mathbf{map}\langle t'_1 \sqcup t'_2, t''_1 \sqcup t''_2 \rangle & \text{if } t_1 = \mathbf{map}\langle t'_1, t''_1 \rangle \wedge t_2 = \mathbf{map}\langle t'_2, t''_2 \rangle \\ \mathbf{value} & \text{otherwise} \end{cases} \end{aligned}$$

Expression Evaluation

The main judgment for Rascal Light expressions has the form

$e; \sigma \xRightarrow{\text{expr}} vres; \sigma'$, where the expression e is evaluated in an initial store $\sigma \in \text{Var} \rightarrow \text{Val}$ —mapping variables to values—returning a result $vres$ and updated store σ' as a side-effect. The result $vres$ is either an ordinary **success** v signifying successful execution or an exceptional result $exres$.

$$vres ::= \mathbf{success} \ v \mid exres$$

An exceptional result is either a control operation (**break**, **continue**, **fail**), an **error** that happened during execution, a thrown exception **throw** v or an early return value **return** v . The difference between **success** v and **return** v is that the latter should propagate directly from sub-expressions to the surrounding function call boundary, while the value in the former can be used further in intermediate computations (Example 6.4).

$$exres ::= \mathbf{return} \ v \mid \mathbf{throw} \ v \mid \mathbf{break} \mid \mathbf{continue} \mid \mathbf{fail} \mid \mathbf{error}$$

Example 6.4 (Early Return). The function that calculates products, uses the early return functionality to short-circuit the rest of the calculation when a factor in the provided list is zero. If this branch is hit during execution, evaluating the expression produces **return** 0 as result, which is then propagated directly to function call boundary, skipping the rest of the loop and sequenced expressions.

```

1  int prod(list[int] xs) {
2    int res = 1;
3    for (x <- xs) {
4      if (x == 0) return 0;
5      else res *= x;
6    };
7    res
8  }

```

▲

Basic values evaluate to their semantic equivalent result values without any side effect (E-VAL).

$$\text{E-VAL} \frac{}{vb; \sigma \xRightarrow{\text{expr}} \mathbf{success} \, vb; \sigma}$$

A variable evaluates to the value it is assigned to in the store if available (E-VAR-SUCS), and otherwise result in an error (E-VAR-ERR).

$$\text{E-VAR-SUCS} \frac{x \in \text{dom } \sigma}{x; \sigma \xRightarrow{\text{expr}} \mathbf{success} \, \sigma(x); \sigma} \quad \text{E-VAR-ERR} \frac{x \notin \text{dom } \sigma}{x; \sigma \xRightarrow{\text{expr}} \mathbf{error}; \sigma}$$

Unary expressions evaluate their operands, applying possible side-effects; if successful the corresponding semantic unary operator $\llbracket \ominus \rrbracket$ is applied on the result value (E-UN-SUCS), and otherwise it propagates the exceptional result (E-UN-EXC).

$$\text{E-UN-SUCS} \frac{e; \sigma \xRightarrow{\text{expr}} \mathbf{success} \, v; \sigma'}{\ominus e; \sigma \xRightarrow{\text{expr}} \llbracket \ominus \rrbracket(v); \sigma'} \quad \text{E-UN-EXC} \frac{e; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'}{\ominus e; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'}$$

Example 6.5 (Unary operator semantics). $\llbracket - \rrbracket(3)$ will evaluate to **success** -3 , while $\llbracket - \rrbracket(\{\})$ will evaluate to **error** ▲

Evaluating binary expressions is similar to unary expressions, requiring both operands to evaluate successfully (E-BIN-SUCS) to apply the corresponding semantic binary operator $\llbracket \oplus \rrbracket$; otherwise the exceptional

results of the operands are propagated in order from left (E-BIN-Exc1) to right (E-BIN-Exc2).

$$\begin{array}{c}
 \text{E-BIN-SUCS} \frac{e_1; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ v_1; \sigma'' \quad e_2; \sigma'' \xRightarrow{\text{expr}} \mathbf{success} \ v_2; \sigma'}{e_1 \oplus e_2; \sigma \xRightarrow{\text{expr}} \llbracket \oplus \rrbracket (v_1, v_2); \sigma'} \\
 \\
 \text{E-BIN-Exc1} \frac{e_1; \sigma \xRightarrow{\text{expr}} \text{exres}_1; \sigma''}{e_1 \oplus e_2; \sigma \xRightarrow{\text{expr}} \text{exres}_1; \sigma'} \\
 \\
 \text{E-BIN-Exc2} \frac{e_1; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ v_1; \sigma'' \quad e_2; \sigma'' \xRightarrow{\text{expr}} \text{exres}_2; \sigma'}{e_1 \oplus e_2; \sigma \xRightarrow{\text{expr}} \text{exres}_2; \sigma'}
 \end{array}$$

Constructor expressions evaluate their arguments first, and if they all successfully evaluate to values, then check whether the types of values match those expected in the declaration. If the result values have the right types and are not \blacksquare , a constructor value is constructed (E-CONS-Sucs), and otherwise a (type) error is produced (E-CONS-ERR). In case any of the arguments has an exceptional result the evaluation of the rest of the arguments halts and the exceptional result is propagated (E-CONS-Exc).

$$\begin{array}{c}
 \text{E-CONS-SUCS} \frac{\mathbf{data} \ at = \dots \mid k(\underline{t}) \mid \dots \quad \underline{e}; \sigma \xRightarrow{\text{expr}^*} \mathbf{success} \ \underline{v}; \sigma' \quad \underline{v} : t' \quad \underline{v} \neq \blacksquare \quad t' <: t}{k(\underline{e}); \sigma \xRightarrow{\text{expr}} \mathbf{success} \ k(\underline{v}); \sigma'} \\
 \\
 \text{E-CONS-ERR} \frac{\mathbf{data} \ at = \dots \mid k(\underline{t}) \mid \dots \quad \underline{e}; \sigma \xRightarrow{\text{expr}^*} \mathbf{success} \ \underline{v}; \sigma' \quad \underline{v} : t' \quad \exists i. v_i = \blacksquare \vee t'_i \not<: t_i}{k(\underline{e}); \sigma \xRightarrow{\text{expr}} \mathbf{error}; \sigma'} \\
 \\
 \text{E-CONS-Exc} \frac{\underline{e}; \sigma \xRightarrow{\text{expr}^*} \text{exres}; \sigma'}{k(\underline{e}); \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'}
 \end{array}$$

Evaluating list expressions also requires evaluating all subexpressions to a series of values; because of sequencing and necessity of early propagation of exceptional results, evaluation of series of subexpressions is done using a mutually recursive sequence evaluation judgment (see page 115). If the evaluation is successful then a list value is constructed

(E-LIST-SUCS), unless any value is undefined (\blacksquare) in which case we produce an error (E-LIST-ERR), and otherwise the exceptional result is propagated (E-LIST-EXC).

$$\begin{array}{c}
 \text{E-LIST-SUCS} \frac{\underline{e}; \sigma \xRightarrow{\text{expr}^*} \mathbf{success} \underline{v}; \sigma' \quad \underline{v} \neq \blacksquare}{[\underline{e}]; \sigma \xRightarrow{\text{expr}} \mathbf{success} [\underline{v}]; \sigma'} \\
 \\
 \text{E-LIST-ERR} \frac{\underline{e}; \sigma \xRightarrow{\text{expr}^*} \mathbf{success} \underline{v}; \sigma' \quad \exists i. v_i = \blacksquare}{[\underline{e}]; \sigma \xRightarrow{\text{expr}} \mathbf{error}; \sigma'} \quad \text{E-LIST-EXC} \frac{\underline{e}; \sigma \xRightarrow{\text{expr}^*} \text{exres}; \sigma'}{[\underline{e}]; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'}
 \end{array}$$

Set expression evaluation mirror the one for lists, except that values are constructed using a set constructor (E-SET-SUCS), which may reorder values and ensures that there are no duplicates. If any contained value was undefined (\blacksquare) then an error is produced instead (E-SET-ERR), and exceptional results are propagated (E-SET-EXC).

$$\begin{array}{c}
 \text{E-SET-SUCS} \frac{\underline{e}; \sigma \xRightarrow{\text{expr}^*} \mathbf{success} \underline{v}; \sigma' \quad \underline{v} \neq \blacksquare}{\{\underline{e}\}; \sigma \xRightarrow{\text{expr}} \mathbf{success} \{\underline{v}\}; \sigma'} \\
 \\
 \text{E-SET-ERR} \frac{\underline{e}; \sigma \xRightarrow{\text{expr}^*} \mathbf{success} \underline{v}; \sigma' \quad \exists i. v_i = \blacksquare}{\{\underline{e}\}; \sigma \xRightarrow{\text{expr}} \mathbf{error}; \sigma'} \quad \text{E-SET-EXC} \frac{\underline{e}; \sigma \xRightarrow{\text{expr}^*} \text{exres}; \sigma'}{\{\underline{e}\}; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'}
 \end{array}$$

Map expressions evaluate their keys and values in the declaration sequence, and if successful construct a map (E-MAP-SUCS). Similarly to other collection expressions, errors are produced if any value is undefined (E-MAP-ERR) and exceptional results are propagated (E-MAP-EXC).

$$\begin{array}{c}
 \text{E-MAP-SUCS} \frac{\underline{e}, \underline{e}'; \sigma \xRightarrow{\text{expr}^*} \mathbf{success} \underline{v}, \underline{v}'; \sigma' \quad \underline{v} \neq \blacksquare \quad \underline{v}' \neq \blacksquare}{(\underline{e} : \underline{e}'); \sigma \xRightarrow{\text{expr}} \mathbf{success} (\underline{v} : \underline{v}'); \sigma'} \\
 \\
 \text{E-MAP-ERR} \frac{\underline{e}, \underline{e}'; \sigma \xRightarrow{\text{expr}^*} \mathbf{success} \underline{v}, \underline{v}'; \sigma' \quad \exists i. v_i = \blacksquare \vee v'_i = \blacksquare}{(\underline{e} : \underline{e}'); \sigma \xRightarrow{\text{expr}} \mathbf{error}; \sigma'} \\
 \\
 \text{E-MAP-EXC} \frac{\underline{e}, \underline{e}'; \sigma \xRightarrow{\text{expr}^*} \text{exres}; \sigma'}{(\underline{e} : \underline{e}'); \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'}
 \end{array}$$

Lookup expressions require evaluating the outer expression to a map—otherwise producing an error (E-LOOKUP-ERR)—and the index expression to a value. If the index is an existing key then the corresponding

value is produced as result (E-LOOKUP-SUCS) and otherwise the nokey exception is thrown (E-LOOKUP-NOKEY); here, I assume that

data NoKey = nokey(value key)

is a built-in data-type definition. Exceptional results are propagated from sub-terms (E-LOOKUP-Exc1, E-LOOKUP-Exc2).

$$\begin{array}{c}
 \begin{array}{c}
 e_1; \sigma \xRightarrow{\text{expr}} \mathbf{success} (\dots, v : v', \dots); \sigma'' \\
 e_2; \sigma'' \xRightarrow{\text{expr}} \mathbf{success} v; \sigma' \\
 \hline
 \text{E-LOOKUP-SUCS} \quad e_1[e_2]; \sigma \xRightarrow{\text{expr}} \mathbf{success} v'; \sigma'
 \end{array} \\
 \\
 \begin{array}{c}
 e_1; \sigma \xRightarrow{\text{expr}} \mathbf{success} (v : v'); \sigma'' \\
 e_2; \sigma'' \xRightarrow{\text{expr}} \mathbf{success} v''; \sigma' \quad \forall i. v'' \neq v_i \\
 \hline
 \text{E-LOOKUP-NOKEY} \quad e_1[e_2]; \sigma \xRightarrow{\text{expr}} \mathbf{throw} \text{nokey}(v''); \sigma'
 \end{array} \\
 \\
 \begin{array}{c}
 e_1; \sigma \xRightarrow{\text{expr}} \mathbf{success} v; \sigma' \quad v \neq (v' : v'') \\
 \hline
 \text{E-LOOKUP-ERR} \quad e_1[e_2]; \sigma \xRightarrow{\text{expr}} \mathbf{error}; \sigma'
 \end{array} \\
 \\
 \begin{array}{c}
 e_1; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma' \\
 \hline
 \text{E-LOOKUP-Exc1} \quad e_1[e_2]; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'
 \end{array} \\
 \\
 \begin{array}{c}
 e_1; \sigma \xRightarrow{\text{expr}} \mathbf{success} (v : v'); \sigma'' \quad e_2; \sigma'' \xRightarrow{\text{expr}} \text{exres}; \sigma' \\
 \hline
 \text{E-LOOKUP-Exc2} \quad e_1[e_2]; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'
 \end{array}
 \end{array}$$

Map update expressions also require the outer expression to evaluate to a map—otherwise producing an error (E-UPDATE-ERR1)—and the index and target expressions to evaluate to values. On successful evaluation of both index and target value the map is updated, overriding the old value of the corresponding index if necessary (E-UPDATE-SUCS) unless the index or target value is equal to \blacksquare in which case an error is produced (E-UPDATE-ERR2). Finally, exceptional results are propagated left-to-right if necessary (E-UPDATE-Exc1, E-UPDATE-Exc2, E-UPDATE-Exc3).

$$\begin{array}{c}
 \begin{array}{c}
 e_1; \sigma \xRightarrow{\text{expr}} \mathbf{success} (v : v'); \sigma''' \quad e_2; \sigma''' \xRightarrow{\text{expr}} \mathbf{success} v''; \sigma'' \\
 e_3; \sigma'' \xRightarrow{\text{expr}} \mathbf{success} v'''; \sigma' \quad v'' \neq \blacksquare \quad v''' \neq \blacksquare \\
 \hline
 \text{E-UPDATE-SUCS} \quad e_1[e_2 = e_3]; \sigma \xRightarrow{\text{expr}} \mathbf{success} (v : v', v'' : v'''); \sigma'
 \end{array}
 \end{array}$$

$$\begin{array}{c}
\text{E-UPDATE-ERR1} \frac{e_1; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ v; \sigma' \quad v \neq (v' : v'')}{e_1[e_2 = e_3]; \sigma \xRightarrow{\text{expr}} \mathbf{error}; \sigma'} \\
\\
\text{E-UPDATE-ERR2} \frac{\begin{array}{cc} e_1; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ (v : v'); \sigma''' & e_2; \sigma''' \xRightarrow{\text{expr}} \mathbf{success} \ v''; \sigma'' \\ e_3; \sigma'' \xRightarrow{\text{expr}} \mathbf{success} \ v''; \sigma' & v'' = \blacksquare \vee v''' = \blacksquare \end{array}}{e_1[e_2 = e_3]; \sigma \xRightarrow{\text{expr}} \mathbf{error}; \sigma'} \\
\\
\text{E-UPDATE-EXC1} \frac{e_1; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'}{e_1[e_2 = e_3]; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'} \\
\\
\text{E-UPDATE-EXC2} \frac{e_1; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ (v : v'); \sigma'' \quad e_2; \sigma'' \xRightarrow{\text{expr}} \text{exres}; \sigma'}{e_1[e_2 = e_3]; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'} \\
\\
\text{E-UPDATE-EXC3} \frac{\begin{array}{cc} e_1; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ (v : v'); \sigma'' & \\ e_2; \sigma''' \xRightarrow{\text{expr}} \mathbf{success} \ v''; \sigma'' & e_3; \sigma'' \xRightarrow{\text{expr}} \text{exres}; \sigma' \end{array}}{e_1[e_2 = e_3]; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'}
\end{array}$$

Evaluation of function calls is more elaborate. Function definitions are statically scoped, and the semantics is eager, so arguments are evaluated using call-by-value. The initial step is thus to evaluate all the arguments to values if possible—propagating the exceptional result otherwise (E-CALL-ARG-EXC)—and then check whether the values have the right type (otherwise producing an error, E-CALL-ARG-ERR). The evaluation proceeds by evaluating the body of the function with a fresh store that contains the values of global variables and the parameters bound to their respective argument values. There are then four cases:

1. If the body successfully evaluates to a correctly typed value, then that value is provided as the result (E-CALL-SUCS).
2. If the body evaluates to a value that does not have the expected type, then it produces an error (E-CALL-RES-ERR1).
3. If the result is a thrown exception or error, then it is propagated (E-CALL-RES-EXC).

4. Otherwise if the result is a control operator, then an error is produced (E-CALL-RES-ERR2).

In all cases the resulting store of executing the body is discarded—since local assignments fall out of scope—except the global variable values which are added to the store that was there before the function call.

$$\text{E-CALL-SUCS} \frac{\frac{\text{global } t_y y \quad \text{fun } t' f(t x) = e'}{e; \sigma \xRightarrow{\text{expr}^*} \text{success } \underline{v}; \sigma''} \quad \frac{v : t'' \quad t'' <: t}{[y \mapsto \sigma''(y), x \mapsto v]; e' \xRightarrow{\text{expr}} vres; \sigma'}}{vres = \text{return } v' \vee vres = \text{success } v' \quad v' : t''' \quad t''' <: t'}{f(\underline{e}); \sigma \xRightarrow{\text{expr}} \text{success } v'; \sigma'' [y \mapsto \sigma'(y)]}$$

$$\text{E-CALL-ARG-ERR} \frac{\text{fun } t' f(t x) = e' \quad e; \sigma \xRightarrow{\text{expr}^*} \text{success } \underline{v}; \sigma' \quad v : t'' \quad t'' \not<: t_i}{f(\underline{e}); \sigma \xRightarrow{\text{expr}} \text{error}; \sigma'}$$

$$\text{E-CALL-ARG-EXC} \frac{e; \sigma \xRightarrow{\text{expr}^*} \text{exres}; \sigma' \quad \text{exres} = \text{throw } v'}{f(\underline{e}); \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'}$$

$$\text{E-CALL-RES-EXC} \frac{\frac{\text{global } t_y y \quad \text{fun } t' f(t x) = e'}{e; \sigma \xRightarrow{\text{expr}^*} \text{success } \underline{v}; \sigma''} \quad \frac{v : t'' \quad t'' <: t}{[y \mapsto \sigma''(y), x \mapsto v]; e' \xRightarrow{\text{expr}} \text{exres}; \sigma' \quad \text{exres} = \text{throw } v'}}{f(\underline{e}); \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'' [y \mapsto \sigma'(y)]}$$

$$\text{E-CALL-RES-ERR1} \frac{\frac{\text{global } t_y y \quad \text{fun } t' f(t x) = e'}{e; \sigma \xRightarrow{\text{expr}^*} \text{success } \underline{v}; \sigma''} \quad \frac{v : t'' \quad t'' <: t}{[y \mapsto \sigma''(y), x \mapsto v]; e' \xRightarrow{\text{expr}} vres; \sigma'} \quad vres = \text{return } v' \vee vres = \text{success } v' \quad v' : t''' \quad t''' \not<: t'}{f(\underline{e}); \sigma \xRightarrow{\text{expr}} \text{error}; \sigma'' [y \mapsto \sigma'(y)]}$$

$$\text{E-CALL-RES-ERR2} \frac{\frac{\text{global } t_y y \quad \text{fun } t' f(t x) = e'}{e; \sigma \xRightarrow{\text{expr}^*} \text{success } \underline{v}; \sigma''} \quad \frac{v : t'' \quad t'' <: t}{[y \mapsto \sigma''(y), x \mapsto v]; e' \xRightarrow{\text{expr}} \text{exres}; \sigma'} \quad \text{exres} \in \{\text{break, continue, fail, error}\}}{f(\underline{e}); \sigma \xRightarrow{\text{expr}} \text{error}; \sigma'' [y \mapsto \sigma'(y)]}$$

The **return**-expression evaluates its argument expression first, and if it successfully produces a value then the result would be an early return with that value (E-RET-SUCS); recall that early returns are treated as exceptional values and so propagated through most evaluation rules, except at function call boundaries (rules E-CALL-SUCS and E-CALL-RES-ERR1). Otherwise, the exceptional result is propagated (E-RET-EXC).

$$\begin{array}{c} \text{E-RET-SUCS} \frac{e; \sigma \xRightarrow{\text{expr}} \mathbf{success} v; \sigma'}{\mathbf{return} e; \sigma \xRightarrow{\text{expr}} \mathbf{return} v; \sigma'} \quad \text{E-RET-EXC} \frac{e; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'}{\mathbf{return} e; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'} \end{array}$$

In Rascal Light, variables must be declared before being assigned, and declarations are unique since shadowing is disallowed. Evaluating an assignment proceeds by evaluating the right-hand side expression—propagating exceptional results (E-ASGN-EXC)—and then checking whether the produced value is compatible with the declared type. If it is compatible then the store is updated (E-ASGN-SUCS), and otherwise an error is produced (E-ASGN-ERR).

$$\begin{array}{c} \text{E-ASGN-SUCS} \frac{\mathbf{local} t x \vee \mathbf{global} t x \quad e; \sigma \xRightarrow{\text{expr}} \mathbf{success} v; \sigma' \quad v : t' \quad t' <: t}{x = e; \sigma \xRightarrow{\text{expr}} \mathbf{success} v; \sigma' [x \mapsto v]} \\ \text{E-ASGN-ERR} \frac{\mathbf{local} t x \vee \mathbf{global} t x \quad e; \sigma \xRightarrow{\text{expr}} \mathbf{success} v; \sigma' \quad v : t' \quad t' \not<: t}{x = e; \sigma \xRightarrow{\text{expr}} \mathbf{error}; \sigma'} \\ \text{E-ASGN-EXC} \frac{e; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'}{x = e; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'} \end{array}$$

The **if**-expression works like other languages: the **then**-branch is evaluated if the condition is true (E-IF-TRUE), otherwise the **else**-branch is evaluated (E-IF-FALSE). If the conditional produces a non-Boolean value then an error is raised (E-IF-ERR) and otherwise exceptional results are propagated (E-IF-EXC).

$$\text{E-IF-TRUE} \frac{e_{\text{cond}}; \sigma \xRightarrow{\text{expr}} \mathbf{success} \text{true}(); \sigma'' \quad e_1; \sigma'' \xRightarrow{\text{expr}} \text{vres}_1; \sigma'}{\mathbf{if} e_{\text{cond}} \mathbf{then} e_1 \mathbf{else} e_2; \sigma \xRightarrow{\text{expr}} \text{vres}_1; \sigma'}$$

$$\begin{array}{c}
\text{E-IF-FALSE} \frac{e_{\text{cond}}; \sigma \xRightarrow{\text{expr}} \mathbf{success} \text{ false}(); \sigma'' \quad e_2; \sigma'' \xRightarrow{\text{expr}} v_{\text{res}_2}; \sigma'}{\mathbf{if} \ e_{\text{cond}} \ \mathbf{then} \ e_1 \ \mathbf{else} \ e_2; \sigma \xRightarrow{\text{expr}} v_{\text{res}_2}; \sigma'} \\
\text{E-IF-ERR} \frac{e_{\text{cond}}; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ v; \sigma' \quad v \neq \text{true}() \quad v \neq \text{false}()}{\mathbf{if} \ e_{\text{cond}} \ \mathbf{then} \ e_1 \ \mathbf{else} \ e_2; \sigma \xRightarrow{\text{expr}} \mathbf{error}; \sigma'} \\
\text{E-IF-EXC} \frac{e_{\text{cond}}; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'}{\mathbf{if} \ e_{\text{cond}} \ \mathbf{then} \ e_1 \ \mathbf{else} \ e_2; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'}
\end{array}$$

The **switch**-expression initially evaluates the scrutinee expression (e), and then proceeds to execute the cases (discussed on page 116) on the result value (E-SWITCH-SUCS). The evaluation of cases is allowed to **fail**, in which case the evaluation is successful and has the special value \blacksquare (E-SWITCH-FAIL); other exceptional results are propagated as usual (E-SWITCH-EXC1, E-SWITCH-EXC2).

$$\begin{array}{c}
\text{E-SWITCH-SUCS} \frac{e; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ v; \sigma'' \quad \underline{cs}; v; \sigma'' \xRightarrow{\text{cases}} \mathbf{success} \ v'; \sigma'}{\mathbf{switch} \ e \ \underline{cs}; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ v'; \sigma'} \\
\text{E-SWITCH-FAIL} \frac{e; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ v; \sigma'' \quad \underline{cs}; v; \sigma'' \xRightarrow{\text{cases}} \mathbf{fail}; \sigma'}{\mathbf{switch} \ e \ \underline{cs}; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ \blacksquare; \sigma'} \\
\text{E-SWITCH-EXC1} \frac{e; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'}{\mathbf{switch} \ e \ \underline{cs}; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'} \\
\text{E-SWITCH-EXC2} \frac{e; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ v; \sigma'' \quad \underline{cs}; v; \sigma'' \xRightarrow{\text{cases}} \text{exres}; \sigma' \quad \text{exres} \neq \mathbf{fail}}{\mathbf{switch} \ e \ \underline{cs}; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'}
\end{array}$$

The **visit**-expression has similar evaluation cases to **switch** (E-VISIT-SUCS, E-VISIT-FAIL, E-VISIT-EXC1, E-VISIT-EXC2), except that the cases are evaluated using the visit relation that traverses the produced value of target expression given the provided strategy.

$$\begin{array}{c}
\text{E-VISIT-SUCS} \frac{e; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ v; \sigma'' \quad \underline{cs}; v; \sigma'' \xRightarrow{\text{visit}} \mathbf{success} \ v'; \sigma'}{st \ \mathbf{visit} \ e \ \underline{cs}; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ v'; \sigma'} \\
\text{E-VISIT-FAIL} \frac{e; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ v; \sigma'' \quad \underline{cs}; v; \sigma'' \xRightarrow{\text{visit}} \mathbf{fail}; \sigma'}{st \ \mathbf{visit} \ e \ \underline{cs}; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ v; \sigma'} \\
\text{E-VISIT-EXC1} \frac{e; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'}{st \ \mathbf{visit} \ e \ \underline{cs}; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'} \\
\text{E-VISIT-EXC2} \frac{e; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ v; \sigma'' \quad \underline{cs}; v; \sigma'' \xRightarrow{\text{visit}} \text{exres}; \sigma' \quad \text{exres} \neq \mathbf{fail}}{st \ \mathbf{visit} \ e \ \underline{cs}; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'}
\end{array}$$

The control operations **break**, **continue** and **fail** evaluate to themselves without any side-effects (E-BREAK, E-CONTINUE, E-FAIL).

$$\begin{array}{c}
\text{E-BREAK} \frac{}{\mathbf{break}; \sigma \xRightarrow{\text{expr}} \mathbf{break}; \sigma} \quad \text{E-FAIL} \frac{}{\mathbf{fail}; \sigma \xRightarrow{\text{expr}} \mathbf{fail}; \sigma} \\
\text{E-CONTINUE} \frac{}{\mathbf{continue}; \sigma \xRightarrow{\text{expr}} \mathbf{continue}; \sigma}
\end{array}$$

Blocks allow evaluating inner expressions using a local declaration of variables, which are then afterwards removed from the resulting store (E-BLOCK-SUCS, E-BLOCK-EXC). Recall, that we consider an implicit definition environment based on scoping, and so the local declarations in the block will be implicitly available in the evaluation of the body subexpression sequence.

$$\begin{array}{c}
\text{E-BLOCK-SUCS} \frac{e; \sigma \xRightarrow{\text{expr}^*} \mathbf{success} \ \underline{v}; \sigma'}{\mathbf{local} \ \underline{t} \ x \ \mathbf{in} \ \underline{e} \ \mathbf{end}; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ \text{last}(\underline{v}); (\sigma' \setminus \underline{x})} \\
\text{E-BLOCK-EXC} \frac{e; \sigma \xRightarrow{\text{expr}^*} \text{exres}; \sigma'}{\mathbf{local} \ \underline{t} \ x \ \mathbf{in} \ \underline{e} \ \mathbf{end}; \sigma \xRightarrow{\text{expr}} \text{exres}; (\sigma' \setminus \underline{x})}
\end{array}$$

The auxiliary function `last` is here used to extract the last element in the sequence (or return \blacksquare if empty).

$$\begin{aligned} \text{last}(v_1, \dots, v_n, v') &= v' \\ \text{last}(\varepsilon) &= \blacksquare \end{aligned}$$

The **for**-loop evaluates target generator expression to a set of possible *environments* that represent possible assignments of variables to values—propagating exceptions if necessary (E-FOR-EXC)—and then it iterates over each possible assignment using the each-relation (E-FOR-SUCS).

$$\begin{array}{c} \text{E-FOR-SUCS} \frac{\begin{array}{c} g; \sigma \xRightarrow{\text{gexpr}} \text{success } \underline{\rho}; \sigma'' \\ e; \underline{\rho}; \sigma'' \xRightarrow{\text{each}} vres; \sigma' \end{array}}{\text{for } g \ e; \sigma \xRightarrow{\text{expr}} vres; \sigma'} \quad \text{E-FOR-EXC} \frac{g; \sigma \xRightarrow{\text{gexpr}} \text{exres}; \sigma'}{\text{for } g \ e; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'} \end{array}$$

The evaluation of **while**-loops is analogous to other imperative languages with control operations, in that the body of the while loop is continuously executed until the target condition does not hold (E-WHILE-FALSE). If the body successfully finishes with an value or **continue** then iteration continues (E-WHILE-TRUE-SUCS), if the body finishes with **break** the iteration stops with value \blacksquare (E-WHILE-TRUE-BREAK), if the conditional evaluates to a non-Boolean value it errors out (E-WHILE-ERR) and otherwise if another kind of exceptional result is produced then it is propagated (E-WHILE-EXC1, E-WHILE-EXC2).

$$\begin{array}{c} \text{E-WHILE-FALSE} \frac{e_{\text{cond}}; \sigma \xRightarrow{\text{expr}} \text{success } \text{false}(); \sigma'}{\text{while } e_{\text{cond}} \ e; \sigma \xRightarrow{\text{expr}} \text{success } \blacksquare; \sigma'} \\ \\ \text{E-WHILE-TRUE-SUCS} \frac{\begin{array}{c} e_{\text{cond}}; \sigma \xRightarrow{\text{expr}} \text{success } \text{true}(); \sigma'' \quad e; \sigma'' \xRightarrow{\text{expr}} vres; \sigma''' \\ vres = \text{success } v \vee vres = \text{continue} \\ \text{while } e_{\text{cond}} \ e; \sigma''' \xRightarrow{\text{expr}} vres'; \sigma' \end{array}}{\text{while } e_{\text{cond}} \ e; \sigma \xRightarrow{\text{expr}} vres'; \sigma'} \\ \\ \text{E-WHILE-TRUE-BREAK} \frac{e_{\text{cond}}; \sigma \xRightarrow{\text{expr}} \text{success } \text{true}(); \sigma'' \quad e; \sigma'' \xRightarrow{\text{expr}} \text{break}; \sigma'}{\text{while } e_{\text{cond}} \ e; \sigma \xRightarrow{\text{expr}} \text{success } \blacksquare; \sigma'} \end{array}$$

$$\begin{array}{c}
\text{E-WHILE-EXC1} \frac{e_{\text{cond}}; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'}{\mathbf{while} \ e_{\text{cond}} \ e; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'} \\
\\
\text{E-WHILE-EXC2} \frac{e_{\text{cond}}; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ \text{true}(); \sigma'' \quad e; \sigma'' \xRightarrow{\text{expr}} \text{exres}; \sigma' \quad \text{exres} \in \{\mathbf{throw} \ v, \mathbf{return} \ v, \mathbf{fail}, \mathbf{error}\}}{\mathbf{while} \ e_{\text{cond}} \ e; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'} \\
\\
\text{E-WHILE-ERR} \frac{e_{\text{cond}}; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ v; \sigma' \quad v \neq \text{true}() \quad v \neq \text{false}()}{\mathbf{while} \ e_{\text{cond}} \ e; \sigma \xRightarrow{\text{expr}} \mathbf{error}; \sigma'}
\end{array}$$

The **solve**-loop keeps evaluating the body expression until the provided variables reach a fixed-point; this is analogous to the TRON **fix**-statement presented in Chapter 5. Initially, the body expression is evaluated and then the values of target variables is compared from before and after iteration; if the values are equal after an iteration, then evaluation stops (E-SOLVE-EQ) and otherwise the iteration continues (E-SOLVE-NEQ). If any of the variables do not have a value assigned, an error is produced (E-SOLVE-ERR), and otherwise if an exceptional result is produced, it is propagated (E-SOLVE-EXC).

$$\begin{array}{c}
\text{E-SOLVE-EQ} \frac{e; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ v; \sigma' \quad \underline{x} \subseteq \text{dom} \ \sigma \cap \text{dom} \ \sigma' \quad \sigma(x) = \sigma'(x)}{\mathbf{solve} \ \underline{x} \ e; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ v; \sigma'} \\
\\
\text{E-SOLVE-NEQ} \frac{e; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ v; \sigma'' \quad \underline{x} \subseteq \text{dom} \ \sigma \cap \text{dom} \ \sigma'' \quad \exists i. \sigma(x_i) \neq \sigma''(x_i) \quad \mathbf{solve} \ \underline{x} \ e; \sigma'' \xRightarrow{\text{expr}} \text{vres}; \sigma'}{\mathbf{solve} \ \underline{x} \ e; \sigma \xRightarrow{\text{expr}} \text{vres}; \sigma'} \\
\\
\text{E-SOLVE-EXC} \frac{e; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'}{\mathbf{solve} \ \underline{x} \ e; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'} \\
\\
\text{E-SOLVE-ERR} \frac{e; \sigma \xRightarrow{\text{expr}} \mathbf{success} \ v; \sigma' \quad x_i \notin \text{dom} \ \sigma \cap \text{dom} \ \sigma'}{\mathbf{solve} \ \underline{x} \ e; \sigma \xRightarrow{\text{expr}} \mathbf{error}; \sigma'}
\end{array}$$

The **throw**-expression, evaluates its inner expression first—propagating exceptional results if necessary (E-THR-EXC)—and then produces a **throw** result with result value (E-THR-SUCS).

$$\text{E-THR-SUCS} \frac{e; \sigma \xRightarrow{\text{expr}} \mathbf{success} v; \sigma'}{\mathbf{throw} e; \sigma \xRightarrow{\text{expr}} \mathbf{throw} v; \sigma'} \quad \text{E-THR-EXC} \frac{e; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'}{\mathbf{throw} e; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'}$$

The **try-finally** expression executes the **try**-body first and then the **finally**-body. If the **finally**-body produces an exceptional result during execution then that result is propagated (E-FIN-EXC) and otherwise the **try**-body result value is used (E-FIN-SUCS).

$$\text{E-FIN-SUCS} \frac{e_1; \sigma \xRightarrow{\text{expr}} \text{vres}_1; \sigma'' \quad e_2; \sigma'' \xRightarrow{\text{expr}} \mathbf{success} v_2; \sigma'}{\mathbf{try} e_1 \mathbf{finally} e_2; \sigma \xRightarrow{\text{expr}} \text{vres}_1; \sigma'}$$

$$\text{E-FIN-EXC} \frac{e_1; \sigma \xRightarrow{\text{expr}} \text{vres}_1; \sigma'' \quad e_2; \sigma'' \xRightarrow{\text{expr}} \text{exres}_2; \sigma'}{\mathbf{try} e_1 \mathbf{finally} e_2; \sigma \xRightarrow{\text{expr}} \text{exres}_2; \sigma'}$$

The **try-catch** expression evaluates the **try**-body and if it produces a thrown value, then it binds the value in the body of **catch** and continues evaluation (E-TRY-CATCH). For all other results, it simply propagates them without evaluating the **catch**-body (E-TRY-ORD).

$$\text{E-TRY-CATCH} \frac{e_1; \sigma \xRightarrow{\text{expr}} \mathbf{throw} v_1; \sigma'' \quad e_2; \sigma''[x \mapsto v_1] \xRightarrow{\text{expr}} \text{vres}_2; \sigma'}{\mathbf{try} e_1 \mathbf{catch} x \Rightarrow e_2; \sigma \xRightarrow{\text{expr}} \text{vres}_2; (\sigma' \setminus x)}$$

$$\text{E-TRY-ORD} \frac{e_1; \sigma \xRightarrow{\text{expr}} \text{vres}_1; \sigma' \quad \text{vres}_1 \neq \mathbf{throw} v_1}{\mathbf{try} e_1 \mathbf{catch} x \Rightarrow e_2; \sigma \xRightarrow{\text{expr}} \text{vres}_1; \sigma'}$$

Expression Sequences Evaluating a sequence of expressions proceeds by evaluating each expression, combining the results if successful (ES-EMP, ES-MORE) and otherwise propagating the first exceptional result encountered (ES-Exc1, ES-Exc2).

$$\text{ES-MORE} \frac{e; \sigma \xRightarrow{\text{expr}} \mathbf{success} v; \sigma'' \quad \underline{e'}; \sigma'' \xRightarrow{\text{expr}^*} \mathbf{success} \underline{v'}; \sigma'}{e, \underline{e'}; \sigma \xRightarrow{\text{expr}^*} \mathbf{success} v, \underline{v'}; \sigma'}$$

$$\begin{array}{c}
\text{ES-EMP} \frac{\text{success } \varepsilon; \sigma \Longrightarrow_{\text{expr}^*} \varepsilon; \sigma}{\text{success } \varepsilon; \sigma \Longrightarrow_{\text{expr}^*} \varepsilon; \sigma} \quad \text{ES-Exc1} \frac{e; \sigma \Longrightarrow_{\text{expr}} \text{exres}; \sigma'}{e, \underline{\quad}; \sigma \Longrightarrow_{\text{expr}^*} \text{exres}; \sigma'} \\
\text{ES-Exc2} \frac{e; \sigma \Longrightarrow_{\text{expr}} \text{success } v; \sigma'' \quad \underline{\quad}; \sigma'' \Longrightarrow_{\text{expr}^*} \text{exres}; \sigma'}{e, \underline{\quad}; \sigma \Longrightarrow_{\text{expr}^*} \text{exres}; \sigma'}
\end{array}$$

Cases The evaluation relation for evaluating a series of cases has the form $\boxed{\underline{\quad}; v; \sigma \Longrightarrow_{\text{cases}} \text{vres}; \sigma'}$, and intuitively proceeds by sequentially evaluating each case (in cs) against value v until one of them produces a non-**fail** result. For each $\overline{\text{case}}$, the first step is to match the given value against target pattern and then evaluate the target expression under the set of possible matches; if the evaluation of the target expression produces a **fail** as result, the rest of the cases are evaluated in a restored initial state (ECS-MORE-FAIL) and otherwise the result is propagated (ECS-MORE-ORD). If all possible cases are exhausted, the result is **fail** (ECS-EMP).

$$\begin{array}{c}
\text{ECS-EMP} \frac{}{\varepsilon; v; \sigma \Longrightarrow_{\text{cases}} \text{fail}; \sigma} \\
\text{ECS-MORE-FAIL} \frac{\sigma \vdash p \stackrel{?}{:=} v \Longrightarrow_{\text{match}} \underline{\rho} \quad \underline{\rho}; e; \sigma \Longrightarrow_{\text{case}} \text{fail}; \sigma'' \quad \underline{\quad}; v; \sigma \Longrightarrow_{\text{cases}} \text{vres}; \sigma'}{\text{case } p \Rightarrow e, \underline{\quad}; v; \sigma \Longrightarrow_{\text{cases}} \text{vres}; \sigma'} \\
\text{ECS-MORE-ORD} \frac{\sigma \vdash p \stackrel{?}{:=} v \Longrightarrow_{\text{match}} \underline{\rho} \quad \underline{\rho}; e; \sigma \Longrightarrow_{\text{case}} \text{vres}; \sigma' \quad \text{vres} \neq \text{fail}}{\text{case } p \Rightarrow e, \underline{\quad}; v; \sigma \Longrightarrow_{\text{cases}} \text{vres}; \sigma'}
\end{array}$$

Evaluating a single case—with relation $\boxed{\underline{\rho}; e; \sigma \Longrightarrow_{\text{case}} \text{vres}; \sigma'}$ —requires trying each possible binding (in $\underline{\rho}$) sequentially, producing **fail** if no binding is available (EC-EMP). If evaluating target expression produces a non-**fail** value then it is propagated (EC-MORE-ORD), otherwise the rest of the possible bindings are tried in a restored initial state (EC-MORE-FAIL).

$$\begin{array}{c}
\text{EC-MORE-FAIL} \frac{e; \sigma \xrightarrow[\text{expr}]{\text{fail}} \sigma'' \quad \underline{\rho'}; e; \sigma \xrightarrow[\text{case}]{vres} \sigma'}{\rho, \underline{\rho'}; e; \sigma \xrightarrow[\text{case}]{vres} \sigma'} \\
\\
\text{EC-EMP} \frac{\varepsilon; e; \sigma \xrightarrow[\text{case}]{\text{fail}} \sigma}{\varepsilon; e; \sigma \xrightarrow[\text{case}]{\text{fail}} \sigma} \quad \text{EC-MORE-ORD} \frac{e; \sigma \xrightarrow[\text{expr}]{vres} \sigma' \quad vres \neq \text{fail}}{\rho, \underline{\rho'}; e; \sigma \xrightarrow[\text{case}]{vres} (\sigma' \setminus \text{dom } \rho)}
\end{array}$$

Traversals One of the key features of Rascal is **visit**-expressions which provide generic traversals over data values and collections, allowing for multiple strategies to determine the traversal ordering and halting conditions. In a traditional object-oriented language or functional language, transforming large structures is cumbersome and requires a great amount of boilerplate, requiring a function for each type of datatype, where each function must deconstruct the input data, applying target changes, recursively calling the right traversal functions for traversal of further contained data and reconstructing the data with new values. Precisely, the first-class handling of these aspects combined with the other features such as the powerful pattern matching makes Rascal particularly suitable as a high-level transformation language.

The main traversal relation $\underline{cs}; v; \sigma \xrightarrow[\text{visit}]{st} vres; \sigma'$ delegates execution to the correct strategy-dependent traversal, and performs fixed-point calculation if necessary. For the **top-down-break** and **top-down** strategies it uses the top-down traversal relation $\underline{cs}; v; \sigma \xrightarrow[\text{td-visit}]{br} vres; \sigma'$ specifying **break** and no **no-break** as breaking strategies respectively (EV-TD and EV-TDB); this works analogously with the **bottom-up-break** (EV-BU and EV-BUB) and **bottom-up** strategies using the $\underline{cs}; v; \sigma \xrightarrow[\text{bu-visit}]{br} vres; \sigma'$ relation.

$$\begin{array}{cc}
\text{EV-TD} \frac{\underline{cs}; v; \sigma \xrightarrow[\text{td-visit}]{\text{no-break}} vres; \sigma'}{\underline{cs}; v; \sigma \xrightarrow[\text{visit}]{\text{top-down}} vres; \sigma'} & \text{EV-TDB} \frac{\underline{cs}; v; \sigma \xrightarrow[\text{td-visit}]{\text{break}} vres; \sigma'}{\underline{cs}; v; \sigma \xrightarrow[\text{visit}]{\text{top-down-break}} vres; \sigma'} \\
\\
\text{EV-BU} \frac{\underline{cs}; v; \sigma \xrightarrow[\text{bu-visit}]{\text{no-break}} vres; \sigma'}{\underline{cs}; v; \sigma \xrightarrow[\text{visit}]{\text{bottom-up}} vres; \sigma'} & \text{EV-BUB} \frac{\underline{cs}; v; \sigma \xrightarrow[\text{bu-visit}]{\text{break}} vres; \sigma'}{\underline{cs}; v; \sigma \xrightarrow[\text{visit}]{\text{bottom-up-break}} vres; \sigma'}
\end{array}$$

The **innermost** strategy evaluates the **bottom-up** traversal as long as it produces a resulting value not equal to the one from the previous iteration (EV-IM-NEQ), returning the result value when a fixed-point is reached (EV-IM-EQ); if any exceptional result happens during evaluation it will be propagated (EV-IM-Exc). Analogous evaluation steps happens with **outermost** strategy and **top-down** traversal (EV-OM-NEQ, EV-OM-EQ, EV-OM-Exc).

$$\begin{array}{c}
 \text{EV-IM-EQ} \frac{\underline{cs}; v; \sigma \xrightarrow{\text{no-break}}_{\text{bu-visit}} \text{success } v; \sigma'}{\underline{cs}; v; \sigma \xrightarrow{\text{innermost}}_{\text{visit}} \text{success } v; \sigma'} \quad \text{EV-IM-Exc} \frac{\underline{cs}; v; \sigma \xrightarrow{\text{no-break}}_{\text{bu-visit}} \text{exres}; \sigma'}{\underline{cs}; v; \sigma \xrightarrow{\text{innermost}}_{\text{visit}} \text{exres}; \sigma'} \\
 \\
 \text{EV-IM-NEQ} \frac{\underline{cs}; v; \sigma \xrightarrow{\text{no-break}}_{\text{bu-visit}} \text{success } v'; \sigma'' \quad v \neq v' \quad \underline{cs}; v'; \sigma'' \xrightarrow{\text{innermost}}_{\text{visit}} \text{vres}; \sigma'}{\underline{cs}; v; \sigma \xrightarrow{\text{innermost}}_{\text{visit}} \text{vres}; \sigma'} \\
 \\
 \text{EV-OM-EQ} \frac{\underline{cs}; v; \sigma \xrightarrow{\text{no-break}}_{\text{td-visit}} \text{success } v; \sigma'}{\underline{cs}; v; \sigma \xrightarrow{\text{outermost}}_{\text{visit}} \text{success } v; \sigma'} \quad \text{EV-OM-Exc} \frac{\underline{cs}; v; \sigma \xrightarrow{\text{no-break}}_{\text{td-visit}} \text{exres}; \sigma'}{\underline{cs}; v; \sigma \xrightarrow{\text{outermost}}_{\text{visit}} \text{exres}; \sigma'} \\
 \\
 \text{EV-OM-NEQ} \frac{\underline{cs}; v; \sigma \xrightarrow{\text{no-break}}_{\text{td-visit}} \text{success } v'; \sigma'' \quad v \neq v' \quad \underline{cs}; v'; \sigma'' \xrightarrow{\text{outermost}}_{\text{visit}} \text{vres}; \sigma'}{\underline{cs}; v; \sigma \xrightarrow{\text{outermost}}_{\text{visit}} \text{vres}; \sigma'}
 \end{array}$$

The top-down traversal strategy starts by executing all cases on the target value, *scrutinee*, applying possible replacements and effects to produce an intermediate result value; the traversal then continues on the sequence of contained values of the this intermediate result, finally reconstructing a new output containing the possible replacement values obtained (ETV-ORD-SUCS1, ETV-ORD-SUCS2). If using the **break** strategy, the traversal will stop at the first value that produces a successful result (ETV-BREAK-SUCS); otherwise, if any sub-result produces a non-**fail** exceptional result it is propagated (ETV-Exc1, ETV-Exc2).

$$\text{ETV-BREAK-SUCS} \frac{\underline{cs}; v; \sigma \xrightarrow{\text{cases}} \text{success } v'; \sigma' \quad br = \text{break}}{\underline{cs}; v; \sigma \xrightarrow{\text{td-visit}}_{br} \text{success } v'; \sigma'}$$

$$\begin{array}{c}
\begin{array}{c}
\frac{cs; v; \sigma \xRightarrow{\text{cases}} vfres; \sigma'' \quad br = \mathbf{break} \Rightarrow vfres = \mathbf{fail}}{v'' = \mathbf{if-fail}(vfres, v) \quad v''' = \mathbf{children}(v'')} \\
\frac{cs; v'''; \sigma'' \xrightarrow[\text{td-visit}]{br} \mathbf{fail}; \sigma'}{} \\
\text{ETV-ORD-Sucs1} \text{-----} \\
\frac{}{cs; v; \sigma \xrightarrow[\text{td-visit}]{br} vfres; \sigma'}
\end{array} \\
\begin{array}{c}
\frac{cs; v; \sigma \xRightarrow{\text{cases}} vfres; \sigma'' \quad br = \mathbf{break} \Rightarrow vfres = \mathbf{fail}}{v'' = \mathbf{if-fail}(vfres, v) \quad v''' = \mathbf{children}(v'')} \\
\frac{cs; v'''; \sigma'' \xrightarrow[\text{td-visit}]{br} \mathbf{success} \ v'''; \sigma' \quad \mathbf{recons} \ v'' \ \mathbf{using} \ v'''' \ \mathbf{to} \ rcres}{} \\
\text{ETV-ORD-Sucs2} \text{-----} \\
\frac{}{cs; v; \sigma \xrightarrow[\text{td-visit}]{br} rcres; \sigma'} \\
\frac{cs; v; \sigma \xRightarrow{\text{cases}} exres; \sigma' \quad exres \neq \mathbf{fail}}{\text{ETV-Exc1} \text{-----}} \\
\frac{}{cs; v; \sigma \xrightarrow[\text{td-visit}]{br} exres; \sigma'} \\
\frac{cs; v; \sigma \xRightarrow{\text{cases}} vfres; \sigma'' \quad br = \mathbf{break} \Rightarrow vfres = \mathbf{fail}}{v'' = \mathbf{if-fail}(vfres, v) \quad v''' = \mathbf{children}(v'')} \\
\frac{cs; v'''; \sigma'' \xrightarrow[\text{td-visit}]{br} exres; \sigma' \quad exres \neq \mathbf{fail}}{\text{ETV-Exc2} \text{-----}} \\
\frac{}{cs; v; \sigma \xrightarrow[\text{td-visit}]{br} exres; \sigma'}
\end{array}
\end{array}$$

Evaluating a sequence of top-down traversals, requires executing a top-down traversal for each element, failing if the input sequence is empty (ETVS-EMP) and otherwise combining the results (ETVS-MORE). If the **break** strategy is used, then the iteration will instead stop at first succesful result (ETVS-BREAK), and any non-**fail** exceptional result is propagated (ETVS-Exc1, ETVS-Exc2).

$$\begin{array}{c}
\text{ETVS-EMP} \text{-----} \\
\frac{}{cs; \varepsilon; \sigma \xrightarrow[\text{td-visit}]{br} \mathbf{fail}; \sigma} \\
\text{ETVS-BREAK} \text{-----} \\
\frac{cs; v; \sigma \xrightarrow[\text{td-visit}]{br} \mathbf{success} \ v''; \sigma' \quad br = \mathbf{break}}{\frac{cs; v, v'; \sigma \xrightarrow[\text{td-visit}]{br} \mathbf{success} \ v'', v'; \sigma'}} \\
\text{ETVS-MORE} \text{-----} \\
\frac{cs; v; \sigma \xrightarrow[\text{td-visit}]{br} vfres; \sigma'' \quad br = \mathbf{break} \Rightarrow vfres = \mathbf{fail}}{\frac{cs; v'; \sigma'' \xrightarrow[\text{td-visit}]{br} vfres^*; \sigma'}} \\
\frac{cs; v, v'; \sigma \xrightarrow[\text{td-visit}]{br} \mathbf{vcombine}(vfres, vfres^*, v, v'); \sigma'}{}
\end{array}$$

$$\begin{array}{c}
\text{ETVS-Exc1} \frac{\underline{cs}; v; \sigma \xrightarrow[\text{td-visit}]{br} \text{exres}; \sigma' \quad \text{exres} \neq \mathbf{fail}}{\underline{cs}; v, \underline{v}'; \sigma \xrightarrow[\text{td-visit}]{br} \text{exres}; \sigma'} \\
\\
\text{ETVS-Exc2} \frac{\underline{cs}; v; \sigma \xrightarrow[\text{td-visit}]{br} \text{vfres}; \sigma'' \quad br = \mathbf{break} \Rightarrow \text{vfres} = \mathbf{fail} \quad \underline{cs}; \underline{v}'; \sigma'' \xrightarrow[\text{td-visit}]{br} \text{exres}; \sigma' \quad \text{exres} \neq \mathbf{fail}}{\underline{cs}; v, \underline{v}'; \sigma \xrightarrow[\text{td-visit}]{br} \text{exres}; \sigma'}
\end{array}$$

Bottom-up traversals work analogously to top-down traversals, except that traversal of children and reconstruction happens before traversing the final reconstructed value (EBU-BREAK-SUCS, EBU-NO-BREAK-SUCS, EBU-FAIL-SUCS). The analogy also holds with propagation of exceptional results and errors (EBU-Exc, EBU-NO-BREAK-ERR, EBU-NO-BREAK-Exc).

$$\begin{array}{c}
\text{EBU-BREAK-SUCS} \frac{\underline{v}'' = \text{children}(v) \quad \underline{cs}; \underline{v}''; \sigma \xrightarrow[\text{bu-visit}]{br} \mathbf{success} \underline{v}'; \sigma' \quad br = \mathbf{break} \quad \mathbf{recons} \ v \ \mathbf{using} \ \underline{v}' \ \mathbf{to} \ r\text{cres}}{\underline{cs}; v; \sigma \xrightarrow[\text{bu-visit}]{br} r\text{cres}; \sigma'} \\
\\
\text{EBU-NO-BREAK-SUCS} \frac{\underline{v}'' = \text{children}(v) \quad \underline{cs}; \underline{v}''; \sigma \xrightarrow[\text{bu-visit}]{br} \mathbf{success} \underline{v}'''; \sigma'' \quad br = \mathbf{no-break} \quad \mathbf{recons} \ v \ \mathbf{using} \ \underline{v}''' \ \mathbf{to} \ \mathbf{success} \ \underline{v}' \quad \underline{cs}; \underline{v}'; \sigma'' \xrightarrow[\text{cases}]{\Rightarrow} \text{vfres}'; \sigma'}{\underline{cs}; v; \sigma \xrightarrow[\text{bu-visit}]{br} \mathbf{success} \ \text{if-fail}(\text{vfres}', \underline{v}'); \sigma'} \\
\\
\text{EBU-FAIL-SUCS} \frac{\underline{v}'' = \text{children}(v) \quad \underline{cs}; \underline{v}''; \sigma \xrightarrow[\text{bu-visit}]{br} \mathbf{fail}; \sigma'' \quad \underline{cs}; v; \sigma'' \xrightarrow[\text{cases}]{\Rightarrow} \text{vres}; \sigma'}{\underline{cs}; v; \sigma \xrightarrow[\text{bu-visit}]{br} \text{vres}; \sigma'} \\
\\
\text{EBU-Exc} \frac{\underline{v}' = \text{children}(v) \quad \underline{cs}; \underline{v}'; \sigma \xrightarrow[\text{bu-visit}]{br} \text{exres}; \sigma' \quad \text{exres} \neq \mathbf{fail}}{\underline{cs}; v; \sigma \xrightarrow[\text{bu-visit}]{br} \text{exres}; \sigma'}
\end{array}$$

$$\text{EBU-NO-BREAK-ERR} \frac{\begin{array}{l} \underline{v''} = \text{children}(v) \quad \underline{cs}; \underline{v''}; \sigma \xrightarrow[\text{bu-visit}\star]{br} \text{success } \underline{v'''}; \sigma' \\ br = \text{no-break} \quad \text{recons } v \text{ using } \underline{v'''} \text{ to error} \end{array}}{\underline{cs}; \underline{v}; \sigma \xrightarrow[\text{bu-visit}]{br} \text{error}; \sigma'}$$

$$\text{EBU-NO-BREAK-EXC} \frac{\begin{array}{l} \underline{v''} = \text{children}(v) \quad \underline{cs}; \underline{v''}; \sigma \xrightarrow[\text{bu-visit}\star]{br} \text{success } \underline{v'''}; \sigma'' \\ br = \text{no-break} \quad \text{recons } v \text{ using } \underline{v'''} \text{ to success } v' \\ \underline{cs}; \underline{v'}; \sigma'' \xrightarrow[\text{cases}]{=} \text{exres}; \sigma' \end{array}}{\underline{cs}; \underline{v}; \sigma \xrightarrow[\text{bu-visit}]{br} \text{exres}; \sigma'}$$

Evaluating a sequence of bottom-up traversals is analogous to evaluating a sequence of top-down traversals. Each element in the sequence is evaluated and their results is combined (EBUS-EMP,EBUS-MORE), stopping at the first succesful result when using the break strategy (EBUS-BREAK). Otherwise, non-**fail** exceptional results are propagated (EBUS-Exc1, EBUS-Exc2).

$$\text{EBUS-EMP} \frac{}{\underline{cs}; \underline{\varepsilon}; \sigma \xrightarrow[\text{bu-visit}\star]{br} \text{fail}; \sigma}$$

$$\text{EBUS-BREAK} \frac{\begin{array}{l} \underline{cs}; \underline{v}; \sigma \xrightarrow[\text{bu-visit}]{br} \text{success } \underline{v''}; \sigma \quad br = \text{break} \\ \underline{cs}; \underline{v}, \underline{v'}; \sigma \xrightarrow[\text{bu-visit}\star]{br} \text{success } \underline{v''}, \underline{v'}; \sigma' \end{array}}{\underline{cs}; \underline{v}; \sigma \xrightarrow[\text{bu-visit}]{br} \text{success } \underline{v''}; \sigma \quad br = \text{break} \Rightarrow \text{v}fres = \text{fail}}$$

$$\text{EBUS-MORE} \frac{\begin{array}{l} \underline{cs}; \underline{v'}; \sigma'' \xrightarrow[\text{bu-visit}\star]{br} \text{v}fres\star'; \sigma' \\ \underline{cs}; \underline{v}, \underline{v'}; \sigma \xrightarrow[\text{bu-visit}\star]{br} \text{vcombine}(\text{v}fres, \text{v}fres\star', \underline{v}, \underline{v'}); \sigma' \end{array}}{\underline{cs}; \underline{v}; \sigma \xrightarrow[\text{bu-visit}]{br} \text{v}fres; \sigma'' \quad br = \text{break} \Rightarrow \text{v}fres = \text{fail}}$$

$$\text{EBUS-Exc1} \frac{\begin{array}{l} \underline{cs}; \underline{v}; \sigma \xrightarrow[\text{bu-visit}]{br} \text{exres}; \sigma' \quad \text{exres} \neq \text{fail} \\ \underline{cs}; \underline{v}, \underline{v'}; \sigma \xrightarrow[\text{bu-visit}\star]{br} \text{exres}; \sigma' \end{array}}{\underline{cs}; \underline{v}; \sigma \xrightarrow[\text{bu-visit}]{br} \text{exres}; \sigma' \quad \text{exres} \neq \text{fail}}$$

$$\begin{array}{c}
 \underline{cs}; v; \sigma \xrightarrow[\text{bu-visit}]{br} \underline{vfres}; \sigma'' \quad br = \mathbf{break} \Rightarrow \underline{vfres} = \mathbf{fail} \\
 \underline{cs}; \underline{v}'; \sigma'' \xrightarrow[\text{bu-visit}^*]{br} \underline{exres}; \sigma' \\
 \hline
 \underline{cs}; v, \underline{v}'; \sigma \xrightarrow[\text{bu-visit}^*]{br} \underline{exres}; \sigma'
 \end{array}$$

EBUS-Exc2

Auxiliary The children function extracts the directly contained values of the given input value.

$$\begin{array}{l}
 \text{children}(vb) = \varepsilon \quad \text{children}(k(\underline{v})) = \underline{v} \\
 \text{children}([\underline{v}]) = \underline{v} \quad \text{children}(\{\underline{v}\}) = \underline{v} \\
 \text{children}((\underline{v} : \underline{v}')) = \underline{v}, \underline{v}' \quad \text{children}(\blacksquare) = \varepsilon
 \end{array}$$

$$\underline{vfres} ::= \mathbf{success} \ v \mid \mathbf{fail}$$

The if-fail function will return a provided default value if the first argument is **fail**, and otherwise it will use the provided value in the first argument.

$$\text{if-fail}(\mathbf{fail}, v) = v \quad \text{if-fail}(\mathbf{success} \ v', v) = v'$$

The vcombine function will combine **success** and **fail** results from visitor, producing **fail** if both result arguments are **fail** otherwise producing **success** result, possibly using default values

$$\text{vcombine}(\underline{vfres}, \underline{vfres}^*, v, \underline{v}') = \begin{cases} \mathbf{fail} & \text{if } \underline{vfres} = \mathbf{fail} \wedge \\ & \underline{vfres}^* = \mathbf{fail} \\ \mathbf{success} \left(\begin{array}{l} \text{if-fail}(\underline{vfres}, v), \\ \text{if-fail}(\underline{vfres}^*, \underline{v}') \end{array} \right) & \text{otherwise} \end{cases}$$

The reconstruction relation $\boxed{\text{recons } v \text{ using } \underline{v}' \text{ to } r\text{cres}}$ tries to update the elements of a value, checking whether provided values are type correct and defined (not \blacksquare), otherwise producing an error.

$$r\text{cres} ::= \mathbf{success} \ v \mid \mathbf{error}$$

$$\begin{array}{c}
\text{RC-VAL-SUCS} \frac{}{\text{recons } vb \text{ using } \varepsilon \text{ to success } vb} \\
\text{RC-VAL-ERR} \frac{}{\text{recons } vb \text{ using } v', \underline{v''} \text{ to error}} \\
\text{RC-CONS-SUCS} \frac{\text{data } at = \dots | k(\underline{t}) | \dots \quad \underline{v'} : \underline{t'} \quad \underline{v'} \neq \blacksquare \quad \underline{t'} <: \underline{t}}{\text{recons } k(\underline{v}) \text{ using } \underline{v'} \text{ to success } k(\underline{v'})} \\
\text{RC-CONS-ERR} \frac{\text{data } at = \dots | k(\underline{t}) | \dots \quad \underline{v'} : \underline{t'} \quad v_i = \blacksquare \vee t'_i \not<: t_i}{\text{recons } k(\underline{v}) \text{ using } \underline{v'} \text{ to error}} \\
\text{RC-LIST-SUCS} \frac{\underline{v'} \neq \blacksquare}{\text{recons } [\underline{v}] \text{ using } \underline{v'} \text{ to success } [\underline{v'}]} \\
\text{RC-LIST-ERR} \frac{\underline{v'_i} = \blacksquare}{\text{recons } [\underline{v}] \text{ using } \underline{v'} \text{ to error}} \\
\text{RC-SET-SUCS} \frac{\underline{v'} = \blacksquare}{\text{recons } \{\underline{v}\} \text{ using } \underline{v'} \text{ to success } \{\underline{v'}\}} \\
\text{RC-SET-ERR} \frac{\underline{v'_i} = \blacksquare}{\text{recons } \{\underline{v}\} \text{ using } \underline{v'} \text{ to error}} \\
\text{RC-MAP-SUCS} \frac{\underline{v''} \neq \blacksquare \quad \underline{v'''} \neq \blacksquare}{\text{recons } (\underline{v} : \underline{v'}) \text{ using } \underline{v''}, \underline{v'''} \text{ to success } (\underline{v''} : \underline{v'''})} \\
\text{RC-MAP-ERR} \frac{\underline{v''_i} = \blacksquare \vee \underline{v'''_i} = \blacksquare}{\text{recons } (\underline{v} : \underline{v'}) \text{ using } \underline{v''}, \underline{v'''} \text{ to error}} \\
\text{RC-BOT-SUCS} \frac{}{\text{recons } \blacksquare \text{ using } \varepsilon \text{ to success } \blacksquare} \\
\text{RC-BOT-ERR} \frac{}{\text{recons } \blacksquare \text{ using } v', \underline{v''} \text{ to error}}
\end{array}$$

Enumeration The enumeration relation $\boxed{e; \underline{\rho}; \sigma \xRightarrow{\text{each}} vres; \sigma'}$ iterates over all provided bindings (EE-MORE-SUCS) until there are none left (EE-EMP) or the result is neither an ordinary value or **continue** from one of the iterations (EE-MORE-EXC); in case the result is **break** the evaluation will terminate early with a successful result (EE-MORE-BREAK).

$$\begin{array}{c}
 \text{EE-EMP} \frac{}{e; \underline{\rho}; \sigma \xRightarrow{\text{each}} \mathbf{success} \blacksquare; \sigma} \\
 \\
 \text{EE-MORE-SUCS} \frac{e; \sigma \rho \xRightarrow{\text{expr}} vres; \sigma'' \quad vres = \mathbf{success} v \vee vres = \mathbf{continue} \quad e; \underline{\rho}'; (\sigma'' \setminus \text{dom } \rho) \xRightarrow{\text{each}} vres'; \sigma'}{e; \underline{\rho}, \underline{\rho}'; \sigma \xRightarrow{\text{each}} vres'; \sigma'} \\
 \\
 \text{EE-MORE-BREAK} \frac{e; \sigma \rho \xRightarrow{\text{expr}} \mathbf{break}; \sigma'}{e; \underline{\rho}, \underline{\rho}'; \sigma \xRightarrow{\text{each}} \mathbf{success} \blacksquare; (\sigma' \setminus \text{dom } \rho)} \\
 \\
 \text{EE-MORE-EXC} \frac{e; \sigma \rho \xRightarrow{\text{expr}} exres; \sigma' \quad exres \in \{\mathbf{throw} v, \mathbf{return} v, \mathbf{fail}, \mathbf{error}\}}{e; \underline{\rho}, \underline{\rho}'; \sigma \xRightarrow{\text{each}} exres; (\sigma' \setminus \text{dom } \rho)}
 \end{array}$$

Generator expressions The evaluation relation for generator expressions has form $\boxed{g; \sigma \xRightarrow{\text{gexpr}} envres; \sigma'}$. For matching assignments the target right-hand side expression is evaluated first—propagating possible exceptional results (G-PAT-EXC) and then the value is matched against target pattern (G-PAT-SUCS).

$$\begin{array}{c}
 envres ::= \mathbf{success} \vec{\rho} \mid exres \\
 \\
 \text{G-PAT-SUCS} \frac{e; \sigma \xRightarrow{\text{expr}} \mathbf{success} v; \sigma' \quad \sigma' \vdash p \stackrel{?}{:=} v \xRightarrow{\text{match}} \vec{\rho}}{p := e; \sigma \xRightarrow{\text{gexpr}} \mathbf{success} \vec{\rho}; \sigma'} \\
 \\
 \text{G-PAT-EXC} \frac{e; \sigma \xRightarrow{\text{expr}} exres; \sigma'}{p := e; \sigma \xRightarrow{\text{gexpr}} exres; \sigma'}
 \end{array}$$

For enumerating assignments, each possible value in a collection is provided as a possible binding to the range variable in the output (G-ENUM-LIST, G-ENUM-SET); for maps in particular, only the keys are bounds

(G-ENUM-MAP). An error is raised if the result value is not a collection (G-ENUM-ERR), and exceptions are propagated as always (G-ENUM-Exc).

$$\begin{array}{c}
 \text{G-ENUM-LIST} \frac{e; \sigma \xRightarrow{\text{expr}} \mathbf{success} [v_1, \dots, v_n]; \sigma'}{x \leftarrow e; \sigma \xRightarrow{\text{gexpr}} \mathbf{success} [x \mapsto v_1], \dots, [x \mapsto v_n]; \sigma'} \\
 \\
 \text{G-ENUM-SET} \frac{e; \sigma \xRightarrow{\text{expr}} \mathbf{success} \{v_1, \dots, v_n\}; \sigma'}{x \leftarrow e; \sigma \xRightarrow{\text{gexpr}} \mathbf{success} [x \mapsto v_1], \dots, [x \mapsto v_n]; \sigma'} \\
 \\
 \text{G-ENUM-MAP} \frac{e; \sigma \xRightarrow{\text{expr}} \mathbf{success} (v_1 : v'_1, \dots, v_n : v'_n); \sigma'}{x \leftarrow e; \sigma \xRightarrow{\text{gexpr}} \mathbf{success} [x \mapsto v_1], \dots, [x \mapsto v_n]; \sigma'} \\
 \\
 \text{G-ENUM-ERR} \frac{e; \sigma \xRightarrow{\text{expr}} \mathbf{success} v; \sigma' \quad v = vb \vee v = k(\underline{v'}) \vee v = \blacksquare}{x \leftarrow e; \sigma \xRightarrow{\text{gexpr}} \mathbf{error}; \sigma'} \\
 \\
 \text{G-ENUM-Exc} \frac{e; \sigma \xRightarrow{\text{expr}} \text{exres}; \sigma'}{x \leftarrow e; \sigma \xRightarrow{\text{gexpr}} \text{exres}; \sigma'}
 \end{array}$$

Pattern Matching

The pattern matching relation $\boxed{v \vdash \sigma \stackrel{?}{\underset{\text{match}}{\Longrightarrow}} \rho}$ takes as input the current store σ , a pattern p and a target value v and produces a sequence of compatible environments that represent possible bindings of variables to values. Pattern matching a basic value against target value produces a single environment that does not bind any variable (\square) if the target value is the same basic value (P-VAL-SUCS) and otherwise does not produce any binding environment (ε), (P-VAL-FAIL).

$$\begin{array}{c}
 \text{P-VAL-SUCS} \frac{}{\sigma \vdash vb \stackrel{?}{\underset{\text{match}}{\Longrightarrow}} \square} \quad \text{P-VAL-FAIL} \frac{v \neq vb}{\sigma \vdash vb \stackrel{?}{\underset{\text{match}}{\Longrightarrow}} \varepsilon}
 \end{array}$$

Pattern matching against a variable depends on whether the variable already exists in the current store. If it is assigned in the current store, then the target value must be the assigned value to return a possible binding (P-VAR-UNI) and otherwise failing with no bindings (P-VAR-FAIL); if it is not in the current store, it will simply bind the variable to the target value (P-VAR-BIND).

$$\begin{array}{c}
 \text{P-VAR-UNI} \frac{x \in \text{dom } \sigma \quad v = \sigma(x)}{\sigma \vdash x := v \xrightarrow[\text{match}]{} \square} \quad \text{P-VAR-FAIL} \frac{x \in \text{dom } \sigma \quad v \neq \sigma(x)}{\sigma \vdash x := v \xrightarrow[\text{match}]{} \varepsilon} \\
 \\
 \text{P-VAR-BIND} \frac{x \notin \text{dom } \sigma}{\sigma \vdash x := v \xrightarrow[\text{match}]{} [x \mapsto v]}
 \end{array}$$

When pattern matching against a constructor pattern, it is first checked whether the target value has the same constructor. If it does, then the sub-patterns are matched against the contained values of the target value, merging their resulting environments (P-CONS-SUCS), and otherwise failing with no bindings (P-CONS-FAIL). The merging procedure is described formally later in this section, but it intuitively it takes the union of all bindings that have consistent assignments to the same variables.

$$\begin{array}{c}
 \text{P-CONS-SUCS} \frac{\sigma \vdash p_1 := v_1 \xrightarrow[\text{match}]{} \underline{\rho_1} \quad \dots \quad \sigma \vdash p_n := v_n \xrightarrow[\text{match}]{} \underline{\rho_n}}{\sigma \vdash k(\underline{p}) := k(\underline{v}) \xrightarrow[\text{match}]{} \text{merge}(\underline{\rho_1}, \dots, \underline{\rho_n})} \\
 \\
 \text{P-CONS-FAIL} \frac{v \neq k(\underline{v'})}{\sigma \vdash k(\underline{p}) := v \xrightarrow[\text{match}]{} \varepsilon}
 \end{array}$$

When pattern matching against a typed labeled pattern, the target value is checked to have a compatible type—failing with no bindings otherwise (P-TYPE-FAIL)—and then then the inner pattern is matched against the same value. The result of the sub-pattern match is merged with the environment where the target value is bound to the label variable (P-TYPE-SUCS).

$$\text{P-TYPE-SUCS} \frac{v : t' \quad t' <: t \quad \sigma \vdash p := v \xrightarrow[\text{match}]{} \underline{\rho}}{\sigma \vdash t x : p := v \xrightarrow[\text{match}]{} \text{merge}([x \mapsto v], \underline{\rho})}$$

$$\text{P-TYPE-FAIL} \frac{v : t' \quad t' \not\prec t}{\sigma \vdash t x : p \stackrel{?}{\underset{\text{match}}{=} v} \Longrightarrow \varepsilon}$$

Pattern matching against a list pattern first checks whether the target value is a list—otherwise failing (P-LIST-FAIL)—and then pattern matches against the sub-patterns returning their result (P-LIST-SUCS).

$$\text{P-LIST-SUCS} \frac{\sigma \vdash \star p \stackrel{?}{\underset{\text{match}\star}{=} \underline{v}} \mid \emptyset \xrightarrow{\parallel} \underline{\rho}}{\sigma \vdash [\star p] \stackrel{?}{\underset{\text{match}}{=} [\underline{v}]} \Longrightarrow \underline{\rho}} \quad \text{P-LIST-FAIL} \frac{v \neq [\underline{v}']}{\sigma \vdash [\star p] \stackrel{?}{\underset{\text{match}}{=} v} \Longrightarrow \varepsilon}$$

Pattern matching against a set pattern is analogous to pattern matching against list patterns (P-SET-SUCS, P-SET-FAIL).

$$\text{P-SET-SUCS} \frac{\sigma \vdash \star p \stackrel{?}{\underset{\text{match}\star}{=} \underline{v}} \mid \emptyset \xrightarrow{\{\}} \underline{\rho}}{\sigma \vdash \{\star p\} \stackrel{?}{\underset{\text{match}}{=} \{\underline{v}\}} \Longrightarrow \underline{\rho}} \quad \text{P-SET-FAIL} \frac{v \neq \{\underline{v}'\}}{\sigma \vdash \{p\} \stackrel{?}{\underset{\text{match}}{=} v} \Longrightarrow \varepsilon}$$

Negation pattern $!p$ matching succeeds with no variables bound if the sub-pattern p produces no binding environment (P-NEG-SUCS), and otherwise fails (P-NEG-FAIL).

$$\text{P-NEG-SUCS} \frac{\sigma \vdash p \stackrel{?}{\underset{\text{match}}{=} v} \Longrightarrow \varepsilon}{\sigma \vdash !p \stackrel{?}{\underset{\text{match}}{=} v} \Longrightarrow []} \quad \text{P-NEG-FAIL} \frac{\sigma \vdash p \stackrel{?}{\underset{\text{match}}{=} v} \Longrightarrow \rho \quad \rho \neq \varepsilon}{\sigma \vdash !p \stackrel{?}{\underset{\text{match}}{=} v} \Longrightarrow \varepsilon}$$

Descendant pattern matching applies the sub-pattern against target value, and keeps applying the deep matching pattern against the children values, concatenating their results (P-DEEP). Notice here the similarities with the deep type-directed match operation in TRON (see Chapter 5); the operation in Rascal Light is however more general allowing use of the complete pattern matching capabilities available, in contrast to only types.

$$\text{P-DEEP} \frac{\sigma \vdash p \stackrel{?}{\underset{\text{match}}{=} v} \Longrightarrow \underline{\rho} \quad v'_1, \dots, v'_n = \text{children}(v) \quad \sigma \vdash /p \stackrel{?}{\underset{\text{match}}{=} v'_1} \Longrightarrow \underline{\rho'_1} \quad \dots \quad \sigma \vdash /p \stackrel{?}{\underset{\text{match}}{=} v'_n} \Longrightarrow \underline{\rho'_n}}{\sigma \vdash /p \stackrel{?}{\underset{\text{match}}{=} v} \Longrightarrow \underline{\rho}, \underline{\rho'_1}, \dots, \underline{\rho'_n}}$$

The star pattern matching relation has form

$$\sigma \vdash \underline{\star p} \stackrel{?}{:=} \underline{v} \mid \mathbb{V} \xrightarrow[\text{match}\star]{\langle \rangle \otimes} \underline{\rho}$$

which tries to match the sequence of patterns $\underline{\star p}$ on the left-hand side with sequence of provided values \underline{v} on the right-hand side; the relations is parameterized over the construction function $\langle \rangle$ and the partition relation \otimes , and because matching arbitrary elements patterns $\star x$ is non-deterministic we keep track of a set of values \mathbb{V} that have already been tried for the latest available variable.

If both the pattern sequence and the value sequence is empty, then we have successfully finished matching (PL-EMP-BOTH). Otherwise if any of the sequences finish while the other is non-empty we produce no possible bindings (PL-EMP-PAT, PL-EMP-VAL); an exemption is made for arbitrary match patterns because they can match empty sequences of values.

$$\begin{array}{c} \text{PL-EMP-BOTH} \frac{}{\sigma \vdash \varepsilon \stackrel{?}{:=} \varepsilon \mid \mathbb{V} \xrightarrow[\text{match}\star]{\langle \rangle \otimes} []} \quad \text{PL-EMP-PAT} \frac{}{\sigma \vdash \varepsilon \stackrel{?}{:=} \underline{v}, \underline{v}' \mid \mathbb{V} \xrightarrow[\text{match}\star]{\langle \rangle \otimes} \varepsilon} \\ \\ \text{PL-EMP-VAL} \frac{}{\sigma \vdash p, \underline{\star p} \stackrel{?}{:=} \varepsilon \mid \mathbb{V} \xrightarrow[\text{match}\star]{\langle \rangle \otimes} \varepsilon} \end{array}$$

When the initial element of the starred pattern sequence is an ordinary pattern then it is matched against the initial element of the value sequence, and the rest of the pattern sequence is matched against the rest of the value sequence. The results of both submatches are then merged together (PL-MORE-PAT).

$$\text{PL-MORE-PAT} \frac{\sigma \vdash p \stackrel{?}{:=} \underline{v} \xrightarrow[\text{match}]{\langle \rangle} \underline{\rho} \quad \sigma \vdash \underline{\star p} \stackrel{?}{:=} \underline{v}' \mid \mathbb{V} \xrightarrow[\text{match}\star]{\langle \rangle \otimes} \underline{\rho}'}{\sigma \vdash p, \underline{\star p} \stackrel{?}{:=} \underline{v}, \underline{v}' \mid \mathbb{V} \xrightarrow[\text{match}\star]{\langle \rangle \otimes} \text{merge}(\underline{\rho}, \underline{\rho}')}$$

Like ordinary variables, arbitrary matching patterns depend on whether the binding variable already exists in the current store. If the variable is assigned in the current store then either there must exist a

partition of values that has a matching subcollection (PL-MORE-STAR-UNI), or the matching fails producing any consistent binding environment (PL-MORE-STAR-PAT-FAIL, PL-MORE-STAR-VAL-FAIL).

$$\begin{array}{c}
\text{PL-MORE-STAR-UNI} \frac{x \in \text{dom } \sigma \quad \sigma(x) = \langle \underline{v'} \rangle \quad \underline{v} = \underline{v'} \otimes \underline{v''} \quad \sigma \vdash \star p \stackrel{?}{:=} \underline{v''} \mid \emptyset \xrightarrow[\text{match}\star]{\langle \rangle \otimes} \underline{\rho}}{\sigma \vdash \star x, \star p \stackrel{?}{:=} \underline{v} \mid \mathbb{V} \xrightarrow[\text{match}\star]{\langle \rangle \otimes} \underline{\rho}} \\
\text{PL-MORE-STAR-PAT-FAIL} \frac{x \in \text{dom } \sigma \quad \sigma(x) = \langle \underline{v'} \rangle \quad \nexists \underline{v''}. \underline{v} = \underline{v'} \otimes \underline{v''}}{\sigma \vdash \star x, \star p \stackrel{?}{:=} \underline{v} \mid \mathbb{V} \xrightarrow[\text{match}\star]{\langle \rangle \otimes} \varepsilon} \\
\text{PL-MORE-STAR-VAL-FAIL} \frac{x \in \text{dom } \sigma \quad \sigma(x) = \underline{v} \quad \underline{v} \neq \langle \underline{v'} \rangle}{\sigma \vdash \star x, \star p \stackrel{?}{:=} \underline{v''} \mid \mathbb{V} \xrightarrow[\text{match}\star]{\langle \rangle \otimes} \varepsilon}
\end{array}$$

If the variable is not in the current store then there are two options: either i) there still exist a partition that is possible to try, or ii) we have exhausted all possible partitions of the value sequence. In the first case, an arbitrary partition is bound to the target variable and the rest of the patterns are matched against the rest of the values, merging their results; additionally, the other partitions are also tried concatenating their results with the merged one (PL-MORE-STAR-RE). In the exhausted case, the pattern match produces no bindings (PL-MORE-STAR-EXH).

$$\begin{array}{c}
\text{PL-MORE-STAR-RE} \frac{x \notin \text{dom } \sigma \quad \underline{v} = \underline{v'} \otimes \underline{v''} \quad \underline{v'} \notin \mathbb{V} \quad \sigma \vdash \star p \stackrel{?}{:=} \underline{v''} \mid \emptyset \xrightarrow[\text{match}\star]{\langle \rangle \otimes} \underline{\rho} \quad \sigma \vdash \star x, \star p \stackrel{?}{:=} \underline{v} \mid \mathbb{V} \cup \underline{v'} \xrightarrow[\text{match}\star]{\langle \rangle \otimes} \underline{\rho'}}{\sigma \vdash \star x, \star p \stackrel{?}{:=} \underline{v} \mid \mathbb{V} \xrightarrow[\text{match}\star]{\langle \rangle \otimes} \text{merge}([x \mapsto \underline{v'}], \underline{\rho}), \underline{\rho'}} \\
\text{PL-MORE-STAR-EXH} \frac{x \notin \text{dom } \sigma \quad \nexists \underline{v'}, \underline{v''}. \underline{v} = \underline{v'} \otimes \underline{v''} \wedge \underline{v'} \notin \mathbb{V}}{\sigma \vdash \star x, \star p \stackrel{?}{:=} \underline{v} \mid \mathbb{V} \xrightarrow[\text{match}\star]{\langle \rangle \otimes} \varepsilon}
\end{array}$$

Merging a sequence of possible variable bindings produces a sequence containing consistent variable bindings from the sequence: that

is, all possible environments that assign consistent values to the same variables are merged.

$$\begin{aligned} \text{merge}(\varepsilon) &= [] \\ \text{merge}(\underline{\rho}, \underline{\rho}'_1, \dots, \underline{\rho}'_n) &= \text{merge-pairs}(\underline{\rho} \times \text{merge}(\underline{\rho}'_1, \dots, \underline{\rho}'_n)) \end{aligned}$$

$$\text{merge-pairs}(\langle \rho_1, \rho'_1 \rangle, \dots, \langle \rho_n, \rho'_n \rangle) = \text{merge-pair}(\rho_1, \rho'_1), \dots, \text{merge-pair}(\rho_n, \rho'_n)$$

$$\text{merge-pair}(\rho, \rho') = \begin{cases} \rho\rho' & \text{if } \forall x \in \text{dom } \rho \cap \text{dom } \rho'. \rho(x) = \rho'(x) \\ \varepsilon & \text{otherwise} \end{cases}$$

6.3 Semantics Properties

Backtracking is pure in Rascal (Light) programs, and so if evaluating a set cases produces **fail** as result, the initial state is restored.

Theorem 6.1 (Backtracking purity). *If $\underline{\text{cs}}; v; \sigma \xrightarrow[\text{cases}]{\text{CS}} \text{fail}; \sigma'$ then $\sigma' = \sigma$*

Strong typing is an important safety property of Rascal, which I capture by specifying a theorems with two result properties: one about the well-typedness of state, and the other about well-typedness of resulting values.

Theorem 6.2 (Strong typing). *Assume that semantic unary $\llbracket \ominus \rrbracket$ and binary*

operators $\llbracket \oplus \rrbracket$ are strongly typed. If $e; \sigma \xrightarrow[\text{expr}]{\mathcal{E}} \text{vres}; \sigma'$ and there exists a type

t such that $\mathcal{T} \frac{}{v : t}$ for each value in the input store $v \in \text{img } \sigma$, then

1. *There exists a type t' such that $v' : t'$ for each value in the result store $v' \in \text{img } \sigma'$.*
2. *If the result value vres is either **success** v'' , **return** v'' , or **throw** v'' , then there exists a type t'' such that $v'' : t''$.*

Consider an augmented version of the operational semantics where each execution relation is annotated with *partiality fuel* [Amin and Rompf, 2017]—represented by a superscript natural number n —which specifies the maximal number of recursive premises allowed in the derivation. The fuel is subtracted on each recursion, resulting in a **timeout** value when it reaches zero, and the rule set is amended by congruence rules which propagate **timeout** results from the recursive premises to the conclusion.

$$vtres ::= vres \mid \mathbf{timeout}$$

For this version of the semantics, we can specify the property that execution will either produce a result or it will timeout; that is, the semantics does not get stuck.

Theorem 6.3 (Partial progress). *It is possible to construct a derivation $e; \sigma \xrightarrow[\text{expr}]^n vtres; \sigma'$ for any input expression e , well-typed store σ and fuel n .*

Finally, consider a subset of the expression language described by syntactic category e_{fin} , where **while**-loops, **solve**-loops and function calls are disallowed, and similarly with all traversal strategies except **bottom-up** and **bottom-up-break**. This subset is known to be terminating:

Theorem 6.4 (Terminating expressions). *There exists n such that derivation \mathcal{E} $e_{\text{fin}}; \sigma \xrightarrow[\text{expr}]^n vres; \sigma'$ has a result $vres$ which is not **timeout** for expression e_{fin} in the terminating subset.*

Remark 6.1. Why is the **top-down** traversal strategy potentially non-terminating while the **bottom-up** traversal strategy is terminating? The answer lies in that the **top-down** traverses the children of the target expression after evaluating the required case, which makes it possible to keep add children ad infinitum as illustrated by the following example:

```

1   data Nat = zero() | succ(Nat pred);
2
3   Nat infincrement(Nat n) =
4     top-down visit(n) {
5       case succ(n) => succ(succ(n))
6     }

```

6.4 Formal Semantics, Types and Intended Semantics

This formal specification of Rascal Light was largely performed when the official type checker of Rascal (developed by CWI Amsterdam) was only at the experimental stage. Useful extension of this formalization include a deductive type system that includes polymorphic aspects of Rascal to prove the consistency of the type system, and type safety with regards to the dynamic semantics provided in this formalization.

The presented semantics was checked against the Rascal implementation in various ways: the existing implementation was used as a reference point, rules for individual were held up against the documentation, and correspondence with the Rascal developers was used to clarify remaining questions. A more formal way to check whether the captured semantics is the intended one, is to construct an independent formal semantics which is related to the natural semantics presented in this chapter. This could be an axiomatic semantics Hoare [1969], which modularly specifies for each construct the intended effects using logical formulae as pre and post-conditions, without necessarily specifying how each construct is evaluated. The natural semantics would then be checked to satisfy the axiomatic semantics for each construct, which will further increase confidence that the captured semantics is the intended one.

6.5 Recap

I presented the formalization of a large subset of the operational part of Rascal Klint et al. [2009, 2011], called Rascal Light. The formalization was primarily based on the available open source implementation⁴ and the accompanying documentation⁵, and personal correspondence with the developers further clarified previous ambiguities and mismatches.

Unlike TRON, Rascal Light is meant to be both a formal and practical language subset, and so supports a much broader range of both ordinary language features—e.g., function definitions (including recursion) loops (including control flow operators), case analysis, and exceptions—and high-level transformation language features such as generic traversals with strategies, backtracking and powerful pattern matching. This subset has been verified against properties specifying strong typing and safety, which provides further trustworthiness.

⁴<https://github.com/usethesource/rascal>

⁵<http://tutor.rascal-mpl.org/>

Chapter 7

Verifying Rascal Transformations using Inductive Refinement Types

The established rich features in Rascal—first-class traversal, powerful pattern matching, backtracking, fixed-point iteration— make it easy to express complex programs in a succinct and maintainable manner. Verification-wise, availability of these features is a double-edged. They provide an opportunity to develop techniques that provide more efficient and precise results for the different constructs, e.g., by better propagating information during traversal and provide specialized abstractions for the different collections available. They however also bring along challenges:

1. The manipulated data represents often large models or programs, relying on large inductive algebraic data types and collections such as lists, sets and maps.
2. There is a non-trivial control flow stemming from the presence of exceptions and backtracking—allowing potentially non-local jumps during execution—and the traversal constructs, where the control flow depends on the type and shape of matched data in addition to the current state.

In this chapter, we will discuss how we combine existing techniques in abstract interpretation in order to deal with the aforementioned challenges. Concretely, we will discuss the following contributions:

- Modular domain design for abstract shape domains, that allows easily extending and replacing abstraction for concrete types of

elements, e.g. extending the abstraction for lists to include length of elements in addition to shapes.

- Formal Schmidt-style abstract *operational* semantics [Schmidt, 1998] for a subset of Rascal, tuning the idea of trace-based memoization to support the various expressive constructs in Rascal.
- An abstract interpretation-based static analysis tool for Rascal Light that supports inferring inductive refinements types that characterise the type and shape of manipulated data.
- An evaluation of the technique on a series of realistic program transformations, including a desugaring transformation and a DSL transpilation transformation taken from real Rascal projects.

These contributions together show the feasibility of using abstract interpretation for constructing static analyses for expressive transformation languages and properties. We hope that our presented adaptation of Schmidt-style abstract interpretation provides inspiration beyond Rascal, and that more language designers will reuse their existing interpreters of other languages to build similar tools.

7.1 Motivating Example

As a motivating example, we will use a multiplicative expression simplification transformation (Figure 7.1) which eliminates multiplications by 0 and 1. The example is inspired by the constant folding transformation common in many compilers, and is implemented in Rascal using bottom-up traversal for the required rewriting of sub-expressions.

```

1  data Nat = zero() | suc(Nat pred);
2  data Expr = var(str nm) | cst(Nat vl) | mult(Expr el, Expr er);
3
4  Expr simplify(Expr expr) =
5    bottom-up visit (expr) {
6      case mult(cst(zero()), y) => cst(zero())
7      case mult(x, cst(zero())) => cst(zero())
8      case mult(cst(suc(zero())), y) => y
9      case mult(x, cst(suc(zero()))) => x
10   };

```

Figure 7.1: Expression Simplification

For similar transformations to this one, there are typically multiple properties we want to verify:

Type properties concern that all executions of the program are type and scope-safe. In our example, this would be that the result of the visitor is of type `Expr` as expected by the `simplify` function.

Shape properties concern that the output data of our program has a particular structure. For our example, this could be no multiplication sub-expression `mult` in its output which has `cst(zero())` or `cst(suc(zero()))` as a sub-expression.

The presented technique allows verifying such type and shape properties by inferring inductive refinements of algebraic data type definitions, e.g., the `Expr` data type used in the motivating example. To illustrate this, we will abstractly interpret our motivating example, providing a high-level overview of the required steps:

1. Let the initial abstract store be $\hat{\sigma} = [\text{expr} \mapsto \text{Expr}]$, so the input variable `expr` abstractly represents any value of type `Expr`.
2. Execution proceeds by evaluating the **visit**-statement on Line 5; since the traversal is bottom-up, it starts traversing the contained values (children) and then the applies the required cases to the top-level value.
 - a) The children of a datatype value dependent on the particular constructor used, and so the interpreter proceed to unfold `expr` to the possible constructor values: `mult(Expr, Expr)`, `cst(Nat)` and `var(str)`.
 - b) The interpreter considers each of the possible constructor values separately, merging their results using the least upper bound operator when execution is finished.
3. In case the analyzed value is `mult(Expr, Expr)` then the children have type `Expr` twice.
 - a) Traversing the first child requires us to recursively consider the type `Expr` again, which would result in non-termination if the above steps were naively repeated.

- b) Instead we will apply a version of trace memoization [Schmidt, 1998] where we return the value of the previous recursive traversal to compute a new approximation of the result, and then recompute the recursive traversal again until a fixed-point is reached.
 - c) Because this is the first recursion in the traversal, we will return the bottom value \perp is used as state and the traversal continues evaluation at the other branches.
4. In case the analyzed value is $\text{cst}(\text{Nat})$ then there is a single child of type Nat .
 - Traversal proceeds with the type Nat , using a similar procedure of unfolding, merging, memoization and fixed-point iteration as above.
 - Since there is no case that affects the type Nat , the traversal produces $\text{fail Nat}; \hat{\sigma}$ as a result, signifying that the evaluation did not match any case. Notice, that the failure result of case evaluation, additionally contains an abstract value Nat —not present in the concrete executor—which signifies a refinement of the input value.
 - We have to use the result to reconstruct an updated value using the cst constructor for the rest of the traversal, which in this case trivially succeeds producing the same value as our input $\text{cst}(\text{Nat})$.
 - Since $\text{cst}(\text{Nat})$ does not match any case either, we get $\text{fail cst}(\text{Nat}); \hat{\sigma}$ as a result for this branch.
 5. In case of $\text{var}(\text{str})$ the traversal proceeds in the same way as the cst case, with none of the cases matching and a result value of $\text{fail var}(\text{str}); \hat{\sigma}$.
 6. We merge the result of the three branches which produces $\perp \sqcup (\text{fail cst}(\text{Nat}); \hat{\sigma}) \sqcup (\text{fail var}(\text{str}); \hat{\sigma}) = \text{fail Expr\#1}; \hat{\sigma}$ as result, where:

$$\text{refine Expr\#1} = \text{cst}(\text{Nat}) \mid \text{var}(\text{str})$$

This refinement of input values on failure is important to maintain precision during traversals, which is what allows us to produce

precise result shapes. In particular, if we did not produce a refinement as above, we would only have the input value shape (Expr) and so we would not be able to infer any useful shape property.

7. Since we got a new result that was more over-approximating than the previous state—i.e., $\perp \sqsubseteq (\mathbf{fail} \text{ Expr\#1}; \hat{\sigma})$ —the trace memoization technique requires us to re-execute the **visit**-statement, in order to maintain soundness (via over-approximation).
8. We unfold to three possible constructor cases again: $\text{mult}(\text{Expr}, \text{Expr})$, $\text{cst}(\text{Nat})$ and $\text{var}(\mathbf{str})$. We only have to re-consider the result for $\text{mult}(\text{Expr}, \text{Expr})$ since the result for the other two cases remain the same.
9. In the case $\text{mult}(\text{Expr}, \text{Expr})$, so the children are of type Expr , twice.
 - a) This time on the recursive traversal of Expr , trace memoization will return the state from the previous iteration $\mathbf{fail} \text{ Expr\#1}; \hat{\sigma}$.
 - b) We reconstruct the expression with the refined children values, to get value $\text{mult}(\text{Expr\#1}, \text{Expr\#1})$
 - c) We unfold and match the cases, applying the right-hand sides and get the following results:
 - i. On the first two cases (Line 6 and Line 7), we get **success** $\text{cst}(\text{zero}()); \hat{\sigma}$ as a result.
 - ii. On the second two cases (Line 8 and Line 9), we get **success** $\text{Expr\#2}; \hat{\sigma}$ as a result, where:

$$\mathbf{refine} \text{ Expr\#2} = \text{cst}(\text{suc}(\text{Nat})) \mid \text{var}(\mathbf{str})$$

since both the left and right component are refined to not be $\text{cst}(\text{zero}())$ after matching the previous patterns, and the multiplication with $\text{cst}(\text{suc}(\text{zero}()))$ is rewritten to only contain the other operand.

- iii. In case of failure, we get **fail** $\text{mult}(\text{Expr\#3}, \text{Expr\#3}); \hat{\sigma}$ as a result, where:

$$\mathbf{refine} \text{ Expr\#3} = \text{cst}(\text{suc}(\text{suc}(\text{Nat}))) \mid \text{var}(\mathbf{str})$$

since both components are also refined to not be $\text{suc}(\text{zero}())$ after matching the previous patterns.

- d) Combining the results for the different cases by taking the least upper bound gives us

$$(\mathbf{success} \text{ Expr\#1}; \hat{\sigma}), (\mathbf{fail} \text{ mult}(\text{Expr\#3}, \text{Expr\#3}); \hat{\sigma})$$

as a result state which keeps track of separate shapes and stores for each type of result (success and failure).

10. Then by combining all constructor branch results we get:

$$(\mathbf{success} \text{ Expr\#1}; \hat{\sigma}), (\mathbf{fail} \text{ Expr\#4}; \hat{\sigma})$$

where we have

$$\mathbf{refine} \text{ Expr\#4} = \text{cst}(\text{Nat}) \mid \text{var}(\mathbf{str}) \mid \text{mult}(\text{Expr\#3}, \text{Expr\#3})$$

11. Since we got a new more over-approximating result from previous execution of **visit**-statement—i.e., $(\mathbf{fail} \text{ Expr\#1}) \sqsubseteq (\mathbf{success} \text{ Expr\#1}; \hat{\sigma}), (\mathbf{fail} \text{ Expr\#4}; \hat{\sigma})$ —the fixed-point iteration must be continued. To ensure termination in a finite number of steps, we use widening between the previous result state and the current result state $(\mathbf{fail} \text{ Expr\#1}; \hat{\sigma}) \nabla (\mathbf{success} \text{ Expr\#1}; \hat{\sigma}), (\mathbf{fail} \text{ Expr\#4}; \hat{\sigma}) = (\mathbf{success} \text{ Expr\#1}; \hat{\sigma}), (\mathbf{fail} \text{ Expr\#4}; \hat{\sigma})$, which produces the same state as the new result because the previous result did not have any recursively defined refinement.

12. After another iteration for the fixed-point computation of **visit**-statement we get $(\mathbf{success} \text{ Expr\#1}; \hat{\sigma}), (\mathbf{fail} \text{ Expr\#5}; \hat{\sigma})$ as result, where:

$$\mathbf{refine} \text{ Expr\#5} = \text{cst}(\text{Nat}) \mid \text{var}(\mathbf{str}) \mid \text{mult}(\text{Expr\#6}, \text{Expr\#6})$$

$$\mathbf{refine} \text{ Expr\#6} = \text{cst}(\text{suc}(\text{suc}(\text{Nat}))) \mid \text{var}(\mathbf{str}) \mid \text{mult}(\text{Expr\#3}, \text{Expr\#3})$$

13. By widening of the new result with the previous one, we get a result state $(\mathbf{success} \text{ Expr\#1}; \hat{\sigma}), (\mathbf{success} \text{ Expr\#6}; \hat{\sigma})$. This uses the widening algorithm presented in Section 7.3, to calculate the widening on refinements: $\text{Expr\#4} \nabla \text{Expr\#5} = \text{Expr\#7}$ with previous inductive refinement definitions:

$$\mathbf{refine} \text{ Expr\#3} = \text{cst}(\text{suc}(\text{suc}(\text{Nat}))) \mid \text{var}(\mathbf{str})$$

$$\mathbf{refine} \text{ Expr\#4} = \text{cst}(\text{Nat}) \mid \text{var}(\mathbf{str}) \mid \text{mult}(\text{Expr\#3}, \text{Expr\#3})$$

$$\mathbf{refine} \text{ Expr\#5} = \text{cst}(\text{Nat}) \mid \text{var}(\mathbf{str}) \mid \text{mult}(\text{Expr\#6}, \text{Expr\#6})$$

$$\mathbf{refine} \text{ Expr\#6} = \text{cst}(\text{suc}(\text{suc}(\text{Nat}))) \mid \text{var}(\mathbf{str}) \mid \text{mult}(\text{Expr\#3}, \text{Expr\#3})$$

and new refinement definitions:

$$\begin{aligned} \mathbf{refine} \text{ Expr\#7} &= \text{cst}(\text{Nat}) \mid \text{var}(\mathbf{str}) \mid \text{mult}(\text{Expr\#8}, \text{Expr\#8}) \\ \mathbf{refine} \text{ Expr\#8} &= \text{cst}(\text{suc}(\text{suc}(\text{Nat}))) \mid \text{var}(\mathbf{str}) \mid \text{mult}(\text{Expr\#8}, \text{Expr\#8}) \end{aligned}$$

The core intuition behind the refinement is to combine the set of constructors, recursively merging the arguments of the common ones; additionally, when merging two inductive refinements, we create a new refinement replacing the references to the old ones with it. In our example, this means that merging Expr#3 and Expr#6 creates Expr#8 which has the same set of constructors, but the recursive references are to Expr#8 instead of Expr#3 or Expr#6.

14. After another iteration, we get the same result $(\mathbf{success} \widehat{\text{Expr\#1}}, (\mathbf{fail} \text{ Expr\#7}; \hat{\sigma}))$ and so we have reached a fixed-point.
15. On failure, the **visit**-expression after traversal acts as a kind of identity transformation producing a successful rules containing the original input value. In our abstract interpreter we will do the same on failure, but use the refined the result; we therefore convert the possible failure state $\mathbf{fail} \text{ Expr\#7}; \hat{\sigma}$ to $\mathbf{success} \text{ Expr\#7}; \hat{\sigma}$ and merge with the other success state $\mathbf{success} \text{ Expr\#1}; \hat{\sigma}$ to get $\mathbf{success} \text{ Expr\#7}; \hat{\sigma}$ as final result for the **visit**-expression.

On the result inductive shape, we can then perform checks whether our target properties are satisfied. For example, we can easily see that all constants have natural numbers of shape $\text{suc}(\text{suc}(\text{Nat}))$ (that is, ≥ 2), and so the simplification procedure correctly eliminates all multiplications by 0 and 1.

7.2 Formal Language

The presented technique is meant to work with all of Rascal Light, but to keep the formal presentation concise we will consider a further subset, called Rascal UltraLight. We will delimitate this subset in this section, and then use it to describe the general technique we use to derive abstract operational semantics from the concrete operational semantics for

Rascal Light.¹ The constructs in this subset get their formal semantics directly from the Rascal Light concrete semantics presented in Chapter 6, but we will recall the rules when appropriate.

We consider the following value types: integers (**int**), finite sets (**set**(t)) and algebraic data types (at). Each algebraic data type, at associates a set of unique constructors, so that each constructor $k(\underline{t})$ has a fixed set of typed parameters.

$$vt \in \text{ValueType} ::= \mathbf{int} \mid \mathbf{set}\langle t \rangle \mid at \in \text{TyCons}$$

Like the types in Rascal, there is support for subtyping where we have **void** and **value** as bottom and top types specifically.

$$t \in \text{Type} ::= \mathbf{void} \mid vt \mid \mathbf{value}$$

The value domain Value is generated by the least fixedpoint of the following semantic equations for the types:

$$\begin{aligned} \llbracket \mathbf{void} \rrbracket &= \emptyset & \llbracket \mathbf{int} \rrbracket &= \mathbb{Z} & \llbracket \mathbf{set}\langle t \rangle \rrbracket &= \wp(\llbracket t \rrbracket) & \llbracket \mathbf{value} \rrbracket &= \bigcup_{vt} \llbracket vt \rrbracket \\ \llbracket at \rrbracket &= \left\{ k(\underline{v}) \mid k(\underline{t}) \text{ is a constructor of } at \wedge \forall i. v_i \in \llbracket t_i \rrbracket \right\} \end{aligned}$$

Expression language We consider the following subset of Rascal expressions:

$$\begin{aligned} e ::= & n \in \mathbb{Z} \mid x \in \text{Var} \mid e \otimes e \mid x = e \mid e; e \mid k(\underline{e}) \mid \{e\} \\ & \mid \mathbf{fail} \mid \mathbf{if } e \mathbf{ then } e \mathbf{ else } e \mid \mathbf{bottom-up visit } e \underline{cs} \\ cs ::= & \mathbf{case } p \Rightarrow e \quad p ::= x \mid k(\underline{p}) \mid \{\star p\} \quad \star p ::= p \mid \star x \end{aligned}$$

The considered subset includes integer constants, variables, binary operators, variable assignments, sequencing, constructor expressions, set literal expressions, matching failure, conditionals and bottom-up visitors. The pattern language consists of either variable patterns x , constructor patterns $k(\underline{p})$ or set patterns $\{\star p\}$ which internally consists of either ordinary patterns or star patterns $\star x$, which match a subcollection and binds the value to the variable x .

¹We will even for the Rascal UltraLight subset focus on presenting the high-level ideas behind the general technique and not mechanically translate every rule and auxiliary function.

7.3 Abstract Domains

Our abstract domain design focuses on providing a modular way to compose individual abstract domains, such that it supports the following properties:

Replacement Allowing changes to the abstract domain for a particular concrete type of elements without affecting the rest of the complete abstract domain, e.g. replacing the interval abstract domain with the congruence abstract domain for integers.

Extension Adding abstract domains for new types of values, e.g. supporting an abstraction for lists of elements.

Operational Compositionality Domain operations (e.g., least upper bound, subsumption checking and widening) can be performed effectively in a compositional fashion, e.g., that the widening for lists does not depend on the definition of the particular widening used for integers.

Our goal is to construct an abstract value domain $\widehat{v} \in \widehat{\text{Value}}$ which captures key shape and type properties for the different values in our concrete domain (integers, set, algebraic data types). As with our motivating example, this includes capturing inductive shapes represented by inductive data type refinements of form:

$$\mathbf{refine} \text{ } at\#r = k_1(\widehat{v}_1) \mid \cdots \mid k_n(\widehat{v}_n)$$

where the left hand side defines a refinement with name $at\#r$ that refines a datatype at with a subset of possible constructors on the right hand side that have abstract value components refining the shape of elements that are abstracted and can possibly inductively refer to $at\#r$ and other refinements.

The modular abstract domain design generalizes parameterized abstract domains initially suggested by Cousot [2003] to follow a design inspired by the modular construction of types and domains suggested in various literature in type theory and domain theory [Scott, 1976; Smyth and Plotkin, 1982; Winskel, 1993b; Löh and Magalhães, 2011; Chapman et al., 2010; Backhouse et al., 2007; Benke et al., 2003]. The idea is to define domains parametrically—i.e. of form $\widehat{F}(\widehat{E})$ so that abstract domains for subcomponents are taken as parameters, and instead of having closed recursion the recursive references are replaced by an abstract

domain parameter which is decided later. After defining the abstract domains parametrically for each concrete domain we wish to abstract, we can then use standard domain combinators [Winskel, 1993b] —products, sums, fixed points—to combine the various domains into our target abstract value domain. We assume it is possible to pass information for operations recursively as long as it is done in a parametric fashion, so the information only affects operations for elements of the same domain; to avoid an overly abstract presentation we will only mention the relevant information to be passed along when needed.

Integer shape domain

To abstract over integer values, we use the classic interval domain [Cousot and Cousot, 1977]. We chose this domain because it is well-defined, familiar and easy to use with all arithmetic operations; because of our modular design, it is easy to replace it when needed with another suitable numeric abstract domain such as sign, parity or congruence².

An interval $[l; u] \in \widehat{\text{Interval}}$ represents a continuous set of integers $\{n \in \mathbb{Z} \mid l \leq n \leq u\}$ and we let $n \in \mathbb{Z}$ refer to the singleton interval $[n; n]$. The $\widehat{\text{Interval}}$ domain forms an infinite lattice, with the following operations:

$$\begin{aligned} \perp_{\widehat{\text{Interval}}} &= [\infty; -\infty] & \top_{\widehat{\text{Interval}}} &= [-\infty; \infty] \\ [l_1; u_1] \sqsubseteq_{\widehat{\text{Interval}}} [l_2; u_2] & \text{ iff } l_2 \leq l_1 \wedge u_1 \leq u_2 \\ [l_1; u_1] \sqcup_{\widehat{\text{Interval}}} [l_2; u_2] &= [\min(l_1, l_2); \max(u_1, u_2)] \\ [l_1; u_1] \sqcap_{\widehat{\text{Interval}}} [l_2; u_2] &= [\max(l_1, l_2); \min(u_1, u_2)] \end{aligned}$$

The $\widehat{\text{Interval}}$ also admits a *widening*, which is an upper bound of elements that ensures finite convergence to a result when iteratively applied:

$$[l_1; u_1] \nabla_{\widehat{\text{Interval}}} [l_2; u_2] = [\text{if}(l_2 < l_1, -\infty, l_1); \text{if}(u_1 < u_2, \infty, u_1)]$$

²In theory one could allow even relational domains such as octagon, but it would require a redesign of the presented modular framework.

Proposition 7.1. *The $\widehat{\text{Interval}}$ domain forms a Galois connection with the power set of integers $\wp(\mathbb{Z}) \xleftrightarrow[\alpha_{\widehat{\text{Interval}}}]{} \gamma_{\widehat{\text{Interval}}}$, where:*

$$\alpha_{\widehat{\text{Interval}}}(N) = [\min(N); \max(N)] \quad \gamma_{\widehat{\text{Interval}}}([l; u]) = \{n \in \mathbb{Z} \mid l \leq n \leq u\}$$

Set Shape Domain

Recall that the goal of our analysis is to capture type and shape properties, and so we must pick an abstraction that can reason about the shape of elements inside sets (and other collections) to preserve precision, especially for inductive refinements; additionally, operations on collections like pattern matching and iteration rely on the cardinality of collections and it is likewise useful to provide an abstraction for this component.

We therefore abstract finite sets using a parameterized abstract domain, $\widehat{\text{SetShape}}(\widehat{E})$, consisting of a reduced product between a component that abstracts the shape of elements $\widehat{\text{SetContent}}(\widehat{E})$ and a component that abstracts the cardinality of the set using non-negative intervals $\widehat{\text{Interval}}^+$. The set shape content abstraction $\widehat{\text{SetContent}}(\widehat{E})$ is defined as follows:

$$\widehat{sc} \in \widehat{\text{SetContent}}(\widehat{E}) ::= \{\widehat{e}\}, \widehat{e} \in \widehat{E}$$

We will use the notation $\{\widehat{e}\}_{[l; u]}$, where $\widehat{e} \in \widehat{E}$ describes the content and $[l; u] \in \widehat{\text{Interval}}^+$ describes the cardinality to represent $(\{\widehat{e}\}, [l; u]) \in \widehat{\text{SetShape}}(\widehat{E})$.

Example 7.1. Using the $\widehat{\text{Interval}}$ domain as content abstract domain, we can abstract sets of integers using $\widehat{\text{SetShape}}(\widehat{\text{Interval}})$ as follows:

$$\alpha_{\widehat{\text{SS}}}(\{\{1, 42\}, 32\}) = \{[1; 42]\}_{[1; 2]} \quad \alpha_{\widehat{\text{SS}}}(\{\{-41\}, \emptyset\}) = \{-41\}_{[0; 1]}$$

▲

If the content parameter domain \widehat{E} forms a lattice (with widening), then we can lift all operations trivially to this domain:

$$\begin{aligned} \perp_{\widehat{\text{SC}}} &= \{\perp_{\widehat{E}}\} & \top_{\widehat{\text{SC}}} &= \{\top_{\widehat{E}}\} & \{\widehat{e}_1\} \sqsubseteq_{\widehat{\text{SC}}} \{\widehat{e}_2\} & \text{iff } \widehat{e}_1 \sqsubseteq_{\widehat{E}} \widehat{e}_2 \\ \{\widehat{e}_1\} \sqcup_{\widehat{\text{SC}}} \{\widehat{e}_2\} &= \{\widehat{e}_1 \sqcup_{\widehat{E}} \widehat{e}_2\} & \{\widehat{e}_1\} \cap_{\widehat{\text{SC}}} \{\widehat{e}_2\} &= \{\widehat{e}_1 \cap_{\widehat{E}} \widehat{e}_2\} \\ \{\widehat{e}_1\} \nabla_{\widehat{\text{SC}}} \{\widehat{e}_2\} &= \{\widehat{e}_1 \nabla_{\widehat{E}} \widehat{e}_2\} \end{aligned}$$

Given a Galois connection for the content domains $\wp(\mathbb{E}) \xleftrightarrow[\alpha_{\widehat{\mathbb{E}}}]^{\gamma_{\widehat{\mathbb{E}}}} \widehat{\mathbb{E}}$ we produce the following abstraction $\alpha_{\widehat{\mathbb{S}\mathbb{S}}}$ and concretization $\gamma_{\widehat{\mathbb{S}\mathbb{S}}}$ between the concrete set of elements domain $\wp(\wp(\mathbb{E}))$ and abstract set shape domain $\widehat{\text{SetShape}}(\widehat{\mathbb{E}})$:

$$\alpha_{\widehat{\mathbb{S}\mathbb{S}}}(ess) = \bigsqcup_{es \in ess} \{\alpha_{\widehat{\mathbb{E}}}(es)\}_{|es|}$$

$$\gamma_{\widehat{\mathbb{S}\mathbb{S}}}(\{\widehat{e}\}_{[l;u]}) = \{es \mid es \subseteq \gamma_{\widehat{\mathbb{E}}}(\widehat{e}) \wedge |es| \in \gamma_{\mathbb{I}}([l;u])\}$$

Theorem 7.1. *Assuming Galois connection for element domains $\wp(\mathbb{E}) \xleftrightarrow[\alpha_{\widehat{\mathbb{E}}}]^{\gamma_{\widehat{\mathbb{E}}}} \widehat{\mathbb{E}}$ then we have a Galois connection for the set shape domain $\wp(\wp(\mathbb{E})) \xleftrightarrow[\alpha_{\widehat{\mathbb{S}\mathbb{S}}}]^{\gamma_{\widehat{\mathbb{S}\mathbb{S}}}} \widehat{\text{SetShape}}(\widehat{\mathbb{E}})$*

Data Shape Domain

For the clear majority of transformations, we want to ensure that the output captures the desired type and shape properties at all levels of the abstract syntax, i.e., that a program is fully normalized or that all subcomponents of a model have been suitably translated. A particularly natural way to capture this is by representing the shape constraints inductively, which we do using inductive refinements.

Inductive refinements are a generalization of refinement types [Freeman and Pfenning, 1991; Xi and Pfenning, 1998; Rushby et al., 1998] using ideas from regular tree abstractions [Aiken and Murphy, 1991; Cousot and Cousot, 1995]. Like refinement types, we allow constraining the constructors and content—in our case using classical abstract domains like the aforementioned interval—and we build on the regular tree operations to make it possible to infer inductive shapes.

Concretely, our abstraction of data type values is done using the abstract domain $\widehat{\text{DataShape}}(\widehat{\mathbb{E}})$, which takes a parameter domain $\widehat{\mathbb{E}}$ that abstracts the shape of elements in constructors. We define the elements of the data shape domain as follows:

$$\widehat{d} \in \widehat{\text{DataShape}}(\widehat{\mathbb{E}}) ::= \perp_{\widehat{\mathbb{D}\mathbb{S}}} \mid at\#r \in \widehat{\text{Refinement}} \mid \top_{\widehat{\mathbb{D}\mathbb{S}}}$$

We have the bottom $\perp_{\widehat{\mathbb{D}\mathbb{S}}}$ and top $\top_{\widehat{\mathbb{D}\mathbb{S}}}$ elements—respectively representing no data types value and all data type values—and references to

inductive refinements $at\#r$ abstracting over the values of algebraic data type at . For our parametrized data shape domain, the corresponding definition of refinement $at\#r$ consists of a choice between a subset of constructors of at where the arguments are abstracted using the content domain \widehat{E} :

$$\mathbf{refine} \text{ } at\#r_1 = k_1(\widehat{e}_1) \mid \dots \mid k_n(\widehat{e}_n)$$

When a refinement only has a single constructor, we may inline the definition to simplify presentation.

Example 7.2. Given $\widehat{\text{Interval}}$ as the base domain, we can abstract values of data type

$$\mathbf{data} \text{ } \text{intoption} = \text{none}() \mid \text{some}(\mathbf{int})$$

as follows:

$$\alpha_{\widehat{\text{DS}}}(\{\text{none}(), \text{some}(-2), \text{some}(3)\}) = \text{intoption}\#m23$$

$$\alpha_{\widehat{\text{DS}}}(\{\text{some}(41), \text{some}(42)\}) = \text{intoption}\#4142$$

with refinement definitions:

$$\mathbf{refine} \text{ } \text{intoption}\#m23 = \text{none}() \mid \text{some}([-2; 3])$$

$$\mathbf{refine} \text{ } \text{intoption}\#4142 = \text{some}([41; 42])$$

We can abstract shape of inductive data types by using the least fixed point of our parameterized data shape domain, i.e., $\text{lfp } \widehat{X}.\widehat{\text{DataShape}}(\widehat{X})$.

For example, we can abstract value of data type

$$\mathbf{data} \text{ } \text{Nat} = \text{zero}() \mid \text{suc}(\text{Nat})$$

as follows:

$$\alpha_{\widehat{\text{DS}}}(\{\text{zero}(), \text{suc}(\text{suc}(\text{zero}()))\}) = \text{Nat}\#02$$

with refinement definitions:

$$\mathbf{refine} \text{ } \text{Nat}\#02 = \text{zero}() \mid \text{suc}(\text{Nat}\#1)$$

$$\mathbf{refine} \text{ } \text{Nat}\#1 = \text{suc}(\text{Nat}\#0) \quad \mathbf{refine} \text{ } \text{Nat}\#0 = \text{zero}()$$



The lattice and widening operations on abstract elements follow directly from the refinement definitions, which we will describe afterwards.

$$\begin{aligned}
\perp_{\widehat{\text{DS}}} \sqsubseteq_{\widehat{\text{DS}}} at\#r \sqsubseteq_{\widehat{\text{DS}}} \top_{\widehat{\text{DS}}} \quad at\#r_1 \sqsubseteq_{\widehat{\text{DS}}} at\#r_2 \text{ iff } at\#r_1 \sqsubseteq_{\widehat{\text{R}}} at\#r_2 \\
\perp \circlearrowleft_{\widehat{\text{DS}}} \widehat{d} = \widehat{d} \circlearrowleft_{\widehat{\text{DS}}} \perp = \widehat{d} \quad \top \circlearrowleft_{\widehat{\text{DS}}} \widehat{d} = \widehat{d} \circlearrowleft_{\widehat{\text{DS}}} \top = \top \\
at_1\#r_1 \circlearrowleft_{\widehat{\text{DS}}} at_2\#r_2 = \text{if}(at_1 = at_2, at_1\#r_1 \circlearrowleft_{\widehat{\text{R}}} at_2\#r_2, \top) \quad \circlearrowleft \in \{\sqcup, \nabla\} \\
\perp \sqcap_{\widehat{\text{DS}}} \widehat{d} = \widehat{d} \sqcap_{\widehat{\text{DS}}} \perp = \perp \quad \top \sqcap_{\widehat{\text{DS}}} \widehat{d} = \widehat{d} \sqcap_{\widehat{\text{DS}}} \top = \widehat{d} \\
at_1\#r_1 \sqcap_{\widehat{\text{DS}}} at_2\#r_2 = \text{if}(at_1 = at_2, at_1\#r_1 \sqcap_{\widehat{\text{R}}} at_2\#r_2, \perp)
\end{aligned}$$

Let the abstract content domain form a Galois connection with a concrete content domain $\wp(\text{E}) \xrightleftharpoons[\alpha_{\widehat{\text{DS}}}]^{\gamma_{\widehat{\text{DS}}}} \widehat{\text{E}}$, then we can create abstract $\alpha_{\widehat{\text{DS}}}$ and concretization $\gamma_{\widehat{\text{DS}}}$ functions between

$$\text{Data} = \left\{ k(\underline{v}) \mid \exists at.k(\underline{v}) \in \llbracket at \rrbracket \right\}$$

and $\widehat{\text{DataShape}}(\widehat{\text{E}})$ as follows:

$$\begin{aligned}
\alpha_{\widehat{\text{DS}}}(ds) &= \bigsqcup_{k(\underline{v}) \in ds} \text{if}(\forall i.v_i \in \text{E}, \alpha_{\widehat{\text{R}}}(\{k(\underline{v})\}), \top_{\widehat{\text{DS}}}) \\
\gamma_{\widehat{\text{DS}}}(\perp_{\widehat{\text{DS}}}) &= \emptyset \quad \gamma_{\widehat{\text{DS}}}(\top_{\widehat{\text{DS}}}) = \text{Data} \quad \gamma_{\widehat{\text{DS}}}(at\#r) = \gamma_{\widehat{\text{R}}}(at\#r)
\end{aligned}$$

Theorem 7.2. Assuming a Galois connection for the leaf elements $\wp(\text{E}) \xrightleftharpoons[\alpha_{\widehat{\text{E}}}]^{\gamma_{\widehat{\text{E}}}} \widehat{\text{E}}$ then we have a Galois connection for the data elements $\wp(\text{Data}) \xrightleftharpoons[\alpha_{\widehat{\text{DS}}}]^{\gamma_{\widehat{\text{DS}}}} \widehat{\text{DataShape}}(\widehat{\text{E}})$

Inductive Refinement Operations Recall that the data shape operations depended on the operations for inductive shape refinements, which we will now define. The lattice operations are based directly on the definitions given by Aiken and Murphy [1991].

Let $\Delta \subseteq \widehat{\text{Refinement}}(\widehat{\text{E}}) \times \widehat{\text{Refinement}}(\widehat{\text{E}})$ be a set of assumptions containing pairs of refinements $(at\#r_1, at\#r_2)$, such that $at\#r_1$ is assumed to be included in $at\#r_2$. Checking inclusion between two refinements is initially done in an empty set of assumptions, which is then possibly extended and passed down when checking recursive references.

We can then define inclusion using the following deductive rules:

$$\frac{(at\#r_1, at\#r_2) \in \Delta}{\Delta \vdash at\#r_1 \sqsubseteq_{\widehat{R}} at\#r_2} \quad \frac{\begin{array}{l} (at\#r_1, at\#r_2) \notin \Delta \\ \mathbf{refine} \ at\#r_1 = k_1(\widehat{e}_1) \mid \cdots \mid k_n(\widehat{e}_n) \\ \mathbf{refine} \ at\#r_2 = k'_1(\widehat{e}'_1) \mid \cdots \mid k'_m(\widehat{e}'_m) \\ \forall i \exists j. k_i = k'_j \wedge \Delta, (at\#r_1, at\#r_2) \vdash \widehat{e}_i \sqsubseteq_{\widehat{E}} \widehat{e}'_j \end{array}}{\Delta \vdash at\#r_1 \sqsubseteq_{\widehat{R}} at\#r_2}$$

The inclusion rules state that either we already can determine the inclusion from the given set of assumptions, or otherwise we must check that $at\#r_2$ has at least the same constructors as in $at\#r_1$ and that the arguments of the same constructors are correspondingly included as well; note, how we extended the set of assumptions on the recursive call, which is required to detect inclusion for recursive references.

Example 7.3. The refinement $\text{Nat}\#1$ is included in $\text{Nat}\#12$, given the following definitions:

$$\begin{array}{l} \mathbf{refine} \ \text{Nat}\#12 = \text{succ}(\text{Nat}\#01) \\ \mathbf{refine} \ \text{Nat}\#01 = \text{zero}() \mid \text{succ}(\text{Nat}\#0) \\ \mathbf{refine} \ \text{Nat}\#1 = \text{succ}(\text{Nat}\#0) \\ \mathbf{refine} \ \text{Nat}\#0 = \text{zero}() \end{array}$$

We check inclusion using the following steps:

1. We start with $\vdash \text{Nat}\#1 \sqsubseteq_{\widehat{R}} \text{Nat}\#12$, and since the environment is empty we proceed to use the second rule, which for the sole succ constructor requires us to further check $(\text{Nat}\#1, \text{Nat}\#12) \vdash \text{Nat}\#0 \sqsubseteq_{\widehat{D}\widehat{S}} \text{Nat}\#01$
2. By definition, checking a data shape requires checking the corresponding refinement $(\text{Nat}\#1, \text{Nat}\#12) \vdash \text{Nat}\#0 \sqsubseteq_{\widehat{R}} \text{Nat}\#01$
3. Checking $(\text{Nat}\#1, \text{Nat}\#12) \vdash \text{Nat}\#0 \sqsubseteq_{\widehat{R}} \text{Nat}\#01$ also requires using the second rule, where we can see that $\text{Nat}\#01$ contains the sole $\text{zero}()$ constructor of $\text{Nat}\#0$, and because it has no arguments we have fulfilled all required assumptions.

▲

The least upper bound and greatest lower bound operations require a memoization map $\Phi \in \widehat{\text{Refinement}}(\widehat{E}) \times \widehat{\text{Refinement}}(\widehat{E}) \rightarrow \widehat{\text{Refinement}}(\widehat{E})$ which keeps track of results for already merged refinements. We define the least upper bound using the following rules:

$$\frac{\Phi(at\#r_1, at\#r_2) = at\#r_3}{\Phi \vdash at\#r_1 \sqcup_{\widehat{R}} at\#r_2 = at\#r_3}$$

$$\frac{\begin{array}{l} (at\#r_1, at\#r_2) \notin \text{dom } \Phi \\ \mathbf{refine } at\#r_1 = k_1(\widehat{e}_1) \mid \cdots \mid k_n(\widehat{e}_n) \mid k'_{n+1}(\widehat{e}_{n+1}') \mid \cdots \mid k'_m(\widehat{e}_m') \\ \mathbf{refine } at\#r_2 = k_1(\widehat{e}_1'') \mid \cdots \mid k_n(\widehat{e}_n'') \mid k''_{n+1}(\widehat{e}_{n+1}'') \mid \cdots \mid k''_m(\widehat{e}_m'') \\ \forall i.j.k'_i \neq k''_j \quad \forall i.\Phi[(at\#r_1, at\#r_2) \mapsto at\#r_3] \vdash \widehat{e}_i \sqcup_{\widehat{E}} \widehat{e}_i'' = \widehat{e}_i''' \\ \mathbf{refine } at\#r_3 = k_1(\widehat{e}_1''') \mid \cdots \mid k_n(\widehat{e}_n''') \mid k'_{n+1}(\widehat{e}_{n+1}') \mid \cdots \mid k'_m(\widehat{e}_m') \\ \quad \mid k'_{n+1}(\widehat{e}_{n+1}') \mid \cdots \mid k''_m(\widehat{e}_m'') \end{array}}{\Phi \vdash at\#r_1 \sqcup_{\widehat{R}} at\#r_2 = at\#r_3}$$

This states that the upper bound is either the one given by the memoization, or a refinement $at\#r_3$ which contains the constructors of both $at\#r_1$ and $at\#r_2$ and where the arguments for the common constructors are merged using the least upper bound with an updated memoization. In practice, if the definition of $at\#r_3$ is not already existing in our considered set of refinements then we will create it.

Example 7.4. We will calculate the least upper bound of $\text{Nat}\#1$ and $\text{Nat}\#2$ given the following existing definitions:

$$\begin{array}{l} \mathbf{refine } \text{Nat}\#2 = \text{suc}(\text{Nat}\#1) \\ \mathbf{refine } \text{Nat}\#1 = \text{suc}(\text{Nat}\#0) \\ \mathbf{refine } \text{Nat}\#0 = \text{zero}() \end{array}$$

The calculation is done using the following steps:

- We first need to compute $\square \vdash \text{Nat}\#1 \sqcup_{\widehat{R}} \text{Nat}\#2 = \text{Nat}\#12$, which requires uses the second rule to create a refinement:

$$\mathbf{refine } \text{Nat}\#12 = \text{suc}(\text{Nat}\#01)$$

which also requires us to calculate: $[(\text{Nat}\#1 \sqcup_{\widehat{R}} \text{Nat}\#2) \mapsto \text{Nat}\#12] \vdash \text{Nat}\#0 \sqcup_{\widehat{DS}} \text{Nat}\#1 = \text{Nat}\#01$

- By definition, this means that we need to calculate $[(\text{Nat}\#1 \sqcup_{\widehat{R}} \text{Nat}\#2) \mapsto \text{Nat}\#12] \vdash \text{Nat}\#0 \sqcup_{\widehat{R}} \text{Nat}\#1 = \text{Nat}\#01$, which we do using the second rule to get refinement:

$$\mathbf{refine} \text{Nat}\#01 = \text{zero} \mid \text{suc}(\text{Nat}\#0)$$

- Since there was no common constructor, there are no more components to merge, finishing the least upper bound calculation.

▲

The greatest lower bound operator on refinements is defined similarly as:

$$\frac{\begin{array}{c} \Phi(\text{at}\#r_1, \text{at}\#r_2) = \text{at}\#r_3 \\ \Phi \vdash \text{at}\#r_1 \sqcap_{\widehat{R}} \text{at}\#r_2 = \text{at}\#r_3 \\ (\text{at}\#r_1, \text{at}\#r_2) \notin \text{dom } \Phi \\ \mathbf{refine} \text{at}\#r_1 = k_1(\widehat{e}_1) \mid \cdots \mid k_n(\widehat{e}_n) \mid k'_{n+1}(\widehat{e}_{n+1}') \mid \cdots \mid k'_m(\widehat{e}_m') \\ \mathbf{refine} \text{at}\#r_2 = k_1(\widehat{e}_1'') \mid \cdots \mid k_n(\widehat{e}_n'') \mid k''_{n+1}(\widehat{e}_{n+1}'') \mid \cdots \mid k''_m(\widehat{e}_m'') \\ \forall i, j. k'_i \neq k''_j \quad \forall i. \Phi[(\text{at}\#r_1, \text{at}\#r_2) \mapsto \text{at}\#r_3] \vdash \widehat{e}_i \sqcap_{\widehat{E}} \widehat{e}_i'' = \widehat{e}_i''' \\ \mathbf{refine} \text{at}\#r_3 = k_1(\widehat{e}_1''') \mid \cdots \mid k_n(\widehat{e}_n''') \end{array}}{\Phi \vdash \text{at}\#r_1 \sqcap_{\widehat{R}} \text{at}\#r_2 = \text{at}\#r_3}$$

The main difference from the least upper bound is that the result refinement only has common constructors and that constructor arguments are recursively merged using greatest lower bounds as required.

The widening operator is more elaborate since it needs to infer inductive rules in order to ensure that calling programs terminate. Our widening operator is an extension of the definition given in Cousot and Cousot [1995], and in particular it may need to transform the existing refinement definitions³. The procedure can be written down fully formally, but it quickly becomes incomprehensible. Instead we will describe the procedure in a step-by-step fashion:

1. Assume we want to widen the two refinements $\text{at}\#r_1$ and $\text{at}\#r_2$ to a refinement $\text{at}\#r_3$ given the following definitions (so the constructors k' and the constructors k'' are disjoint):

$$\begin{array}{l} \mathbf{refine} \text{at}\#r_1 = k_1(\widehat{e}_1) \mid \cdots \mid k_n(\widehat{e}_n) \mid k'_{n+1}(\widehat{e}_{n+1}') \mid \cdots \mid k'_m(\widehat{e}_m') \\ \mathbf{refine} \text{at}\#r_2 = k_1(\widehat{e}_1'') \mid \cdots \mid k_n(\widehat{e}_n'') \mid k''_{n+1}(\widehat{e}_{n+1}'') \mid \cdots \mid k''_m(\widehat{e}_m'') \end{array}$$

³To avoid breaking other abstract values, we in practice perform the transformation on a local copy of the rule set and then add back the local refinement definitions that are relevant for the result

2. The first merging step is to create a refinement definition for $at\#r_3$ that contains both sets of constructor definitions:

$$\begin{aligned} \mathbf{refine} \text{ } at\#r_3 = & k_1(\widehat{e}_1) \mid \cdots \mid k_n(\widehat{e}_n) \mid k'_{n+1}(\widehat{e}_{n+1}') \mid \cdots \mid k'_m(\widehat{e}_m') \\ & \mid k_1(\widehat{e}_1'') \mid \cdots \mid k_n(\widehat{e}_n'') \mid k''_{n+1}(\widehat{e}_{n+1}'') \mid \cdots \mid k''_m(\widehat{e}_m'') \end{aligned}$$

3. This may have multiple alternative definitions for the same constructor if there were any in common, so the second step is to recursively merge the arguments of the common constructors.
4. On refinements that should be merged recursively, we will this time not continue eagerly creating new definitions but instead update the memoization environment Φ , so that the refinements to-be-merged map to fresh target refinement names. For example, if we need to recursively merge refinements $at\#r_4$ and $at\#r_5$ we produce a fresh refinement $at\#r_6$ that is used as the new value (without a definition) along with environment $\Phi[(at\#r_4, at\#r_5) \mapsto at\#r_6]$.
5. If the definition of $at\#_1$ (or $at\#_2$) has more the choice between more than one constructor then we replace all references to it in the existing refinement definitions and memoization environment Φ with $at\#_3$.

6. We will then transform the definition for $at\#r_3$ to use the merged result arguments \widehat{e}_i''' , so we get:

$$\begin{aligned} \mathbf{refine} \text{ } at\#r_3 = & k_1(\widehat{e}_1''') \mid \cdots \mid k_n(\widehat{e}_n''') \mid k'_{n+1}(\widehat{e}_{n+1}''') \mid \cdots \mid k'_m(\widehat{e}_m''') \\ & \mid k''_{n+1}(\widehat{e}_{n+1}''') \mid \cdots \mid k''_m(\widehat{e}_m''') \end{aligned}$$

7. Finally, if Φ is not empty, then we repeat the same widening procedure to a new pair $(at\#r'_1, at\#r'_2) \in \text{dom } \Phi$, merging it to the target refinement $\Phi(at\#r'_1, at\#r'_2)$.

Example 7.5. We will calculate the widening of $\text{Nat}\#12$ and $\text{Nat}\#123$ to $\text{Nat}\#s$, given definitions

$$\begin{aligned} \mathbf{refine} \text{ } \text{Nat}\#12 &= \text{suc}(\text{Nat}\#01) \\ \mathbf{refine} \text{ } \text{Nat}\#123 &= \text{suc}(\text{Nat}\#012) \\ \mathbf{refine} \text{ } \text{Nat}\#01 &= \text{zero}() \mid \text{suc}(\text{Nat}\#0) \\ \mathbf{refine} \text{ } \text{Nat}\#012 &= \text{zero}() \mid \text{suc}(\text{Nat}\#01) \\ \mathbf{refine} \text{ } \text{Nat}\#0 &= \text{zero}() \end{aligned}$$

We proceed by following the steps of the widening procedure:

- The first step is to create the definition for $\text{Nat}\#s$:

refine $\text{Nat}\#s = \text{suc}(\text{Nat}\#01) \mid \text{suc}(\text{Nat}\#012)$

- Because there was a common constructor suc , we proceed by widening the constructor arguments $\text{Nat}\#01 \nabla_{\widehat{\text{DS}}} \text{Nat}\#012$.
- Since by definition the constructor arguments require recursively widening two refinements, we create a fresh refinement $\text{Nat}\#r$ as required and update the memoization map so we get $[(\text{Nat}\#01, \text{Nat}\#012) \mapsto \text{Nat}\#r]$
- We then transform the definition of $\text{Nat}\#s$ to get:

refine $\text{Nat}\#s = \text{suc}(\text{Nat}\#r)$

- Since both $\text{Nat}\#12$ and $\text{Nat}\#123$ only had a single constructor, there is no replacement with $\text{Nat}\#s$ necessary in the rest of the definitions.
- We now continue merging $\text{Nat}\#01$ and $\text{Nat}\#012$ to $\text{Nat}\#r$, by creating the definition:

refine $\text{Nat}\#r = \text{zero}() \mid \text{suc}(\text{Nat}\#0) \mid \text{suc}(\text{Nat}\#01)$

- Because there again was a common constructor suc , we need to merge recursively merge $\text{suc}(\text{Nat}\#0)$ and $\text{suc}(\text{Nat}\#01)$ to fresh refinement $\text{Nat}\#p$, getting new memoization environment $[(\text{Nat}\#0, \text{Nat}\#01) \mapsto \text{Nat}\#p]$ and transformed definition:

refine $\text{Nat}\#r = \text{zero}() \mid \text{suc}(\text{Nat}\#p)$

- Since $\text{Nat}\#01$ had more than one constructor, we replace all occurrences of $\text{Nat}\#01$ with $\text{Nat}\#r$ in the memoization environment getting $[(\text{Nat}\#0, \text{Nat}\#r) \mapsto \text{Nat}\#p]$ and similarly transform the definitions to:

refine $\text{Nat}\#s = \text{suc}(\text{Nat}\#r)$
refine $\text{Nat}\#r = \text{zero}() \mid \text{suc}(\text{Nat}\#p)$
refine $\text{Nat}\#0 = \text{zero}()$

where we dropped definitions irrelevant for the rest of the widening.

- We proceed create a definition for $\text{Nat}\#p$, merging $\text{Nat}\#0$ and $\text{Nat}\#r$ to get:

refine $\text{Nat}\#p = \text{zero}() \mid \text{suc}(\text{Nat}\#p)$

- There are no duplicate constructors, and since $\text{Nat}\#r$ had more than one possible constructor, we need to replace it in the definition with $\text{Nat}\#p$, to get the following set of refinement definitions:

refine $\text{Nat}\#s = \text{suc}(\text{Nat}\#p)$

refine $\text{Nat}\#p = \text{zero}() \mid \text{suc}(\text{Nat}\#p)$

- Here we can see that $\text{Nat}\#p$ is simply the Nat datatype, so we replace it and get the following final refinements:

refine $\text{Nat}\#s = \text{suc}(\text{Nat})$

▲

Given a Galois connection $\wp(\mathbb{E}) \xleftrightarrow[\alpha_{\mathbb{E}}]{\gamma_{\mathbb{E}}} \widehat{\mathbb{E}}$, we can define abstraction $\alpha_{\widehat{\mathbb{R}}}$ and concretization $\gamma_{\widehat{\mathbb{R}}}$ functions between

$$\text{PData}(\mathbb{E}) = \left\{ k(\underline{v}) \mid k(\underline{v}) \in \text{Data} \wedge \forall i. v_i \in \mathbb{E} \right\}$$

and refinements $\widehat{\text{Refinement}}(\widehat{\mathbb{E}})$ as follows:

$$\alpha_{\widehat{\mathbb{R}}}(ds) = \bigsqcup_{k(\underline{v}) \in ds} \text{at}\#rk \quad (\text{where refine } \text{at}\#rk = k(\underline{\alpha_{\widehat{\mathbb{E}}}(\{v\})}))$$

$$\gamma_{\widehat{\mathbb{R}}}(\text{at}\#r) = \left\{ k_i(\underline{v}_i) \mid i \in [1;n] \wedge \underline{v}_i \in \gamma_{\widehat{\mathbb{E}}}(e_i) \wedge k_i(\underline{v}_i) \in \llbracket \text{at} \rrbracket \right\}$$

(where refine $\text{at}\#r = k_1(\underline{\widehat{e}}_1) \mid \dots \mid k_n(\underline{\widehat{e}}_n)$)

Theorem 7.3. Given a Galois connection $\wp(\mathbb{E}) \xleftrightarrow[\alpha_{\mathbb{E}}]{\gamma_{\mathbb{E}}} \widehat{\mathbb{E}}$, then there is a Galois connection $\text{PData}(\mathbb{E}) \xleftrightarrow[\alpha_{\widehat{\mathbb{R}}}]{\gamma_{\widehat{\mathbb{R}}}} \widehat{\text{Refinement}}(\widehat{\mathbb{E}})$

Choice Domain

To allow a choice between multiple data-types we use the classical choice domain [Scott, 1976; Winskel, 1993a] enriched with top, defined as follows:

$$ea \in \bigoplus(\underline{A}) ::= \perp_{\oplus} \mid \top_{\oplus} \mid \text{in}_i(a_i), \quad i \in \mathbb{N}, a_i \in A_i$$

Given that injections are exclusive $\text{in}_i(a) \neq \text{in}_j(b)$ for $i \neq j$, then the ordering is:

$$\begin{aligned} \perp_{\oplus} \sqsubseteq_{\oplus} \text{in}_i(a) \sqsubseteq_{\oplus} \top_{\oplus} \\ \text{in}_i(a) \sqsubseteq_{\oplus} \text{in}_i(a') \text{ iff } a \sqsubseteq_{A_i} a' \end{aligned}$$

The least upper-bound is:

$$\begin{aligned} \perp_{\oplus} \sqcup_{\oplus} ea = ea \sqcup_{\oplus} \perp = ea \quad \top_{\oplus} \sqcup_{\oplus} ea = ea \sqcup_{\oplus} \top_{\oplus} = \top_{\oplus} \\ \text{in}_i(a) \sqcup_{\oplus} \text{in}_i(a') = \text{in}_i(a \sqcup_{\hat{A}_i} a') \quad \text{in}_i(a) \sqcup_{\oplus} \text{in}_j(b) = \top_{\oplus} \text{ iff } i \neq j \end{aligned}$$

Given a Galois connection for each option in the choice $\wp(A_i) \xleftrightarrow[\alpha_{\hat{A}_i}]{\gamma_{\hat{A}_i}} \hat{A}_i$, we can define abstraction α_{\oplus} and concretization γ_{\oplus} functions between the powerset of disjoint union of concrete domains $\wp(\bigsqcup \underline{A})$ and the choice domain $\bigoplus(\hat{\underline{A}})$ as follows:

$$\begin{aligned} \alpha_{\oplus}(eas) &= \bigsqcup_{\text{in}_i(a) \in eas} \text{in}_i(\alpha_{\hat{A}_i}(\{a\})) \\ \gamma_{\oplus}(\perp_{\oplus}) &= \emptyset \quad \gamma_{\oplus}(\text{in}_i(\hat{a})) = \{\text{in}_i(a) \mid a \in \gamma_{\oplus}(\hat{a})\} \\ \gamma_{\oplus}(\top_{\oplus}) &= \bigsqcup \underline{A} \end{aligned}$$

Theorem 7.4. Given $\forall i \left(\wp(A_i) \xleftrightarrow[\alpha_{\hat{A}_i}]{\gamma_{\hat{A}_i}} \hat{A}_i \right)$ then $\wp(\bigsqcup \underline{A}) \xleftrightarrow[\alpha_{\oplus}]{\gamma_{\oplus}} \bigoplus(\hat{\underline{A}})$

Fixedpoint Domain

For any parameterized domain $\forall E.F(E)$, we can apply the least fixed-point to get an inductive domain $\text{FX} = \text{lfp } X.F(X)$ [Scott, 1976; Smyth and Plotkin, 1982; Winskel, 1993b]. If we additionally have that for all domains E that form a lattice then $F(E)$ forms a lattice, then we get that the fixedpoint domain FX also forms a lattice.

Now let an abstract base domain form a Galois connection $\wp(F(E)) \xleftrightarrow[\alpha_F]{\gamma_F} \widehat{F}(\widehat{E})$ for any abstract base element domain \widehat{E} with Galois connection $\wp(E) \xleftrightarrow[\alpha_E]{\gamma_E} \widehat{E}$, then we straightforwardly have the following abstraction α_{Fix} and concretization $\gamma_{\widehat{\text{Fix}}}$ functions between $\wp(\text{lfp } X.F(X))$ and $\wp(\text{lfp } \widehat{X}.\widehat{F}(\widehat{X}))$:

$$\alpha_{\text{Fix}}(fs) = \alpha_F(fs) \quad \gamma_{\widehat{\text{Fix}}}(\widehat{f}) = \gamma_{\widehat{F}}(\widehat{f})$$

Theorem 7.5. *If that for continuous parametrized domains F, \widehat{F} it holds that for all E, \widehat{E} such that $\wp(E) \xleftrightarrow[\alpha_E]{\gamma_E} \widehat{E}$ then $\wp(F(E)) \xleftrightarrow[\alpha_{\widehat{F}}]{\gamma_{\widehat{F}}} \widehat{F}(\widehat{E})$, then $\wp(\text{lfp } X.F(X)) \xleftrightarrow[\alpha_{\widehat{\text{Fix}}}]^{\gamma_{\widehat{\text{Fix}}}} \text{lfp } \widehat{X}.\widehat{F}(\widehat{X})$*

Value Domains

We presented the required components for abstracting abstractions, and now all that is left is putting everything together using the fixedpoint domain.

We first define a parameterized value shape domain that combine all our different shape domains together:

$$\widehat{\text{PValueShape}}(\widehat{E}) = \widehat{\text{Interval}} \oplus \widehat{\text{SetShape}}(\widehat{E}) \oplus \widehat{\text{DataShape}}(\widehat{E})$$

We define a corresponding parameterized concrete domain:

$$\text{PValue}(E) = \mathbb{Z} \uplus \wp(E) \uplus \text{Data}$$

Corollary 7.1. *If given a Galois connection for the parameter element domain $\wp(E) \xleftrightarrow[\alpha_E]{\gamma_E} \widehat{E}$, then we have a Galois connection for the parameterized value domain $\wp(\text{PValue}(E)) \xleftrightarrow[\alpha_{\widehat{\text{PVS}}}]^{\gamma_{\widehat{\text{PVS}}}} \widehat{\text{PValueShape}}(\widehat{E})$*

Finally, we define our complete value shape abstract domain using the least fixedpoint domain and the parameterized value shape domain: $\widehat{\text{ValueShape}} = \text{lfp } \widehat{X}.\widehat{\text{PValueShape}}(\widehat{X})$. Similarly, we have the equivalence $\text{Value} = \text{lfp } X.\text{PValue}(X)$.

Corollary 7.2. *We have a Galois connection $\wp(\text{Value}) \xleftrightarrow[\alpha_{\widehat{\text{VS}}}]^{\gamma_{\widehat{\text{VS}}}} \widehat{\text{ValueShape}}$*

Abstract State Domains

Given our shape and relational domains, we can construct abstractions for our stores and results (control operators) which represent the state that Rascal programs manipulate.

Abstract Store Domain Our abstract store maps variables to a pair consisting of a Boolean indicating whether the variable is assigned and the target abstract value:

$$\widehat{\sigma} \in \widehat{\text{Store}} = \text{Var} \rightarrow \{\text{ff}, \text{tt}\} \times \widehat{\text{ValueShape}}$$

We lift our orderings and lattice operations from the value shape domain to abstract stores (resp. $\sqsubseteq_{\widehat{\text{Store}}}$, $\sqcup_{\widehat{\text{Store}}}$, $\sqcap_{\widehat{\text{Store}}}$, $\nabla_{\widehat{\text{Store}}}$). We define abstraction $\alpha_{\widehat{\text{Store}}}$ and concretization $\gamma_{\widehat{\text{Store}}}$ operations between concrete sets of stores $\wp(\text{Store})$ and abstract stores $\widehat{\text{Store}}$:

$$\alpha_{\widehat{\text{Store}}}(\Sigma) = \bigsqcup_{\sigma \in \Sigma} \lambda x. \text{if}(x \in \mathbf{dom} \sigma, (\text{ff}, \alpha_{\widehat{\text{V}}}(\{\sigma(x)\})), (\text{tt}, \perp_{\widehat{\text{V}}}))$$

$$\gamma_{\widehat{\text{Store}}}(\widehat{\sigma}) = \{\sigma \mid \forall x. \widehat{\sigma}(x) = (b, \widehat{v}_s) \wedge \text{if}(x \in \mathbf{dom} \sigma, \sigma(x) \in \gamma_{\widehat{\text{V}}}(\widehat{v}_s), b)\}$$

Theorem 7.6. *We have a Galois connection $\wp(\text{Store}) \xleftrightarrow[\alpha_{\widehat{\text{Store}}}]{\gamma_{\widehat{\text{Store}}}} \widehat{\text{Store}}$*

Abstract Result Domain A particularly interesting side-effect of Schmidt-style abstract interpretation is that it enables to handle abstraction of control flow more directly. Traditionally, control flow is often handled using a denotational collecting semantics or by explicitly constructing a control flow graph, but both these approaches are non-trivial to apply for a rich language like Rascal.

Our idea is to model different kind of results—success, failure, errors—that affect control flow directly in a $\widehat{\text{ResultSet}}$ domain which keeps track of possible results for a particular execution case, in addition to managing a separate store for each possible result kind. Keeping separate stores is cheap, since there are only three kinds of results, and important to maintain precision, since different kind of results lead to different paths in concrete programs.

$$\begin{aligned} \widehat{vres} \in \widehat{\text{ValueResult}} &::= \mathbf{success} \widehat{v} \mid \widehat{exres} \\ \widehat{exres} \in \widehat{\text{ExceptionalResult}} &::= \mathbf{fail} \mid \mathbf{error} \\ \widehat{res} \in \widehat{\text{Result}} &::= \cdot \mid \widehat{vres}; \widehat{\sigma} \quad \widehat{Res} \in \widehat{\text{ResultSet}} ::= \widehat{res} \end{aligned}$$

The ordering of result sets checks is dependent on the presence of result kinds and the ordering of contained values and stores. Formally, we can define it as follows:

$$\begin{aligned}
& \widehat{Res}_1 \sqsubseteq_{\widehat{RS}} \widehat{Res}_2 \\
& \text{iff} \\
& (\mathbf{success} \ v_1; \sigma_1) \in \widehat{Res}_1 \Rightarrow \exists v_2, \sigma_2. v_1 \sqsubseteq_{\widehat{VS}} v_2 \wedge \sigma_1 \sqsubseteq_{\widehat{Store}} \sigma_2 \wedge \\
& \qquad \qquad \qquad (\mathbf{success} \ v_2; \sigma_2) \in \widehat{Res}_2 \\
& \qquad \qquad \qquad \wedge \\
& (\widehat{exres}; \widehat{\sigma}_1) \in \widehat{Res}_1 \Rightarrow \exists \widehat{\sigma}_2. \widehat{\sigma}_1 \sqsubseteq_{\widehat{Store}} \widehat{\sigma}_2 \wedge (\widehat{exres}; \widehat{\sigma}_2) \in \widehat{Res}_2
\end{aligned}$$

The least upper bound collects the different kinds of results available in both input result sets, merging the components—result value, store—of the common ones. We formally define the least upper bound as follows (widening analogous):

$$\begin{aligned}
& \widehat{Res}_1 \sqcup_{\widehat{RS}} \widehat{Res}_2 = \text{merge}(\widehat{Res}_1, \widehat{Res}_2) \\
& \text{merge}(\widehat{Res}) = \begin{cases} (\mathbf{success} \ (\widehat{v}_1 \sqcup_{\widehat{VS}} \widehat{v}_2); \widehat{\sigma}_1 \sqcup_{\widehat{Store}} \widehat{\sigma}_2), \text{merge}(\widehat{Res}') & \text{if } \widehat{Res} = (\mathbf{success} \ \widehat{v}_1; \widehat{\sigma}_1), (\mathbf{success} \ \widehat{v}_2; \widehat{\sigma}_2), \widehat{Res}' \\ (\mathbf{fail}; \widehat{\sigma}_1 \sqcup_{\widehat{Store}} \widehat{\sigma}_2), \text{merge}(\widehat{Res}') & \text{if } \widehat{Res} = (\mathbf{fail}; \widehat{\sigma}_1), (\mathbf{fail}; \widehat{\sigma}_2), \widehat{Res}' \\ (\mathbf{error}; \widehat{\sigma}_1 \sqcup_{\widehat{Store}} \widehat{\sigma}_2), \text{merge}(\widehat{Res}') & \text{if } \widehat{Res} = (\mathbf{error}; \widehat{\sigma}_1), (\mathbf{error}; \widehat{\sigma}_2), \widehat{Res}' \\ \widehat{Res} & \text{otherwise} \end{cases}
\end{aligned}$$

The greatest lower bound contains the result kinds that were present in both input results sets, and then calculates the greatest lower bound of their components. Formally, we define the greatest lower bound as follows:

$$\begin{aligned}
& \widehat{Res}_1 \sqcap_{\widehat{RS}} \widehat{Res}_2 = \bigsqcup_{\substack{\widehat{res}_1 \in \widehat{Res}_1 \wedge \\ \widehat{res}_2 \in \widehat{Res}_2}} (\widehat{res}_1 \sqcap_{\widehat{RS}} \widehat{res}_2) \\
& \widehat{res}_1 \sqcap_{\widehat{RS}} \widehat{res}_2 = \begin{cases} \mathbf{success} \ (\widehat{v}_1 \sqcap_{\widehat{VS}} \widehat{v}_2); \widehat{\sigma}_1 \sqcap_{\widehat{Store}} \widehat{\sigma}_2 & \text{if } \widehat{res}_i = \mathbf{success} \ \widehat{v}_i; \widehat{\sigma}_i \\ \mathbf{fail}; \widehat{\sigma}_1 \sqcap_{\widehat{Store}} \widehat{\sigma}_2 & \text{if } \widehat{res}_i = \mathbf{fail}; \widehat{\sigma}_i \\ \mathbf{error}; \widehat{\sigma}_1 \sqcap_{\widehat{Store}} \widehat{\sigma}_2 & \text{if } \widehat{res}_i = \mathbf{error}; \widehat{\sigma}_i \\ \varepsilon & \text{otherwise} \end{cases}
\end{aligned}$$

We define abstraction $\alpha_{\widehat{RS}}$ and concretization $\gamma_{\widehat{RS}}$ operations between the concrete domain $\wp(\text{ValueResult} \times \text{Store})$ and abstract result sets $\widehat{\text{ResultSet}}$ as follows:

$$\begin{aligned} \alpha_{\widehat{RS}}(\text{Res}) &= \bigsqcup_{(vres, \sigma) \in \text{Res}} \alpha_{\widehat{VR}}(\{vres\}); \alpha_{\widehat{\text{Store}}}(\{\sigma\}) \\ \alpha_{\widehat{VR}}(Vres) &= \bigsqcup_{\text{value } v \in Vres} (\text{value } \alpha_{\widehat{VS}}(v)) \sqcup \bigsqcup_{\widehat{exres} \in Vres} \widehat{exres} \\ \gamma_{\widehat{RS}}(\widehat{vres}; \widehat{\sigma}) &= \bigcup_i \{(vres, \sigma) \mid vres \in \gamma_{\widehat{VR}}(\widehat{vres}_i) \wedge \sigma \in \gamma_{\widehat{\text{St}}}(\widehat{\sigma}_i)\} \\ \gamma_{\widehat{VR}}(\text{value } \widehat{vs}) &= \{\text{value } v \mid v \in \gamma_{\widehat{VS}}(\widehat{vs})\} \quad \gamma_{\widehat{VR}}(\widehat{exres}) = \{\widehat{exres}\} \end{aligned}$$

Theorem 7.7. *We have a Galois connection*

$$\wp(\text{ValueResult} \times \text{Store}) \xleftrightarrow[\alpha_{\widehat{RS}}]{\gamma_{\widehat{RS}}} \widehat{\text{ResultSet}}$$

7.4 Abstract Semantics

The genius of Schmidt-style abstract interpretation is that it allows us to get abstract operational rules in a way that closely corresponds to our existing concrete operational rules. The operational structure of the rules can be derived almost mechanically [Schmidt, 1998; Bodin et al., 2015], which we exemplify by translation of the operational rules for the Rascal UltraLight subset. The creative work is therefore providing abstract definitions for conditions and semantic operations such as pattern matching, and defining *trace memoization strategies* for non-structurally recursive operational rules in order to make it possible to finitely approximate an infinite number of concrete traces and produce a terminating static analysis.

Abstracting Operational Rules

In Schmidt-style abstract interpretation, we have to for each concrete judgment create an over-approximative abstract judgment that captures the result for a set of possible traces. Recall, that the concrete operational rules for expressions have a judgment of form $e; \sigma \xrightarrow[\text{expr}]{\text{}} vres; \sigma'$ which states that evaluating an expression e in store σ produces the value result $vres$ and store σ' . The corresponding abstract semantics judgment

has form $e; \widehat{\sigma} \xrightarrow[\text{a-expr}]{} \widehat{Res}$ which evaluates expression e in store σ producing a result set $\widehat{Res} = (\widehat{vres}_1; \widehat{\sigma}_1), \dots, (\widehat{vres}_n; \widehat{\sigma}_n)$ consisting of pairs of possible abstract value results and corresponding abstract stores. These sets of pairs are meant to finitely over-approximate the possibly infinite sets of concrete results from evaluating e by providing abstractions of possible values and stores for each different kind of result (success, failure, errors).

Collecting premise notation In concrete semantics we usually only require that premises are satisfied by a single rule, but in abstract semantics we have to consider all possible applicable rules to preserve soundness. Thus, to reason about the generalized sets of results we introduce the notation

$$\{i \Rightarrow O\} \triangleq O = \bigsqcup \{o \mid i \Rightarrow o\}$$

to collect all derivations with input i into a single output O that is equal to the least upper bound of the output of each individual derivation o .

Translating semantics To illustrate the steps required in Schmidt-style translation of operational rules, we will use the operational rules for variable accesses. Recall that the concrete semantics contains two rules for variable accesses, E-V-S for successful lookup, and E-V-ER for producing errors when accessing unassigned variables.

$$\text{E-V-S} \frac{x \in \text{dom } \sigma}{x; \sigma \xrightarrow[\text{expr}]{} \mathbf{value} \ \sigma(x); \sigma} \quad \text{E-V-ER} \frac{x \notin \text{dom } \sigma}{x; \sigma \xrightarrow[\text{expr}]{} \mathbf{error}; \sigma}$$

To translate the concrete rules for a syntactic element to abstract operational rules, we follow three steps in general:

1. For each concrete rule, create an abstract rule that uses an judgment for evaluation of a concrete syntactic form, e.g., variables.
2. Replace the concrete conditions and semantic operations with the equivalent abstract conditions and semantic operations for target abstract values.
3. Create a general abstract rule that collects all possible evaluations of the syntax-specific judgment rules, over-approximating all possible result values.

The translated abstract operational rules for evaluating variable accesses, are presented below:

$$\begin{array}{c}
 \text{AE-V-S} \frac{\widehat{\sigma}(x) = (b, \widehat{vs})}{x; \widehat{\sigma} \xrightarrow{\text{a-expr-v}} \mathbf{value} \widehat{vs}; \widehat{\sigma}} \quad \text{AE-V-ER} \frac{\widehat{\sigma}(x) = (\mathbf{tt}, \widehat{vs})}{x; \widehat{\sigma} \xrightarrow{\text{a-expr-v}} \mathbf{error}; \widehat{\sigma}} \\
 \text{AE-V} \frac{\{x; \widehat{\sigma} \xrightarrow{\text{a-expr-v}} \widehat{Res}'\}}{x; \widehat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res}'}
 \end{array}$$

In the first step, we create two abstract syntax-specific rules AE-V-S and AE-V-ER corresponding to E-V-S and E-V-ER respectively. In the second step, we must translate the conditions and semantic operations which for our example requires:

- Translating the test $x \in \text{dom } \sigma$ with a corresponding condition on $\widehat{\sigma}$. Recall that in our abstract store, we have assigned a pair (b, \widehat{vs}) for each variable x , where the first component b is a Boolean that indicates whether the variable x is potentially unassigned; thus, we look at this component to decide whether we have $x \notin \text{dom } \sigma$ (if $b = \mathbf{tt}$) and otherwise we may only have $x \in \text{dom } \sigma$.
- Translating the lookup $\sigma(x)$ in the result of E-V-S to a lookup of values in $\widehat{\sigma}$, which requires extracting the second component assigned to variable x .

Finally, in the last step we use abstract rule AE-V to collect the two syntax-specific abstract operational rules AE-V-S and AE-V-ER.

The possible shapes of the result value depend on the pair assigned to x in the abstract store; if the resulting value is \perp , we propagate the emptiness dropping the complete result value. We can use the following examples, to illustrate the possible outcome result shapes:

Store	Value Results	Rules
$\widehat{\sigma}[x \mapsto (\mathbf{ff}, \perp_{\widehat{vs}})]$	\emptyset	AE-V-S
$\widehat{\sigma}[x \mapsto (\mathbf{ff}, [1;3])]$	$\mathbf{value} [1;3]$	AE-V-S
$\widehat{\sigma}[x \mapsto (\mathbf{tt}, [1;3])]$	$\{\mathbf{value} [1;3], \mathbf{error}\}$	AE-V-S, AE-V-ER
$\widehat{\sigma}[x \mapsto (\mathbf{tt}, \perp_{\widehat{vs}})]$	\mathbf{error}	AE-V-S, AE-V-ER

The abstract operational semantics rules for the rest of the basic expressions (excluding visitors) are presented in Figure 7.2, Figure 7.3 and

Figure 7.4. The rules are generally translated from the corresponding concrete semantics rules in Chapter 6, using the steps we presented; note that sequencing expressions are a special case of block expressions which have no variable declarations and only two expressions in the body. There are a few things to note regarding the abstract conditions and semantic operations:

- When the concrete semantics checks that a result of evaluating an expression has a particular shape, the abstract semantics checks whether the particular shape is possibly included in the result set of evaluating the same expression. E.g., to check that evaluation of e has succeeded the abstract semantics uses $e; \hat{\sigma} \xrightarrow[\text{a-expr}]{} \widehat{Res}$ and $(\mathbf{success} \ \hat{vs}; \hat{\sigma}') \in \widehat{Res}$, as compared to $e; \sigma \xrightarrow[\text{expr}]{} \mathbf{success} \ v; \sigma'$ in the concrete semantics.
- Typing is now done using an abstract judgment $\hat{vs} \hat{:} t$ which works analogously to the concrete one except t is over-approximates the concrete type of values, because the given value shape \hat{vs} is itself over-approximated. This means that the corresponding abstract subtyping $t \hat{<}: t'$ relation must take into account the over-approximation of t and thus that it may be possible that both $t \hat{<}: t'$ and $t \not\hat{<}: t'$ is true.
- To check whether a particular constructor is possible, we use the auxiliary function $\widehat{\text{unfold}}(\hat{vs}, t)$ which produces a refined value of type t if possible—splitting alternative constructors in case of refinements—and otherwise produces **error** if the value is possibly not an element of t .
- We generalize the $\widehat{\text{ResultSet}}$ domain to work with sequences of values on successful evaluation for representing results of expression sequences e . This is straightforward since a sequence of expressions must produce a sequence of values of the same size, and a fixed product of lattice domains forms a lattice (essentially the iterated Cartesian product).

Expressions (General)	
$\text{AE-VAL} \frac{n; \hat{\sigma} \xrightarrow{\text{a-expr}} \text{success } n; \hat{\sigma}}{n; \hat{\sigma} \xrightarrow{\text{a-expr}} \text{success } n; \hat{\sigma}}$	$\text{AE-BIN} \frac{\{\{e_1 \otimes e_2; \hat{\sigma} \xrightarrow{\text{a-expr-bin}} \widehat{Res}\}\}}{e_1 \otimes e_2; \hat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res}}$
$\text{AE-ASGN} \frac{\{\{x = e; \hat{\sigma} \xrightarrow{\text{a-expr-asgn}} \widehat{Res}\}\}}{x = e; \hat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res}}$	$\text{AE-SEQ} \frac{\{\{x = e; \hat{\sigma} \xrightarrow{\text{a-expr-seq}} \widehat{Res}\}\}}{x = e; \hat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res}}$
$\text{AE-CONS} \frac{\{\{k(\underline{e}); \hat{\sigma} \xrightarrow{\text{a-expr-cons}} \widehat{Res}\}\}}{k(\underline{e}); \hat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res}}$	$\text{AE-SET} \frac{\{\{\underline{e}\}; \hat{\sigma} \xrightarrow{\text{a-expr-set}} \widehat{Res}\}}{\{\underline{e}\}; \hat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res}}$
$\text{AE-FAIL} \frac{\text{fail}; \hat{\sigma} \xrightarrow{\text{a-expr}} \text{fail}; \hat{\sigma}}{\text{fail}; \hat{\sigma} \xrightarrow{\text{a-expr}} \text{fail}; \hat{\sigma}}$	$\text{AE-IF} \frac{\{\{\text{if } e_{\text{cond}} \text{ then } e_1 \text{ else } e_2; \hat{\sigma} \xrightarrow{\text{a-expr-if}} \widehat{Res}\}\}}{\text{if } e_{\text{cond}} \text{ then } e_1 \text{ else } e_2; \hat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res}}$
Expression Sequences	
$\text{AES} \frac{\{\{\underline{e}; \hat{\sigma} \xrightarrow{\text{a-expr}^{\star-1}} \widehat{Res}\}\}}{e; \hat{\sigma} \xrightarrow{\text{a-expr}^{\star}} \widehat{Res}}$	$\text{AES-EMP} \frac{\varepsilon; \hat{\sigma} \xrightarrow{\text{a-expr}^{\star-1}} \text{success } \varepsilon; \hat{\sigma}}{e; \hat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res} \quad (\text{success } \widehat{vs}; \hat{\sigma}') \in \widehat{Res}}$
$\text{AES-MORE} \frac{\underline{e}'; \hat{\sigma}'' \xrightarrow{\text{a-expr}^{\star}} \widehat{Res}^{\star'}}{e, \underline{e}'; \hat{\sigma} \xrightarrow{\text{a-expr}^{\star-1}} \text{success } \widehat{vs}, \widehat{vs}'; \hat{\sigma}' \quad (\text{success } \widehat{vs}'; \hat{\sigma}') \in \widehat{Res}^{\star'}}$	$\text{AES-Exc1} \frac{e, \underline{e}'; \hat{\sigma} \xrightarrow{\text{a-expr}^{\star-1}} \text{success } \widehat{vs}, \widehat{vs}'; \hat{\sigma}' \quad (\widehat{exres}; \hat{\sigma}') \in \widehat{Res}}{e; \hat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res} \quad (\widehat{exres}; \hat{\sigma}') \in \widehat{Res}}$
$\text{AES-Exc2} \frac{\underline{e}'; \hat{\sigma}'' \xrightarrow{\text{a-expr}^{\star}} \widehat{Res}^{\star'} \quad (\widehat{exres}; \hat{\sigma}') \in \widehat{Res}^{\star'}}{e, \underline{e}'; \hat{\sigma} \xrightarrow{\text{a-expr}^{\star-1}} \widehat{exres}; \hat{\sigma}'}$	

Figure 7.2: Abstract Operational Semantics Rules for Basic Expressions

Binary Expression

$$\begin{array}{c}
\frac{
\begin{array}{l}
e_1; \hat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res} \quad (\mathbf{success} \ \widehat{v}_1; \hat{\sigma}') \in \widehat{Res} \\
e_2; \hat{\sigma}'' \xrightarrow{\text{a-expr}} \widehat{Res}' \quad (\mathbf{success} \ \widehat{v}_2; \hat{\sigma}') \in \widehat{Res}'' \\
\widehat{vs}_1 \otimes \widehat{vs}_2 = \widehat{VRes} \quad \widehat{Res}' = \sqcup \left\{ \widehat{vres}; \hat{\sigma}' \mid \widehat{vres} \in \widehat{VRes} \right\}
\end{array}
}{
\text{AE-BIN-SUCS} \quad e_1 \otimes e_2; \hat{\sigma} \xrightarrow{\text{a-expr-bin}} \widehat{Res}'
} \\
\frac{
\text{AE-BIN-EXC1} \quad \frac{
e_1; \hat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res} \quad (\widehat{exres}; \hat{\sigma}') \in \widehat{Res}
}{
e_1 \otimes e_2; \hat{\sigma} \xrightarrow{\text{a-expr-bin}} \widehat{exres}; \hat{\sigma}'
}
}{
\text{AE-BIN-SUCS} \quad \frac{
e_1; \hat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res} \quad (\mathbf{success} \ \widehat{v}_1; \hat{\sigma}') \in \widehat{Res} \\
e_2; \hat{\sigma}'' \xrightarrow{\text{a-expr}} \widehat{Res}' \quad (\widehat{exres}; \hat{\sigma}') \in \widehat{Res}'
}{
e_1 \otimes e_2; \hat{\sigma} \xrightarrow{\text{a-expr-bin}} \widehat{exres}; \hat{\sigma}'
}
}
\end{array}$$

Assignment Expression

$$\begin{array}{c}
\text{AE-ASGN-SUCS} \quad \frac{
\mathbf{local} \ t \ x \vee \mathbf{global} \ t \ x \quad e; \hat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res} \\
(\mathbf{success} \ \widehat{vs}; \hat{\sigma}') \in \widehat{Res} \quad \widehat{vs} \hat{=} t' \quad t' \hat{<} t
}{
x = e; \hat{\sigma} \xrightarrow{\text{a-expr-asgn}} \mathbf{success} \ \widehat{vs}; \hat{\sigma}' [x \mapsto (\mathbf{ff}, \widehat{vs})]
} \\
\text{AE-ASGN-ERR} \quad \frac{
\mathbf{local} \ t \ x \vee \mathbf{global} \ t \ x \quad e; \hat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res} \\
(\mathbf{success} \ \widehat{vs}; \hat{\sigma}') \in \widehat{Res} \quad \widehat{vs} \hat{=} t' \quad t' \not\hat{<} t
}{
x = e; \hat{\sigma} \xrightarrow{\text{a-expr-asgn}} \mathbf{error}; \hat{\sigma}'
} \\
\text{AE-ASGN-EXC} \quad \frac{
e; \hat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res} \quad (\widehat{exres}; \hat{\sigma}') \in \widehat{Res}
}{
x = e; \hat{\sigma} \xrightarrow{\text{a-expr-asgn}} \widehat{exres}; \hat{\sigma}'
}
\end{array}$$

Sequencing Expression

$$\begin{array}{c}
\text{AE-SEQ-SUCS} \quad \frac{
e_1, e_2; \hat{\sigma} \xrightarrow{\text{a-expr}^*} \widehat{Res}^* \quad (\mathbf{success} \ \widehat{vs}_1, \widehat{vs}_2; \hat{\sigma}') \in \widehat{Res}^*
}{
e_1; e_2; \hat{\sigma} \xrightarrow{\text{a-expr-seq}} \mathbf{success} \ \widehat{vs}_2; \hat{\sigma}'
} \\
\text{AE-SEQ-EXC} \quad \frac{
e_1, e_2; \hat{\sigma} \xrightarrow{\text{a-expr}^*} \widehat{Res}^* \quad (\widehat{exres}; \hat{\sigma}') \in \widehat{Res}^*
}{
e_1; e_2; \hat{\sigma} \xrightarrow{\text{a-expr-seq}} \widehat{exres}; \hat{\sigma}'
}
\end{array}$$

Figure 7.3: Abstract Operational Semantics Rules for Basic Expressions (Cont.)

Constructor Expression

$$\begin{array}{c}
\text{AE-CONS-SUCS} \frac{\text{data } at = \dots |k(\underline{t})| \dots \quad \underline{e}; \hat{\sigma} \xrightarrow{\text{a-expr}\star} \widehat{Res}\star \\
(\text{success } \underline{vs}; \hat{\sigma}') \in \widehat{Res}\star \quad \underline{vs} \hat{=} t' \quad t' \hat{<} t}{k(\underline{e}); \hat{\sigma} \xrightarrow{\text{a-expr}} \text{success } k(\underline{vs}); \hat{\sigma}'} \\
\text{AE-CONS-ERR} \frac{\text{data } at = \dots |k(\underline{t})| \dots \quad \underline{e}; \hat{\sigma} \xrightarrow{\text{a-expr}\star} \widehat{Res}\star \\
(\text{success } \underline{vs}; \hat{\sigma}') \in \widehat{Res}\star \quad \underline{vs} \hat{=} t' \quad \exists i.t'_i \hat{<} t_i}{k(\underline{e}); \hat{\sigma} \xrightarrow{\text{a-expr}} \text{error}; \hat{\sigma}'} \\
\text{AE-CONS-EXC} \frac{\underline{e}; \hat{\sigma} \xrightarrow{\text{a-expr}\star} \widehat{Res}\star \quad (\underline{exres}; \hat{\sigma}') \in \widehat{Res}\star}{k(\underline{e}); \hat{\sigma} \xrightarrow{\text{a-expr}} \underline{exres}; \hat{\sigma}'}
\end{array}$$

Set Expression

$$\begin{array}{c}
\text{AE-SET-SUCS} \frac{\underline{e}; \hat{\sigma} \xrightarrow{\text{a-expr}\star} \widehat{Res}\star \quad (\text{success } \underline{vs}; \hat{\sigma}') \in \widehat{Res}\star}{\{\underline{e}\}; \hat{\sigma} \xrightarrow{\text{a-expr-set}} \text{success } \{\underline{vs}\}; \hat{\sigma}'} \\
\text{AE-SET-EXC} \frac{\underline{e}; \hat{\sigma} \xrightarrow{\text{a-expr}\star} \widehat{Res}\star \quad (\underline{exres}; \hat{\sigma}') \in \widehat{Res}\star}{\{\underline{e}\}; \hat{\sigma} \xrightarrow{\text{a-expr-set}} \underline{exres}; \hat{\sigma}'}
\end{array}$$

If Expression

$$\begin{array}{c}
\text{AE-IF-T} \frac{e_{\text{cond}}; \hat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res} \quad (\text{success } \underline{vs}; \hat{\sigma}'') \in \widehat{Res} \\
\text{success true}() \in \widehat{\text{unfold}}(\underline{vs}, \text{Bool}) \quad e_1; \hat{\sigma}'' \xrightarrow{\text{a-expr}} \widehat{Res}'}{\text{if } e_{\text{cond}} \text{ then } e_1 \text{ else } e_2; \hat{\sigma} \xrightarrow{\text{a-expr-if}} \widehat{Res}'} \\
\text{AE-IF-F} \frac{e_{\text{cond}}; \hat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res} \quad (\text{success } \underline{vs}; \hat{\sigma}'') \in \widehat{Res} \\
\text{success false}() \in \widehat{\text{unfold}}(\underline{vs}, \text{Bool}) \quad e_2; \hat{\sigma}'' \xrightarrow{\text{a-expr}} \widehat{Res}'}{\text{if } e_{\text{cond}} \text{ then } e_1 \text{ else } e_2; \hat{\sigma} \xrightarrow{\text{a-expr-if}} \widehat{Res}'} \\
\text{AE-IF-ER} \frac{e_{\text{cond}}; \hat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res} \quad (\text{success } \underline{vs}; \hat{\sigma}') \in \widehat{Res} \\
\text{error} \in \widehat{\text{unfold}}(\underline{vs}, \text{Bool})}{\text{if } e_{\text{cond}} \text{ then } e_1 \text{ else } e_2; \hat{\sigma} \xrightarrow{\text{a-expr-if}} \text{error}; \hat{\sigma}'} \\
\text{AE-IF-EX} \frac{e_{\text{cond}}; \hat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res} \quad (\underline{exres}; \hat{\sigma}') \in \widehat{Res}}{\text{if } e_{\text{cond}} \text{ then } e_1 \text{ else } e_2; \hat{\sigma} \xrightarrow{\text{a-expr-if}} \underline{exres}; \hat{\sigma}'}
\end{array}$$

Figure 7.4: Abstract Operational Semantics Rules for Basic Expressions (Cont. 2)

Abstract Semantic Operators

We have already seen how the rules for basic expressions have been systematically translated, including the required checks. As discussed the same is also required for semantic operations, where we in particular will discuss operators (including unary and binary), and pattern matching.

Semantic Operators In general for each semantic operator $\llbracket \oplus \rrbracket$ corresponding to a language operator \oplus , we need to have a corresponding abstract semantic operator $\llbracket \hat{\oplus} \rrbracket$ which works with the abstracted values of the target type, e.g. for $\llbracket + \rrbracket \in \mathbb{Z} \rightarrow \mathbb{Z} \rightarrow \mathbb{Z}$ we need $\llbracket \hat{+} \rrbracket \in \widehat{\text{Interval}} \rightarrow \widehat{\text{Interval}} \rightarrow \widehat{\text{Interval}}$ and for $\llbracket \cup \rrbracket \in \wp(\text{Value}) \rightarrow \wp(\text{Value}) \rightarrow \wp(\text{Value})$ we need $\llbracket \hat{\cup} \rrbracket \in \widehat{\text{SetShape}}(\widehat{\text{ValueShape}}) \rightarrow \widehat{\text{SetShape}}(\widehat{\text{ValueShape}}) \rightarrow \widehat{\text{SetShape}}(\widehat{\text{ValueShape}})$. Due to the over-approximative semantics, operands might get evaluated to values that are less precise than values in the supported range for the target abstract semantics operator, and in that case we need to consider both cases of where the abstract value lies in target range and where it lies outside of target range. We can do this by introducing rules that check these potential need for coercion when evaluating the syntactic operators, e.g. the following rules for binary operators:

$$\begin{array}{c}
 \text{BOR} \frac{\{\widehat{v}s_1 \oplus \widehat{v}s_2 =_1 \widehat{VRes}\}}{\widehat{v}s_1 \oplus \widehat{v}s_2 = \widehat{VRes}} \\
 \\
 \text{BOP1-1} \frac{\text{value } \widehat{v}s_1' \in \widehat{\text{unfold}}(\widehat{v}s_1, t_1) \quad \text{value } \widehat{v}s_2' \in \widehat{\text{unfold}}(\widehat{v}s_2, t_2) \quad \llbracket \hat{\oplus} \rrbracket \in \llbracket \widehat{t}_1 \rrbracket \times \llbracket \widehat{t}_2 \rrbracket \rightarrow \llbracket \widehat{t}' \rrbracket}}{\widehat{v}s_1 \oplus \widehat{v}s_2 =_1 \text{value } \llbracket \hat{\oplus} \rrbracket(\widehat{v}s_1', \widehat{v}s_2')} \\
 \\
 \text{BOP1-2} \frac{\llbracket \hat{\oplus} \rrbracket \in \llbracket \widehat{t}_1 \rrbracket \times \llbracket \widehat{t}_2 \rrbracket \rightarrow \llbracket \widehat{t}' \rrbracket \quad \text{error} \in \widehat{\text{unfold}}(\widehat{v}s_1, t_1) \vee \text{error} \in \widehat{\text{unfold}}(\widehat{v}s_2, t_2)}{\widehat{v}s_1 \oplus \widehat{v}s_2 =_1 \text{error}}
 \end{array}$$

Pattern matching The concrete pattern matching judgments have forms $\sigma \vdash p \stackrel{?}{:=} v \xrightarrow[\text{match}]{} \underline{\rho}$ for a concrete pattern p matched against value v under store σ —used for matching target value against assigned variable values—producing a sequence of binding environments (each $\rho \in \text{BindingEnv} = \text{Var} \rightarrow \text{Value}$), and $\sigma \vdash \star p \stackrel{?}{:=} v \mid \mathbb{V} \xrightarrow[\text{match}\star]{} \underline{\rho}$ for a

sequence pattern $\star p$ with similar construction but also having a set of already matched values \mathbb{V} to track already the set of already matched partitions of values for non-deterministic matches of collection elements.

The corresponding abstract matching judgment for ordinary patterns has form $\hat{\sigma} \vdash p \stackrel{?}{:=} \hat{v}s \xrightarrow[\text{a-match}]{} \hat{Q}_R$ where we pattern match on value shapes $\hat{v}s$ and produce an output \hat{Q}_R representing the set of possible match results. Each match result triple $(\hat{\sigma}', \hat{v}s', \hat{Q}) \in \hat{Q}_R$, has an abstract store $\hat{\sigma}'$ and abstract value shape $\hat{v}s'$ that soundly refine the input store $\hat{\sigma}$ and value $\hat{v}s$ —so $\hat{\sigma}' \sqsubseteq \hat{\sigma}$ and $\hat{v}s' \sqsubseteq \hat{v}s$ —assuming that matching the pattern p produced the abstract set of binding environments \hat{Q} ; this refinement is important to maintain precision during success and failure of individual pattern matches. More formally, we can define refined binding environments as follows:

$$\begin{aligned} \hat{Q}_R \in \widehat{\text{RBindingEnvs}} &= \wp \left(\widehat{\text{Store}} \times \widehat{\text{ValueShape}} \times \wp \left(\widehat{\text{BindingEnv}} \right) \right) \\ \hat{\rho} \in \widehat{\text{BindingEnv}} &= \text{Var} \rightarrow \widehat{\text{ValueShape}} \end{aligned}$$

The AP-VAR rule collects three possible rules for abstract pattern matching against variable patterns: AP-VAR-UNI, AP-VAR-FAIL and AP-VAR-BIND. AP-VAR-UNI pattern matches the values shape against the possibly assigned value shape in the store, refining both shapes assuming that they are equal. The AP-VAR-FAIL rule works analogously, but captures the case where they are possibly not equal. AP-VAR-BIND captures the case where the variable is unassigned in the store, and thus returns a binding of the variable to the input value shape.

$$\begin{aligned} & \frac{\{\hat{\sigma} \vdash x \stackrel{?}{:=} \hat{v}s \xrightarrow[\text{a-match-v1}]{} \hat{Q}_R\}}{\text{AP-VAR} \quad \hat{\sigma} \vdash x \stackrel{?}{:=} \hat{v}s \xrightarrow[\text{a-match}]{} \hat{Q}_R} \\ & \frac{\hat{\sigma}(x) = (b, \hat{v}s') \quad \hat{v}s' \neq \perp_{\widehat{\text{VS}}} \quad \hat{v}s'' \in (\hat{v}s \hat{=} \hat{v}s')}{\text{AP-VAR-UNI} \quad \hat{\sigma} \vdash x \stackrel{?}{:=} \hat{v}s \mid \hat{v}s' \xrightarrow[\text{a-match-v}]{} (\hat{\sigma}[x \mapsto (\text{ff}, \hat{v}s'')], \hat{v}s'', \{\ \})} \\ & \frac{\hat{\sigma}(x) = (b, \hat{v}s') \quad \hat{v}s' \neq \perp_{\widehat{\text{VS}}} \quad (\hat{v}s'', \hat{v}s''') \in (\hat{v}s \hat{\neq} \hat{v}s')}{\text{AP-VAR-FAIL} \quad \hat{\sigma} \vdash x \stackrel{?}{:=} \hat{v}s \mid \hat{v}s' \xrightarrow[\text{a-match-v}]{} (\hat{\sigma}[x \mapsto (\text{ff}, \hat{v}s''')], \hat{v}s'', \emptyset)} \end{aligned}$$

$$\text{AP-VAR-BIND} \frac{\widehat{\sigma}(x) = (\text{tt}, \widehat{vs}')}{\widehat{\sigma} \vdash x := \widehat{vs} \xrightarrow[\text{a-match-v}]{} (\widehat{\sigma}[x \mapsto (\text{tt}, \perp_{\widehat{vs}})], \widehat{vs}, \{[x \mapsto \widehat{vs}]\})}$$

The rule AP-CONS collects two rules for pattern matching against constructors: AP-CONS-SUCS and AP-CONS-FAIL. AP-CONS-SUCS captures the case where the target value shape can be unfolded to the constructor required by the pattern, and then continues to match the inner patterns and merge their results. AP-CONS-FAIL in contrast considers the case where target value can be unfolded to some other constructor or is not a value of target type (producing **error**); in this case, the set of binding environments is empty, and the input value is refined to exclude the target pattern constructor from the result value shape if possible.

$$\text{AP-CONS} \frac{\{\widehat{\sigma} \vdash k(\underline{p}) := \widehat{vs} \mid at \xrightarrow[\text{a-match-cons}]{} \widehat{Q}_R\}}{\widehat{\sigma} \vdash k(\underline{p}) := \widehat{vs} \xrightarrow[\text{a-match}]{} \widehat{Q}_R}$$

$$\text{AP-CONS-SUCS} \frac{\begin{array}{l} \mathbf{data} \text{ at} = \dots \mid k(\underline{t}) \mid \dots \quad (\mathbf{value} \ k(\widehat{vs}')) \in \widehat{\text{unfold}}(\widehat{vs}, at) \\ \widehat{\sigma} \vdash p_1 := \widehat{vs}'_1 \xrightarrow[\text{a-match}]{} \widehat{Q}_{R1} \quad \dots \quad \widehat{\sigma} \vdash p_n := \widehat{vs}'_n \xrightarrow[\text{a-match}]{} \widehat{Q}_{Rn} \\ (\widehat{\sigma}'_1, \widehat{vs}'_1, \widehat{Q}_1) \in \widehat{Q}_{R1} \quad \dots \quad (\widehat{\sigma}'_n, \widehat{vs}'_n, \widehat{Q}_n) \in \widehat{Q}_{Rn} \end{array}}{\widehat{\sigma} \vdash k(\underline{p}) := \widehat{vs} \mid at \xrightarrow[\text{a-match-cons}]{} (\prod_i \widehat{\sigma}_i, k(\widehat{vs}'), \widehat{\text{merge}}(\widehat{Q}))}$$

$$\text{AP-CONS-FAIL} \frac{\begin{array}{l} \mathbf{data} \text{ at} = \dots \mid k(\underline{t}) \mid \dots \\ (\mathbf{value} \ k'(\widehat{vs}')) \in \widehat{\text{unfold}}(\widehat{vs}, at) \wedge k' \neq k \vee \mathbf{error} \in \widehat{\text{unfold}}(\widehat{vs}, at) \end{array}}{\widehat{\sigma} \vdash k(\underline{p}) := \widehat{vs} \mid at \xrightarrow[\text{a-match-cons}]{} (\widehat{\sigma}, \widehat{\text{exclude}}(\widehat{vs}, k), \emptyset)}$$

The AP-SET rule collects two rules: AP-SET-SUCS and AP-SET-FAIL. AP-SET-SUCS tries to unfold the target value shape as a set, and then proceeds by pattern matching the against the inner sequence of patterns, using the content shape and cardinality of the set. AP-SET-FAIL produces a failure result with an empty binding environment set, if the target value shape is possibly not a set.

$$\text{AP-SET} \frac{\{\widehat{\sigma} \vdash \{\star p\} := \widehat{vs} \xrightarrow[\text{a-match-set}]{} \widehat{Q}_R\}}{\widehat{\sigma} \vdash \{\star p\} := \widehat{vs} \xrightarrow[\text{a-match}]{} \widehat{Q}_R}$$

$$\begin{array}{c}
\text{AP-SET-SUCS} \frac{\mathbf{value} \{\widehat{vs}'\}_{[l;u]} \in \widehat{\text{unfold}}(\widehat{vs}, \mathbf{set}\langle \mathbf{value} \rangle) \quad \widehat{\sigma} \vdash \underline{\star p} \stackrel{?}{:=} \widehat{vs} \mid [l;u] \xrightarrow{\text{a-match}\star} \widehat{Q}_R}{\widehat{\sigma} \vdash \{\underline{\star p}\} \stackrel{?}{:=} \widehat{vs} \xrightarrow{\text{a-match-set}} \widehat{Q}_R} \\
\text{AP-SET-FAIL} \frac{\mathbf{error} \in \widehat{\text{unfold}}(\widehat{vs}, \mathbf{set}\langle \mathbf{value} \rangle)}{\widehat{\sigma} \vdash \{\underline{\star p}\} \stackrel{?}{:=} \widehat{vs} \xrightarrow{\text{a-match-set}} (\widehat{\sigma}, \widehat{vs}, \emptyset)}
\end{array}$$

The rule APL collects nine abstract rules for matching sequences of patterns, where each rule corresponding to a concrete pattern matching rule for sequence of patterns. Notice that our abstract rules do not back-track like our concrete rules since we only have a single fixed shape abstraction for all the elements in a collection.

$$\text{APL} \frac{\{\widehat{\sigma} \vdash \underline{\star p} \stackrel{?}{:=} \widehat{vs} \mid [l;u] \xrightarrow{\text{a-match}\star-1} \widehat{Q}_R\}}{\widehat{\sigma} \vdash \underline{\star p} \stackrel{?}{:=} \widehat{vs} \mid [l;u] \xrightarrow{\text{a-match}\star} \widehat{Q}_R}$$

There are two possible abstract rules for pattern matching against an empty sequence of patterns: APL-EMP-BOTH and APL-EMP-PAT. APL-EMP-BOTH accounts for the case where the abstracted sequence is possibly empty (has cardinality 0) producing a set containing an empty binding environment (indicating a successful match with no variable bindings). APL-EMP-PAT conversely accounts for the case where abstracted sequence is possibly non-empty, and therefore produces an empty set of binding environments (indicating a failing match).

$$\text{APL-EMP-BOTH} \frac{l \leq u \quad l = 0}{\widehat{\sigma} \vdash \varepsilon \stackrel{?}{:=} \widehat{vs} \mid [l;u] \xrightarrow{\text{a-match}\star-1} (\widehat{\sigma}, \{\perp_{\widehat{vs}}\}_0, \{\{\}\})}$$

$$\text{APL-EMP-PAT} \frac{l \leq u \quad u \neq 0}{\widehat{\sigma} \vdash \varepsilon \stackrel{?}{:=} \widehat{vs} \mid [l;u] \xrightarrow{\text{a-match}\star-1} (\widehat{\sigma}, \{\widehat{vs}\}_{[\max(1,l),u]}, \emptyset)}$$

There is similarly two rules for pattern matching against a non-empty pattern sequence starting, where the first element is an ordinary pattern: APL-EMP-VAL and APL-MORE-PAT. APL-EMP-VAL accounts for the case where the abstracted set of elements is possibly empty, and so produces an empty set of bindings environments. APL-MORE-PAT accounts for the case where the abstracted set of elements is non-empty, and therefore i) pattern matches the shape of elements against the first pattern,

ii) matches the rest of the pattern sequence against the element shape where the cardinality has been decreased by one, and iii) merges the resulting binding set.

$$\begin{array}{c}
\text{APL-EMP-VAL} \frac{l \leq u \quad l = 0}{\hat{\sigma} \vdash p, \star p := \hat{v}s \mid [l; u] \xrightarrow{\text{a-match}_{\star-1}} (\hat{\sigma}, \{\hat{v}s\}_0, \emptyset)} \\
\\
\text{APL-MORE-PAT} \frac{\begin{array}{c} l \leq u \quad u \neq 0 \quad \hat{\sigma} \vdash p := \hat{v}s \xrightarrow{\text{a-match}} \hat{Q}_R' \\ \hat{\sigma} \vdash \star p := \hat{v}s \mid [l-1; u-1] \xrightarrow{\text{a-match}_{\star}} \hat{Q}_R'' \\ (\hat{\sigma}', \hat{v}s', \hat{Q}') \in \hat{Q}_R' \quad (\hat{\sigma}'', \hat{v}s'', \hat{Q}'') \in \hat{Q}_R'' \\ \hat{Q}_R''' = \{(\hat{\sigma}' \sqcup \hat{\sigma}'', \{\hat{v}s'\}_1 \sqcup \hat{v}s'', \widehat{\text{merge}}(\hat{Q}', \hat{Q}''))\} \end{array}}{\hat{\sigma} \vdash p, \star p := \hat{v}s \mid [l; u] \xrightarrow{\text{a-match}_{\star-1}} \hat{Q}_R'''}
\end{array}$$

Finally, there are five rules for matching against star patterns. Three of the rules are applicable when a value is bound in the store: APL-STAR-UNI, APL-STAR-PAT-FAIL and APL-STAR-VAL-FAIL. APL-STAR-UNI is the most elaborate rule, since it has to account for the case where a subset matched the assigned value in the store; essentially, it follows these steps:

1. It checks whether the assigned value to the variable x can possibly unfold to a set value.
2. It calculates the overlapping shape and cardinality, if there is any overlap.
3. It matches the rest of the pattern sequence with a refined store, and interval abstract the cardinality of the rest of elements (essentially uses the interval difference).
4. It additionally takes into account the case where the concrete sets do not match even though their value shapes and cardinality were compatible and thus adds an empty set of binding environments as a possibility. For example, the assigned value shape and the analyzed value shape could both be $\{\{1;2\}\}_1$ which would have both $\{1\}$ and $\{2\}$ as valid concretizations, which would not match in a concrete evaluation.

The APL-STAR-PAT-FAIL rule accounts for the case where either the assigned analyzed shapes are incompatible, or where the cardinalities are incompatible, producing an empty set of binding environments. Finally, APL-STAR-VAL-FAIL accounts for the case where the assigned value shape is not a set at all.

$$\begin{array}{c}
\text{APL-STAR-UNI} \\
\frac{
\begin{array}{c}
l \leq u \qquad \hat{\sigma}(x) = (b, \hat{v}s') \\
\mathbf{value} \{ \hat{v}s'' \}_{[l''; u'']} \in \widehat{\text{unfold}}(\hat{v}s, \mathbf{set}(\mathbf{value})) \\
\hat{v}s''' \in (\hat{v}s \hat{=} \hat{v}s'') \qquad [l'; u'] = [l''; u''] \sqcap [l; u] \\
[l'; u'] \neq \perp \qquad [l'''; u'''] = [\min(0, l - u'); u - l'] \\
\hat{\sigma}[x \mapsto (\mathbf{ff}, \{ \hat{v}s'' \}_{[l'; u']})] \vdash \star p \stackrel{?}{:=} \hat{v}s \mid [l'''; u'''] \xrightarrow{\text{a-match}\star} \hat{Q}_R'' \\
\hat{Q}_R''' = \{ (\hat{\sigma}'', \{ \hat{v}s''' \sqcup \hat{v}s'''' \}_{[l; u]}, \hat{Q}'') \mid (\hat{\sigma}'', \{ \hat{v}s'''' \}_{[l'''; u''']}, \hat{Q}'') \in \hat{Q}_R'' \} \\
\hat{Q}_R' = \hat{Q}_R''' \cup \{ (\hat{\sigma}[x \mapsto (\mathbf{ff}, \{ \hat{v}s'' \}_I)], \{ \hat{v}s \}_{[l; u]}, \emptyset) \}
\end{array}
}{
\hat{\sigma} \vdash \star x, \star p \stackrel{?}{:=} \hat{v}s \mid [l; u] \xrightarrow{\text{a-match}\star-1} \hat{Q}_R'
}
\\
\\
\text{APL-STAR-PAT-FAIL} \\
\frac{
\begin{array}{c}
l \leq u \qquad \hat{\sigma}(x) = (b, \hat{v}s') \\
\mathbf{value} \{ \hat{v}s'' \}_{[l''; u'']} \in \widehat{\text{unfold}}(\hat{v}s, \mathbf{set}(\mathbf{value})) \\
(\hat{v}s \hat{\neq} \hat{v}s'') \neq \emptyset \vee [l''; u''] \sqcap [l; u] = \perp \\
\hat{Q}_R' = \{ (\hat{\sigma}[x \mapsto (\mathbf{ff}, \{ \hat{v}s'' \}_I)], \{ \hat{v}s \}_{[l; u]}, \emptyset) \}
\end{array}
}{
\hat{\sigma} \vdash \star x, \star p \stackrel{?}{:=} \hat{v}s \mid [l; u] \xrightarrow{\text{a-match}\star} \hat{Q}_R'
}
\\
\\
\text{APL-STAR-VAL-FAIL} \\
\frac{
l \leq u \quad \hat{\sigma}(x) = (b, \hat{v}s') \quad \mathbf{error} \in \widehat{\text{unfold}}(\hat{v}s, \mathbf{set}(\mathbf{value}))
}{
\hat{\sigma} \vdash \star x, \star p \stackrel{?}{:=} \hat{v}s \mid [l; u] \xrightarrow{\text{a-match}\star} (\hat{\sigma}[x \mapsto (\mathbf{ff}, \hat{v}s')], \{ \hat{v}s \}_{[l; u]}, \emptyset)
}
\end{array}$$

The two final rules for star patterns—APL-STAR-RE and APL-STAR-EXH—concern when the target variable is possibly not assigned in the store. APL-STAR-RE tries to bind a subcollection of elements and continue matching the rest of the sequence, where APL-STAR-EXH abstracts over the case where backtracking has exhausted all possible partitions and thus produces no binding environment.

$$\begin{array}{c}
\text{APL-STAR-RE} \frac{l \leq u \quad \widehat{\sigma}(x) = (\text{tt}, \widehat{vs}') \quad \widehat{\sigma}[x \mapsto (\text{tt}, \perp_{\widehat{vs}})] \vdash \star p \stackrel{?}{:=} \widehat{vs} \mid [0; u] \xrightarrow{\text{a-match}\star} \widehat{Q}_R'' \quad (\widehat{\sigma}'', \{\widehat{vs}''\}_{[l'; u'']}, \widehat{Q}'') \in \widehat{Q}_R''}{\widehat{Q}'_R = (\widehat{\sigma}'', \{\widehat{vs}\}_{[l; u]}, \widehat{\text{merge}}(\{[x \mapsto \{\widehat{vs}\}_{[l'; u']}\}], \widehat{Q}''))} \\
\widehat{\sigma} \vdash \star x, \star p \stackrel{?}{:=} \widehat{vs} \mid [l; u] \xrightarrow{\text{a-match}\star-1} \widehat{Q}'_R
\end{array}$$

$$\text{APL-STAR-EXH} \frac{l \leq u \quad \widehat{\sigma}(x) = (\text{tt}, \widehat{vs}')}{\widehat{\sigma} \vdash \star x, \star p \stackrel{?}{:=} \widehat{vs} \mid [l; u] \xrightarrow{\text{a-match}\star-1} (\widehat{\sigma}, \{\widehat{vs}\}_{[l; u]}, \emptyset)}$$

Trace Memoization and Traversals

In abstract interpretation and static program analysis in general, the main challenge is usually how to perform fixed-point calculation for unbounded loops and non-structural recursion. In Schmidt-style abstract interpretation, the main technique to handle this is *trace memoization* [Schmidt, 1998]. The core idea of *trace memoization* is to detect non-structural re-evaluation of the same program element—i.e., where the evaluation of a program element is recursively dependent on the evaluation of itself, like a **while**-loop or traversal—and in such case apply a widening on the input values and state in order to get an over-approximating finite circular dependency in the trace of the program element. This circular dependent trace can be seen as a continuous function [Schmidt, 1998; Rosendahl, 2013], and to get an output we perform classical fixed-point iteration.

The trace memoization strategies suggested by existing work, do however not directly apply to the constructs present in Rascal. Schmidt [1998] presents a simple technique which only works for constructs like **while**-loops where the self-recursive evaluation happens in a tail position and so the resulting output is the same as the resulting output as recursive dependency. The technique suggested by Rosendahl [2013] supports more general recursive dependencies, but only works for finite input domains, where our input domains are all infinite.

We have therefore further extended the trace memoization strategy suggested by Rosendahl [2013] to work with abstract inputs that are

from infinite domains. The extension is still terminating and sound—so the function representing the potentially circular trace is continuous—and additionally it allows calculating results with good precision. The core idea is to partition the infinite input domain using a finite domain of elements, and on recursion degrade input values using previously met input values from the same partition. In particular, our technique works as follows:

- Assume all our domains are widening lattices. Consider a each non-structurally recursive operational semantics judgment $i \Longrightarrow o$, with i being an input from domain $\widehat{\text{Input}}$, and o being the output from domain $\widehat{\text{Output}}$. For this judgment, we associate a memoization map $\widehat{M} \in \widehat{\text{PInput}} \rightarrow \widehat{\text{Input}} \times \widehat{\text{Output}}$ where $\widehat{\text{PInput}}$ is a finite partitioning domain that has a Galois connection with our actual input, i.e. $\widehat{\text{Input}} \xrightleftharpoons[\alpha_{\widehat{\text{PI}}}]^{\gamma_{\widehat{\text{PI}}}} \widehat{\text{PInput}}$. The memoization map keeps track of the previously seen input and corresponding output for values in the partition domain. For example, for input from our value domain $\widehat{\text{Value}}$ we can for example use the corresponding type from the domain Type as input to the memoization map.⁴ So for values 1 and [2;3] we would use `int`, while for `fieldaccessexpr(Expr, str)` we would use the defining data type `Expr`.
- We perform a fixed-point calculation over the evaluation of input i . Initially, the memoization map \widehat{M} is empty, and during evaluation we check whether there was already a value from the same partition as i , i.e., $\alpha_{\widehat{\text{PI}}}(i) \in \text{dom } \widehat{M}$. At each iteration, there is then three possibilities:

Hit The corresponding input partition key is in the memoization map and a less precise input is stored, so $\widehat{M}(\alpha_{\widehat{\text{PI}}}(i)) = (i', o')$ where $i \sqsubseteq_{\widehat{\text{Input}}} i'$. In this case we return o' directly as result without further evaluation; this correctly preserves continuity of the complete semantics since we are returning a less precise result, the one for the less precise input i' , than the one that is expected for the current input i .

Widen The corresponding input partition key is in the memoization map, but an unrelated or more precise input is stored, i.e.,

⁴Provided that we put a fixed-bound on the depth of type parameters of collections.

$\widehat{M}(\alpha_{\widehat{PI}}(i)) = (i'', o'')$ where $i \not\sqsubseteq_{\widehat{\text{Input}}} i''$. In this case we continue evaluation but with a widened input $i' = i'' \nabla_{\widehat{\text{Input}}}(i'' \sqcup i)$ and an updated map $\widehat{M}' = [\alpha_{\widehat{PI}}(i) \mapsto (i', o_{\text{prev}})]$, where o_{prev} is the output of the last fixed-point iteration for i' . If the output of the evaluation of i', o' is different than the previous output o_{prev} , we continue fixed-point iteration.

Miss The corresponding input partition key is not in the memoization map, so $\alpha_{\widehat{PI}}(i) \notin \text{dom } \widehat{M}$ and so we continue evaluation of i with an updated map $\widehat{M}' = \widehat{M}[\alpha_{\widehat{PI}}(i) \mapsto (i, o_{\text{prev}})]$ passed down recursively that assigns the input i and the output of the previous iteration for i in the fixed-point calculation o_{prev} (initially $\perp_{\widehat{\text{Output}}}$) to the relevant input partition key. If the resulting output o is equivalent to the previous output o_{prev} then we stop the fixed-point calculation, and otherwise we perform a new iteration with an updated previous output value $o'_{\text{prev}} = o_{\text{prev}} \nabla o$.

We will now show how we constructor our abstract operational semantics rules for visitors. The required steps to get the abstract rules for non-structurally recursive constructs are the same as the ones for translating rules of basic constructs, but we must additionally take into account the memoization map and the fixed-point calculation.

The first step is to construct the top-level abstract evaluation rules for evaluating the visit expressions, each translated from a concrete evaluation rule for the visit expression. This is done as usual using the technique presented in the section. The AE-VISIT rule collects the four possible cases: AE-VISIT-SUCS, AE-VISIT-FAIL, AE-VISIT-Exc1 and AE-VISIT-Exc2.

$$\text{AE-VISIT} \frac{\{\text{bottom-up visit } e \text{ cs}; \widehat{\sigma} \xrightarrow{\text{a-expr-visit}} \widehat{Res}\}}{\text{bottom-up visit } e \text{ cs}; \widehat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res}}$$

AE-VISIT-SUCS handles the case where both the evaluation of the target expression and the traversal succeeds. AE-VISIT-FAIL handles the case where the traversal fails, and produces a success value with the refined value shape given by **fail**; using the refined value is important to maintain precision, since using the initial input to the traversal like the concrete evaluation would not take into account information learned by pattern matching. Finally, AE-VISIT-Exc1 and AE-VISIT-Exc2 handle

the exceptional cases during evaluation of the target expression and the traversal, respectively.

$$\text{AE-VISIT-SUCS} \frac{\begin{array}{l} e; \hat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res} \quad (\text{success } \widehat{vs}; \hat{\sigma}'') \in \widehat{Res} \\ \underline{cs}; \widehat{vs}; \hat{\sigma}'' \xrightarrow{\text{a-bu-visit}} \widehat{Res}' \quad (\text{success } \widehat{vs}'; \hat{\sigma}') \in \widehat{Res}' \end{array}}{\text{bottom-up visit } e \underline{cs}; \hat{\sigma} \xrightarrow{\text{a-expr-visit}} \text{success } \widehat{vs}'; \hat{\sigma}'}$$

$$\text{AE-VISIT-FAIL} \frac{\begin{array}{l} e; \hat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res} \quad (\text{success } \widehat{vs}; \hat{\sigma}'') \in \widehat{Res} \\ \underline{cs}; \widehat{vs}; \hat{\sigma}'' \xrightarrow{\text{a-bu-visit}} \widehat{Res}' \quad (\text{fail } \widehat{vs}'; \hat{\sigma}') \in \widehat{Res}' \end{array}}{\text{bottom-up visit } e \underline{cs}; \hat{\sigma} \xrightarrow{\text{a-expr-visit}} \text{success } \widehat{vs}'; \hat{\sigma}'}$$

$$\text{AE-VISIT-EXC1} \frac{e; \hat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res} \quad (\widehat{exres}; \hat{\sigma}') \in \widehat{Res}}{\text{bottom-up visit } e \underline{cs}; \hat{\sigma} \xrightarrow{\text{a-expr-visit}} \widehat{exres}; \hat{\sigma}'}$$

$$\text{AE-VISIT-EXC2} \frac{\begin{array}{l} e; \hat{\sigma} \xrightarrow{\text{a-expr}} \widehat{Res} \quad (\text{success } \widehat{vs}; \hat{\sigma}'') \in \widehat{Res} \\ \underline{cs}; \widehat{vs}; \hat{\sigma}'' \xrightarrow{\text{a-bu-visit}} \widehat{Res}' \quad (\text{error}; \hat{\sigma}') \in \widehat{Res}' \end{array}}{\text{bottom-up visit } e \underline{cs}; \hat{\sigma} \xrightarrow{\text{a-expr-visit}} \text{error}; \hat{\sigma}'}$$

The next step is to translate the actual traversal rules for visit using the memoization technique. This requires setting up rules that follow the extended steps for trace memoization and passing around the memoization map between mutually recursive judgments. The AEBU rule initializes memoization process, by delegating to the memoized judgment with an initially empty memoization map.

$$\text{AEBU} \frac{\boxed{} \vdash \underline{cs}; \widehat{vs}; \hat{\sigma} \xrightarrow{\text{a-bu-visit-memo}} \widehat{Res}}{\underline{cs}; \widehat{vs}; \hat{\sigma} \xrightarrow{\text{a-bu-visit}} \widehat{Res}}$$

There are three rules for the memoized evaluation judgment—AEBUM-HIT, AEBUM-WIDEN and AEBUM-MISS—each corresponding to one of the described cases. The AEBUM-HIT finds a less precise input stored at the memoization partition key, and can therefore soundly use the memoized output result set. The AEBUM-WIDEN conversely finds a

more precise input memoized at the partition key, and so must use the current input for widening to a less precise value shape; we then do a fixed-point calculation of the result of traversing the widened input with the bottom value $\perp_{\widehat{VS}}$ as initial pre-result. The AEBUM-MISS rule handles the case where the target value shape does not have a corresponding partition key in the memoization map, and so simply uses the current value for fixed-point calculation with $\perp_{\widehat{VS}}$ as initial pre-result.

$$\begin{array}{c}
\text{AEBUM-HIT} \frac{\widehat{M}(\alpha_{\text{Type}}(\widehat{vs})) = (\widehat{vs}', \widehat{Res}') \quad \widehat{vs} \sqsubseteq_{\widehat{RS}} \widehat{vs}'}{\widehat{M} \vdash \underline{cs}; \widehat{vs}; \widehat{\sigma} \xrightarrow{\text{a-bu-visit-memo}} \widehat{Res}'} \\
\\
\text{AEBUM-WIDEN} \frac{\widehat{M}(\alpha_{\text{Type}}(\widehat{vs})) = (\widehat{vs}'', \widehat{Res}') \quad \widehat{vs} \not\sqsubseteq_{\widehat{RS}} \widehat{vs}'' \quad \widehat{vs}' = \widehat{vs} \nabla_{\widehat{RS}} (\widehat{vs} \sqcup \widehat{vs}'')}{\widehat{M}; \perp_{\widehat{RS}} \vdash \underline{cs}; \widehat{vs}'; \widehat{\sigma} \xrightarrow{\text{a-bu-visit-memo-fix}} \widehat{Res}'} \\
\hline
\widehat{M} \vdash \underline{cs}; \widehat{vs}; \widehat{\sigma} \xrightarrow{\text{a-bu-visit-memo}} \widehat{Res}' \\
\\
\text{AEBUM-MISS} \frac{\alpha_{\text{Type}}(\widehat{vs}) \notin \text{dom } \widehat{M} \quad \widehat{M}; \perp_{\widehat{RS}} \vdash \underline{cs}; \widehat{vs}; \widehat{\sigma} \xrightarrow{\text{a-bu-visit-memo-fix}} \widehat{Res}'}{\widehat{M} \vdash \underline{cs}; \widehat{vs}; \widehat{\sigma} \xrightarrow{\text{a-bu-visit-memo}} \widehat{Res}'}
\end{array}$$

The AEBUMF-STOP and AEBUMF-MORE rules are used for fixed-point calculation over the traversal result. The AEBUMF-STOP rule stops when evaluating the traversal body produces the same result as before iteration. Conversely, the AEBUMF-MORE takes an iteration where the result is strictly decreasing in precision, and so it needs to widen the output result states and repeat iteration. Because of monotonicity, it is not possible to get a result that gets more precise after a fixed-point iteration (otherwise, termination and the existence of fixed-point is not certain).

$$\begin{array}{c}
\widehat{M}[\alpha_{\text{Type}}(\widehat{vs}) \mapsto (\widehat{vs}, \widehat{Res}_{\text{prev}})] \vdash \underline{cs}; \widehat{vs}; \widehat{\sigma} \xrightarrow{\text{a-bu-visit-memo-go}} \widehat{Res} \\
\widehat{Res}_{\text{prev}} = \widehat{Res} \\
\hline
\widehat{M}; \widehat{Res}_{\text{prev}} \vdash \underline{cs}; \widehat{vs}; \widehat{\sigma} \xrightarrow{\text{a-bu-visit-memo-fix}} \widehat{Res} \\
\\
\text{AEBUMF-MORE} \frac{\widehat{M}[\alpha_{\text{Type}}(\widehat{vs}) \mapsto (\widehat{vs}, \widehat{Res}_{\text{prev}})] \vdash \underline{cs}; \widehat{vs}; \widehat{\sigma} \xrightarrow{\text{a-bu-visit-memo-go}} \widehat{Res} \quad \widehat{Res}_{\text{prev}} \sqsubset \widehat{Res} \quad \widehat{Res}'' = \widehat{Res}_{\text{prev}} \nabla_{\widehat{RS}} \widehat{Res}}{\widehat{M}; \widehat{Res}'' \vdash \underline{cs}; \widehat{vs}; \widehat{\sigma} \xrightarrow{\text{a-bu-visit-memo-fix}} \widehat{Res}'} \\
\hline
\widehat{M}; \widehat{Res}_{\text{prev}} \vdash \underline{cs}; \widehat{vs}; \widehat{\sigma} \xrightarrow{\text{a-bu-visit-memo-fix}} \widehat{Res}'
\end{array}$$

Having set-up the elaborate memoization machinery, we proceed by translating the traversal rules as usual. The only difference is that the memoization environment must be carried around and the memoized judgment must be used on recursion. The AEBUMG collect the four different bottom-up traversal rules each roughly corresponding to a single rule in the concrete semantics.⁵

$$\text{AEBUMG} \frac{\{\widehat{M} \vdash \underline{cs}; \widehat{vs}; \widehat{\sigma} \xrightarrow{\text{a-bu-visit-go}} \widehat{Res}\}}{\widehat{M} \vdash \underline{cs}; \widehat{vs}; \widehat{\sigma} \xrightarrow{\text{a-bu-visit-memo-go}} \widehat{Res}}$$

The AEBUG-SUCS rule handles the case where the traversal succeeds both for the contained values⁶ and the reconstructed top-level value. The AEBUG-FAIL-SUCS rule handles the case where the traversal fails for the contained values and thus must be carried out only on the top-level value; here, we expect the reconstruction of values to always succeed, because the values returned in failure should be refinements of the original values. The AEBUG-EXC rule handles errors when traversing contained values, and AEBUG-ERR rule handles the case where reconstruction produces an error.

$$\text{AEBUG-SUCS} \frac{\begin{array}{l} (\widehat{vs}', \widehat{vs}'') \in \widehat{\text{children}}(\widehat{vs}) \quad \widehat{M} \vdash \underline{cs}; \widehat{vs}''; \widehat{\sigma} \xrightarrow{\text{a-bu-visit}\star} \widehat{Res}\star \\ (\text{success } \widehat{vs}'''; \widehat{\sigma}') \in \widehat{Res}\star \quad \text{recons } \widehat{vs}' \text{ using } \widehat{vs}''' \text{ to } \widehat{RCRes} \\ (\text{success } \widehat{vs}'''' \in \widehat{RCRes} \quad \widehat{M} \vdash \underline{cs}; \widehat{vs}''''; \widehat{\sigma}' \xrightarrow{\text{a-bu-visit-memo}} \widehat{Res}' \end{array}}{\widehat{M} \vdash \underline{cs}; \widehat{vs}; \widehat{\sigma} \xrightarrow{\text{a-bu-visit-go}} \widehat{Res}'}$$

$$\text{AEBUG-FAIL-SUCS} \frac{\begin{array}{l} (\widehat{vs}', \widehat{vs}'') \in \widehat{\text{children}}(\widehat{vs}) \quad \widehat{M} \vdash \underline{cs}; \widehat{vs}''; \widehat{\sigma} \xrightarrow{\text{a-bu-visit}\star} \widehat{Res}\star \\ (\text{fail } \widehat{vs}'''; \widehat{\sigma}') \in \widehat{Res}\star \quad \text{recons } \widehat{vs}' \text{ using } \widehat{vs}''' \text{ to success } \widehat{vs}'''' \\ \widehat{M} \vdash \underline{cs}; \widehat{vs}''''; \widehat{\sigma}' \xrightarrow{\text{a-bu-visit-memo}} \widehat{Res}' \end{array}}{\widehat{M} \vdash \underline{cs}; \widehat{vs}; \widehat{\sigma} \xrightarrow{\text{a-bu-visit-go}} \widehat{Res}'}$$

$$\text{AEBUG-Exc} \frac{\begin{array}{l} (\widehat{vs}', \widehat{vs}'') \in \widehat{\text{children}}(\widehat{vs}) \quad \widehat{M} \vdash \underline{cs}; \widehat{vs}''; \widehat{\sigma} \xrightarrow{\text{a-bu-visit}\star} \widehat{Res}\star \\ (\text{error}; \widehat{\sigma}') \in \widehat{Res}\star \end{array}}{\widehat{M} \vdash \underline{cs}; \widehat{vs}; \widehat{\sigma} \xrightarrow{\text{a-bu-visit-go}} \text{error}; \widehat{\sigma}'}$$

⁵Not counting the additional rules for the **break** version.

⁶We are here considering the case where the `children` function was applied to data values and so produces a fixed sequence of heterogeneous value shapes. In case of collections or \top , it may produce a value shape-cardinality pair instead.

$$\begin{array}{c}
(\widehat{vs}', \widehat{vs}'') \in \widehat{\text{children}}(\widehat{vs}) \quad \widehat{M} \vdash \underline{cs}; \widehat{vs}''; \widehat{\sigma} \xrightarrow{\text{a-bu-visit}^*} \widehat{Res}^* \\
(\text{success } \widehat{vs}'''; \widehat{\sigma}') \in \widehat{Res}^* \quad \widehat{\text{recons}} \widehat{vs}' \text{ using } \widehat{vs}''' \text{ to } \widehat{RCRes} \\
\text{AEBUSG-ERR} \frac{\text{error} \in \widehat{RCRes}}{\widehat{M} \vdash \underline{cs}; \widehat{vs}; \widehat{\sigma} \xrightarrow{\text{a-bu-visit-go}} \text{error}; \widehat{\sigma}'}
\end{array}$$

We do a similar translation with bottom-up traversal of a sequence of values. The AEBUS rule collects the four rules for traversing a sequence of expressions: AEBUSG-EMP, AEBUSG-MORE, AEBUSG-Exc1 and AEBUSG-Exc2. Note that in order we still need to carry around the memoization environment and recursively use the memoized judgment.

$$\text{AEBUS} \frac{\{\widehat{M} \vdash \underline{cs}; \widehat{vs}; \widehat{\sigma} \xrightarrow{\text{a-bu-visit}^*} \widehat{Res}^*\}}{\widehat{M} \vdash \underline{cs}; \widehat{vs}; \widehat{\sigma} \xrightarrow{\text{a-bu-visit}^*} \widehat{Res}^*}$$

The AEBUSG-EMP rule handles the case where the given abstract value shape is sequence, and so fails to match any case. The AEBUSG-MORE rule evaluates the first value shape in the sequence using the memoized bottom-up traversal rule, continues traversal of the rest of the sequences and combines the results. Finally, the AEBUSG-Exc1 and AEBUSG-Exc2 rules handle the cases where sub-traversals produced errors.

$$\begin{array}{c}
\text{AEBUSG-EMP} \frac{}{\widehat{M} \vdash \underline{cs}; \varepsilon; \widehat{\sigma} \xrightarrow{\text{a-bu-visit}^*} \text{fail } \varepsilon; \widehat{\sigma}} \\
\text{AEBUSG-MORE} \frac{\begin{array}{c} \widehat{M} \vdash \underline{cs}; \widehat{vs}; \widehat{\sigma} \xrightarrow{\text{a-bu-visit-memo}} \widehat{Res} \quad (\widehat{vfres}; \widehat{\sigma}'') \in \widehat{Res} \\ \widehat{M} \vdash \underline{cs}; \widehat{vs}'; \widehat{\sigma}'' \xrightarrow{\text{a-bu-visit}^*} \widehat{Res}^* \quad (\widehat{vfres}^*; \widehat{\sigma}') \in \widehat{Res}^* \end{array}}{\widehat{M} \vdash \underline{cs}; \widehat{vs}, \widehat{vs}'; \widehat{\sigma} \xrightarrow{\text{a-bu-visit}^*} \text{vcombine}(\widehat{vfres}, \widehat{vfres}^*); \widehat{\sigma}'} \\
\text{AEBUSG-Exc1} \frac{\widehat{M} \vdash \underline{cs}; \widehat{vs}; \widehat{\sigma} \xrightarrow{\text{a-bu-visit-memo}} \widehat{Res} \quad (\text{error}; \widehat{\sigma}') \in \widehat{Res}}{\widehat{M} \vdash \underline{cs}; \widehat{vs}, \widehat{vs}'; \widehat{\sigma} \xrightarrow{\text{a-bu-visit}^*} \text{error}; \widehat{\sigma}'} \\
\text{AEBUSG-Exc2} \frac{\begin{array}{c} \widehat{M} \vdash \underline{cs}; \widehat{vs}; \widehat{\sigma} \xrightarrow{\text{a-bu-visit-memo}} \widehat{Res} \quad (\widehat{vfres}; \widehat{\sigma}'') \in \widehat{Res} \\ \widehat{M} \vdash \underline{cs}; \widehat{vs}'; \widehat{\sigma}'' \xrightarrow{\text{a-bu-visit}^*} \widehat{Res}^* \quad (\text{error}; \widehat{\sigma}') \in \widehat{Res}^* \end{array}}{\widehat{M} \vdash \underline{cs}; \widehat{vs}, \widehat{vs}'; \widehat{\sigma} \xrightarrow{\text{a-bu-visit}^*} \text{error}; \widehat{\sigma}'}
\end{array}$$

As observed, the translation of the traversal rules with memoization was done fairly systematically. This provides hope that the presented trace memoization technique can be automated for operational semantics rules of certain shapes, extending the technique presented by Bodin et al. [2015].

Soundness of Abstract Interpreter

The main goal of abstract interpretation is to provide *sound* analyses, where the properties shown on the approximated semantics can be depended upon to hold for the concrete semantics. For our particular abstract interpreter, the main meta-property we want to be satisfied is that the abstract expression evaluation produces values that truly over-approximate the results from all possible concrete evaluations. Formally, this can be stated as:

Theorem 7.8 (Soundness). *For all valid expressions e , concrete stores σ and over-approximating abstract stores $\widehat{\sigma}$, so $\sigma \in \gamma_{\text{Store}}(\widehat{\sigma})$, where we have a concrete evaluation derivation $e; \widehat{\sigma} \xrightarrow[\text{expr}]{} vres; \sigma'$ and corresponding abstract evaluation derivation $e; \widehat{\sigma} \xrightarrow[a\text{-expr}]{} \widehat{Res}$ then it holds that the abstract result set properly over-approximates the concrete result value result and store, i.e., $(vres, \sigma') \in \gamma_{RS}(\widehat{Res})$.*

7.5 Evaluation

The objective of our evaluation is to examine whether our presented abstract interpretation technique works on a wide variety of the transformations written in Rascal.

Research Questions To fulfill our objective, we aim to answer the following research questions:

- Is it possible to abstractly interpret realistic Rascal programs using type and inductive shape analysis, to verify properties of interest?
- What interesting properties require more than inductive shapes to capture, and what extensions are needed?

Subject Selection Our selection of subject transformations is based on three criteria:

1. the evaluated transformations must be realistic and preferably from independent code bases,
2. they must be translatable to our supported subset of Rascal,
3. they must exercise important Rascal constructs like visitors and the expressive pattern matching capabilities.

The first criteria is there to reduce experimental bias from the authors, the second criteria is there to ensure that our tool is able to successfully run on the subjects, and the third is there to ensure that we are not testing trivial or non-interesting transformations.

Subject Transformations We consider four subject transformations in our evaluation, where two of the transformations are extracted directly from real open source Rascal projects. The source code of the transformations is included in Appendix G, but we will discuss them in a high-level fashion here.

NNF The Negation Normal Form transformation [Harrison, 2009, Section 2.5] is a classical transformation in automated reasoning, which translates propositional formulae to a form that has conjunction and disjunction as the only connectives (e.g., no implication), and where negations are only applied to atomic propositions.

RSF The Rename Struct Field refactoring changes the name of a field in a struct, and correspondingly ensures that all field access expressions of that field are renamed correctly as well.

DSO0 The Desugar Oberon-0 transformation, translates for-loops and switch-statements—which are seen as syntactic sugar—to while-loops and nested if-statements respectively. The transformation is part of the Oberon-0 [Wirth, 1996] implementation in Rascal [Basten et al., 2015], which contains all the necessary stages for compiling a structured imperative programming language.

G2PE Glagol⁷ is a DSL for REST-like web development, generating PHP code as output. The Glagol to PHP Expression transpilation generates PHP code for input Glagol expressions.

⁷<https://github.com/BulgariaPHP/glagol-dsl>

The size and target verification properties for each transformation are presented in Table Table 7.1. While the size of the transformation varies, it should be noted that they all use the expressive high-level features in Rascal and so are significantly more succinct than comparable code in non-high-level transformation languages.

The properties checked mainly focus that the output of each transformation satisfies required shape constraints. For the NNF transformation we check that the output indeed is in negation normal form, so it does not have implications (P1) and that all negations are in front of atoms (P2). For the structure field renaming transformation, we focus on properties that show that the old field has correctly been renamed to the new field name and so the old field name is not defined in a structure anymore (P3) and similarly with all field accesses (P4); note, that we need to convert field names to finite enumerations—with constructors representing the old field name, new field name and other names—instead of strings in order for our executor to be able to infer required properties. For the Oberon-0 desugaring, we check that there are no syntactic sugar constructs in the output, i.e. for-loops (P5) and switch-statements (P6), and that auxiliary data values used intermediately in the transformation are eliminated in the output (P7). Finally, we check that the Glagol to PHP translation works as expected, by checking that given inputs with a particular shape the output has a similar shape, e.g. that translating simple arithmetic-logic expressions produces similar arithmetic expressions in PHP (P8)—with no object creation, field or array access, or method invocation—and that Glagol expressions not using negation or unary plus produce no unary PHP expressions.

Threats to Validity We have selected a sample of available Rascal transformations to evaluate on, and so it can be hard to generalize whether our technique is successful for other transformations. We mitigated this by selecting our subject transformations such that they are from realistic projects and vary in authors, programming style and purpose.

We have ported our subject transformations to the Rascal Light subset in a self-contained manner to be readily executed by our tool; this process could potentially introduce some minor differences in the semantics. We mitigated this by only performing minimal changes required to the code, and ensuring that the changes done do not critically affect the credibility of our overall abstract interpretation technique.

Implementation

We have implemented our abstract interpretation technique⁸ as a prototype tool for the complete Rascal Light subset presented in Chapter 6. The abstract semantics of the additional constructs follows is implemented using the process described in Section 7.4; the major challenges included fine-tuning the trace memoization strategies to the different looping constructs and functions (to handle recursion), handling extended result state with more control flow constructs, and taking into account possibly undefined values (■) during evaluation. Using types for partitioning the input during trace memoization can be imprecise, so we use a more precise partitioning strategy when needed, which for data refinements additionally uses the set of present constructors.

⁸Currently with a simple one-value domain for integers instead of intervals, which were not strictly needed for our evaluation subjects.

Transformation	LOCs	Target Property	
NNF	15	P1	Implication is not used as a connective in the result
		P2	All negations in the result are in front of atoms
RSF	35	P3	Structures should not define a fields with the old field name
		P4	There should not be any field access expression to the old field name
DSO0	125	P5	For-loops should be correctly desugared to while-loops
		P6	Switch-statements should be correctly desugared to if-statements
		P7	No auxiliary data in output
G2PE	350	P8	Only have simple PHP expressions in output given simple Glagol expressions in input
		P9	Not have unary PHP expressions if there were no sign markers or negations in the input Glagol expression

Table 7.1: Subject Transformations Statistics and Target Properties

Results

We ran our abstract interpretation tool on the evaluation subjects using Scala 2.12.2 on a 2012 Core i5 MacBook Pro. We present an overview of the results in Table 7.2 detailing the runtime and verified properties; the complete output of each run is included in Appendix H.

The runtime of the analysis is reasonable given the expressiveness of the available features, running around 4-39 seconds; separate runtimes for the G2PE transformation were included because the shape of input varies between the specified properties. There is room for improving performance, since during fixed-point calculation subterms are re-evaluated despite providing the same results.

Transformation	Runtime (s)	Property	Verified
NNF	31	P1	✓
		P2	✓
RSF	6	P3	✗
		P4	✓
DSO0	36	P5	✓
		P6	✓
		P7	✗
G2PE	4	P8	✓
	39	P9	✓

Table 7.2: Results of abstractly interpreting subject programs. Runtime is median of three runs.

All the specified target properties were successfully verified, except for two where the result shape was not precise enough; we provide a further discussion later to describe possible extensions that would allow enough precision to capture these properties.

As an example of the the results, we show the inferred output shape for our NNF transformation in Figure 7.5. To check that our properties were satisfied, we simply had to check that implication is not part of the set of available constructors (P1), and that the negation constructor only allows atoms as subformulae instead of any general formula (P2).

Discussion

There were two properties, P3 and P7, which could not be verified from the resulting shape of the inferred inductive data refinements. Using

```

1 refine Formula#nnf = and(Formula#nnf, Formula#nnf) | atom(str)
2                       | neg(atom(str)) | or(Formula#nnf, Formula#nnf)

```

Figure 7.5: Inferred shape for Negation Normal Form transformation

these properties as a starting point, we discuss three possible extensions to our current abstract domain that could help us get more precise results: relational constraints, abstract attributes and complement domains.

Relational Constraints Relational abstract interpretation [Mycroft and Jones, 1985] allows specifying constraints that relate values across different variables, and even inside and across substructures [Chang and Rival, 2008; Halbwachs and Péron, 2008; Liu and Rival, 2017]. This allows both maintaining a greater precision across results, and inferring properties that are non-local, e.g., that two values not only have similar shape but are actually equal.

Figure 7.6 shows the inferred refinement for the structure field renaming transformation: we can see that the structures contained in target packages remain unrefined and thus P3 does not hold since old field name should not be in definitions. We assume that structures and fields are stored in the containing maps using their names as key, but without relational constraints it is not possible to express this property; this disallows strong map updates and the update to the field definitions is thus not captured. Relational properties would also allow supporting the version of the refactoring that works on classes and additionally takes typing into account.

Abstract Attributes Property P7 provides more interesting verification challenges that pertain particularly to transformations. Recall that the property concerns proving elimination of auxiliary data from the result; concretely the auxiliary statement `begin` is used to intermediately allow returning a list of statements where a single statement is required, and a separate pass is used to flatten the contained list of statements contained with the surrounding list of statements (see Figure 7.7).

Verifying properties that only hold after fixed-point iteration, seems suitably tackled using *abstract attributes* which extract additional information about target structures. For this particular property, a generalization of the multiset abstraction suggested by Perrelle and Halbwachs

```

1  refine Package#rnf =
2    package(map[str, Struct], map[str, Function#rnf])
3  data Struct = struct(str name, map[Nominal, Field] fields);
4  data Field = field(Nominal name, str typ);
5  refine Function#rnf =
6    function(str, str, list[Parameter], Stmt#rnf)
7  refine Stmt#rnf =
8    assignstmt(Expr#rnf, Expr#rnf) | block(list[Stmt#rnf])
9    | ifstmt(Expr#rnf, Stmt#rnf, Stmt#rnf) | returnstmt(Expr#rnf)
10 refine Expr#rnf =
11   fieldaccessexpr(Expr#rnf, Nominal#rnf)
12   | functioncallexpr(Expr#rnf, str, list[Expr#rnf]) | varexpr(str)
13 data Nominal = nfn() | ofn() | other();
14 refine Nominal#rnf = nfn() | other()

```

Figure 7.6: Inferred shape for Rename Struct Field transformation

```

1 public list[Statement] flattenBegin(list[Statement] stats) {
2   return innermost visit (stats) {
3     case [*Statement s1, begin(b), *Statement s2] =>
4       s1 + b + s2
5   }
6 }

```

Figure 7.7: Flattening auxiliary begin statement

[2010] for data types, could be useful to track the number of elements of various sub-elements, e.g., the number of begin statements. Using techniques from term rewriting [Dershowitz and Manna, 1979] one can then show a decrease in number of elements each iteration, inferring that they get eliminated when the fixed-point is reached. Other abstract attributes—such as suggested by Bouajjani et al. [2012] and Pham and Whalen [2013]—can capture other properties of interest like contained set of elements or size of a data type, and in conjunction with relational constraints show meaningful relations between input and output structures, e.g., that the output uses the same variables as the input.

Complement Domains The output for DSO0 and G2PE additionally produced false positive errors, due to values being inferred as potentially undefined ■. The main culprit was that our analysis did not precisely enough capture complementary information—i.e. absence of type or shape—for non-data type values during pattern matching, and so some **switch**-statements fell-through. This can be solved by having a complementary domain component of disallowed types and values for pattern matching, similarly to what we had for TRON in Chapter 5.

7.6 Related Work

Data and Shape Domains Our discussion presented future directions that combine our inductive refinement with ideas from existing work that support more expressive constraints. Bouajjani et al. [2012] allows specifying a flexible set of abstract attributes on lists, both structural such as acyclicity and data-relational such as equality of contained multi-sets of elements. Techniques suggested by Chang and Rival [2008], Vazou et al. [2013], and Albarghouthi et al. [2015] support inferring inductive relational properties for general data-types—such as the binary tree property—but require a pre-specified structure specifying the possible places relational refinement could happen.

Modular Program Analysis Cousot and Cousot [2002] present a general framework for modularly constructing program analyses, but it requires a language with a compositional control flow which Rascal does not have. The framework suggests using symbolic relational domains for domain composition, but it is generally undecidable to infer inductive properties for arbitrary sets of constraints [Padon et al., 2016]. Calcagno et al. [2011] suggests using *bi-abduction* for compositional analysis of programs, but their framework is primarily focused towards pointer programs with only a fixed set of non-nested data types supported. Toubhans et al. [2013] and Rival et al. [2014] further develops the ideas of modular domain design for pointer-manipulating programs supporting a rich set of data abstractions; our domain construction focuses more on pure data-structures with a more heterogeneous construction—i.e., where we have different kinds of algebraic data types, collections, and basic values—still allowing modular and automated inference of inductive refinements for these types.

Verification of Transformations There are various approaches [Jackson et al., 2011; Büttner et al., 2012; Wang et al., 2014] in verification of model transformations that encode transformations as declarative formulae for automated solvers, but such techniques only work for transformations of limited expressiveness and are due to the black-box nature of automated solvers hard to extend. Techniques for model transformation verification based on static analysis such as in Cuadrado et al. [2017] scale to more expressive transformation features,

but are currently focused on verification of rule errors based on types and undefinedness.

Semantic typing has been used by Castagna and Nguyen [2008] and Benzaken et al. [2013] to infer recursive type and shape properties for language with high-level constructs for querying and iteration. The languages considered are small calculi with limited expressiveness compared to the fully feature transformation language subset of Rascal we consider, we have a more flexible domain and abstract interpreter design, and an extensive evaluation showing that our works well with realistic Rascal programs.

Another way to verify transformations using abstract interpretation has been presented by Rival [2004], which uses translation validation and symbolic transfer functions to verify equivalence between a program annotated with abstract invariants and its translated assembly output; The technique is complementary to ours, since it can verify very rich semantic properties but only for concrete input, while our technique verifies simpler properties as shapes but for all possible input.

7.7 Recap

We have presented a formal Schmidt-style abstract interpreter for Rascal Light, that supports verifying type and inductive shape properties. Our interpreter supported all the required high-level transformation features including traversals, backtracking, exceptions and fixed-point iteration. We further presented a practical technique for modularly constructing abstract domains and adapted the idea of *trace memoization* to work with data from infinite abstract domains by using finite input-determined partitioning to widen input on recursion and ensure termination of our algorithm. Finally, we evaluated our algorithm on a variety of realistic transformations showing that it is possible to efficiently and effectively verify inductive shape properties of interest.

Chapter 8

Conclusion

The primary goal for this dissertation was to use traditional techniques in programming languages as a basis for developing foundational theories and tools for transformation languages with high-level features such as traversals, backtracking and fixed-point iteration. The goal of this thesis has been successfully reached by a series of contributions (C1-C5) which together fulfilled the stated objectives (O1-O4).

In Chapter 3, I presented a detailed analysis of a wide selection of declarative transformation languages (C1) using available literature to identify key features and clarify how they contributed to the target languages' computational expressiveness (O1). My analysis showed that all examined declarative transformation languages, except the bidirectional ones, were Turing-complete and so could not be treated specially with regards to verification. This compelled me to focus on general purpose high-level transformation languages like TXL and Rascal, which had the advantage of having and explicit control flow and more commonalities with existing programming languages, making it easier to adapt target programming language verification techniques.

Chapter 4 presented a joint effort to design and validate an industrial transformation (C2) modernizing an industrial configuration tool from a C++ imperative implementation to logical constraints. The effort contributed to understanding how the high-level transformation features were used in practice, and provided first-hand experience with translation validation and symbolic execution (O2). A major concrete outcome was that it displayed the necessity of formal verification techniques for transformations, catching 50 serious bug cases in an expert-written otherwise well-tested transformation.

To allow for effective automated white-box testing of transformations (O3), I had in Chapter 5 developed a formal symbolic execution algorithm for a core transformation language, TRON, which captured key high-level transformation language features such including deep type-directed traversal (C3). The algorithm was implemented as a prototype tool and evaluated using a set of realistic refactorings and model transformations against a baseline black-box algorithm optimizing for literature-suggested metrics; the evaluation showed that our algorithm compared very favorably to the baseline black-box algorithm, achieving almost full coverage on all subject programs, albeit at a cost of slower—but still reasonable—execution time.

I formalized an extensive subset of the operational part of the Rascal transformation language—called Rascal Light—in Chapter 6, which is the first formalization for a large subset of the Rascal language we know of (C4). The formalization has been based on the Rascal implementation, documentation and personal correspondence with developers, it has been implemented as a prototype and tested with a set of Rascal programs, and a series of semantic properties regarding strong typing, safety and termination have been formally proven. The availability of this formalization allows principled development of foundational verification techniques including abstract interpretation-based static analysis (O4), and other complementary semantics like a type system or axiomatic semantics.

Finally, I implemented a Schmidt-style abstract interpreter (C5) for Rascal in Chapter 7 that allowed verification of types and inductive shape properties (O4). The work provided a practical path towards modular design of abstract domains, and an adaption of the idea of *trace memoization*—which allows discovering fixed-points in operational semantics—to support working precisely with input from infinite domains, by utilising input-determined finite partitioning. The technique has been implemented as a prototype tool, which have been evaluated on realistic transformations of different kinds written in Rascal, and was able to infer complex shape properties such as that the output of the negation normal form transformation only had negations in front of atoms, and that a transpilation of an expression with no negative or positive sign operators produced no output unary PHP expressions.

Bibliography

- Aditya Agrawal, Zsolt Kalmar, Gabor Karsai, Feng Shi, and Attila Vizhanyo. *GReAT User Manual*, 2003. URL <http://www.escherinstitute.org/Plone/tools/suites/mic/great/GReAT%20User%20Manual.pdf>.
- Alfred V. Aho, Ravi Sethi, and Jeffrey D. Ullman. *Compilers: Principles, Techniques, and Tools*. Addison-Wesley series in computer science / World student series edition. Addison-Wesley, 2nd edition, 1986a. ISBN 0-201-10088-6. URL <http://www.worldcat.org/oclc/12285707>.
- Alfred V. Aho, Ravi Sethi, and Jeffrey D. Ullman. *Compilers: Principles, Techniques, and Tools*, chapter 8. In *Addison-Wesley series in computer science / World student series edition* Aho et al. [1986a], 2nd edition, 1986b. ISBN 0-201-10088-6. URL <http://www.worldcat.org/oclc/12285707>.
- Alfred V. Aho, Ravi Sethi, and Jeffrey D. Ullman. *Compilers: Principles, Techniques, and Tools*, chapter 9-11. In *Addison-Wesley series in computer science / World student series edition* Aho et al. [1986a], 2nd edition, 1986c. ISBN 0-201-10088-6. URL <http://www.worldcat.org/oclc/12285707>.
- Alexander Aiken and Brian R. Murphy. Implementing regular tree expressions. In *Functional Programming Languages and Computer Architecture, 5th ACM Conference, Cambridge, MA, USA, August 26-30, 1991, Proceedings*, pages 427–447, 1991. URL https://doi.org/10.1007/3540543961_21.
- Robert L. Akers, Ira D. Baxter, Michael Mehlich, Brian J. Ellis, and Kenn R. Luecke. Case study: Re-engineering C++ component models via automatic program transformation. *Information & Software Technology*, 49(3):275–291, 2007. doi: 10.1016/j.infsof.2006.10.012. URL <http://dx.doi.org/10.1016/j.infsof.2006.10.012>.

- Ahmad Salim Al-Sibahi. On the computational expressiveness of model transformation languages. *IT University. Technical Report Series*, January 2015. ISSN 1600-6100.
- Ahmad Salim Al-Sibahi. The Formal Semantics of Rascal Light. *CoRR*, abs/1703.02312, 2017a. URL <http://arxiv.org/abs/1703.02312>.
- Ahmad Salim Al-Sibahi. Abstract interpretation of high-level transformations. In *Student Research Competition at 44th ACM SIGPLAN Symposium on Principles of Programming Languages 2017, POPL SRC 2017, Paris, France, January 15-21, 2017*, 2017b. Poster.
- Ahmad Salim Al-Sibahi, Aleksandar Dimovski, and Andrzej Wasowski. SymexTRON: Symbolic execution of high-level transformation languages: Symbolic execution of high-level transformations. *IT University. Technical Report Series*, September 2016. ISSN 1600-6100.
- Ahmad Salim Al-Sibahi, Aleksandar S. Dimovski, and Andrzej Wasowski. Symbolic execution of high-level transformations. In *Proceedings of the 2016 ACM SIGPLAN International Conference on Software Language Engineering, Amsterdam, The Netherlands, pages 207–220, 2016*. doi: 10.1145/2997364.2997382.
- Ahmad Salim Al-Sibahi, Aleksandar S. Dimovski, Thomas P. Jensen, and Andrzej Wasowski. Verifying Transformations using Inductive Shape Analysis. 2017.
- Jesse Alama. The lambda calculus. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Summer 2014 edition, 2014.
- Aws Albarghouthi, Josh Berdine, Byron Cook, and Zachary Kincaid. Spatial interpolants. In *Programming Languages and Systems - 24th European Symposium on Programming, ESOP 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015. Proceedings*, pages 634–660, 2015. doi: 10.1007/978-3-662-46669-8_26. URL https://doi.org/10.1007/978-3-662-46669-8_26.
- Freddy Allilaire and Frédéric Jouault. Families to persons: A simple illustration of model-to-model transformation, 2007. URL https://www.eclipse.org/at1/documentation/old/ATLUseCase_Families2Persons.pdf.

- Nada Amin and Tiark Rompf. Type soundness proofs with definitional interpreters. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*, pages 666–679, 2017. URL <http://dl.acm.org/citation.cfm?id=3009866>.
- Moussa Amrani, Levi Lucio, Gehan M. K. Selim, Benoît Combemale, Jürgen Dingel, Hans Vangheluwe, Yves Le Traon, and James R. Cordy. A tridimensional approach for studying the formal verification of model transformations. In *Fifth IEEE International Conference on Software Testing, Verification and Validation, ICST 2012, Montreal, QC, Canada, April 17-21, 2012*, pages 921–928, 2012. doi: 10.1109/ICST.2012.197. URL <https://doi.org/10.1109/ICST.2012.197>.
- Franz Baader and Tobias Nipkow. *Term rewriting and all that*. Cambridge University Press, 1998. ISBN 978-0-521-45520-6.
- Roland Carl Backhouse, Jeremy Gibbons, Ralf Hinze, and Johan Jeuring, editors. *Datatype-Generic Programming - International Spring School, SSDGP 2006, Nottingham, UK, April 24-27, 2006, Revised Lectures*, volume 4719 of *Lecture Notes in Computer Science*, 2007. Springer. ISBN 978-3-540-76785-5. doi: 10.1007/978-3-540-76786-2. URL <https://doi.org/10.1007/978-3-540-76786-2>.
- Daniel Balasubramanian, Anantha Narayanan, Christopher van Buskirk, and Gabor Karsai. The graph rewriting and transformation language: Great. *Electronic Communications of the EASST*, 1, 2007.
- András Balogh and Dániel Varró. Advanced model transformation language constructs in the viatra2 framework. In *Proceedings of the 2006 ACM Symposium on Applied Computing, SAC '06*, pages 1280–1287, New York, NY, USA, 2006. ACM. ISBN 1-59593-108-2. doi: 10.1145/1141277.1141575. URL <http://doi.acm.org/10.1145/1141277.1141575>.
- Hans Barendt. *The Lambda Calculus: Its Syntax and Semantics*. North Holland, 2 edition, 1985. ISBN 978-0444867483.
- Bruno Barroca, Levi Lucio, Vasco Amaral, Roberto Félix, and Vasco Sousa. Dsltrans: A turing incomplete transformation language. In Brian A. Malloy, Steffen Staab, and Mark van den Brand, editors, *Software Language Engineering - Third International Conference, SLE 2010*,

- Eindhoven, The Netherlands, October 12-13, 2010, Revised Selected Papers*, volume 6563 of *Lecture Notes in Computer Science*, pages 296–305. Springer, 2010. ISBN 978-3-642-19439-9. doi: 10.1007/978-3-642-19440-5_19. URL http://dx.doi.org/10.1007/978-3-642-19440-5_19.
- Bas Basten, Jeroen van den Bos, Mark Hills, Paul Klint, Arnold Lankamp, Bert Lisser, Atze van der Ploeg, Tijs van der Storm, and Jurgen J. Vinju. Modular language implementation in rascal - experience report. *Sci. Comput. Program.*, 114:7–19, 2015. URL <http://dx.doi.org/10.1016/j.scico.2015.11.003>.
- I.D. Baxter, C. Pidgeon, and M. Mehlich. DMS ®: program transformations for practical scalable software evolution. In *Software Engineering, 2004. ICSE 2004. Proceedings. 26th International Conference on*, pages 625–634, May 2004. doi: 10.1109/ICSE.2004.1317484.
- Marcin Benke, Peter Dybjer, and Patrik Jansson. Universes for generic programs and proofs in dependent type theory. 10(4):265–289, 2003. ISSN 1236-6064.
- Jon Louis Bentley. Little languages. *Commun. ACM*, 29(8):711–721, 1986. doi: 10.1145/6424.315691. URL <http://doi.acm.org/10.1145/6424.315691>.
- Véronique Benzaken, Giuseppe Castagna, Kim Nguyen, and Jérôme Siméon. Static and dynamic semantics of nosql languages. In *The 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '13, Rome, Italy - January 23 - 25, 2013*, pages 101–114, 2013. doi: 10.1145/2429069.2429083. URL <http://doi.acm.org/10.1145/2429069.2429083>.
- Armin Biere. Yet another local search solver and lingeling and friends entering the sat competition 2014. In *SAT Competition 2014, Vienna, Austria, July 14-17, Proceedings*, page 2, 2014.
- G.M. Bierman, M.J. Parkinson, and A.M. Pitts. An imperative core calculus for Java and Java with effects. Technical Report 563, University of Cambridge Computer Laboratory, April 2003.
- D Bisztray and Reiko Heckel. Combining Termination Criteria by Isolating Deletion. *Graph Transformations*, pages 203–217, 2010. URL http://link.springer.com/chapter/10.1007/978-3-642-15928-2_14.

- Sandrine Blazy and Xavier Leroy. Mechanized semantics for the Clight subset of the C language. *J. Autom. Reasoning*, 43(3):263–288, 2009. doi: 10.1007/s10817-009-9148-3. URL <http://dx.doi.org/10.1007/s10817-009-9148-3>.
- Jakob Blumer, Rubino Geiß, and Edgar Jakumeit. *The GrGen.NET User Manual*, oct 2010. URL <https://pp.info.uni-karlsruhe.de/svn/grgen-public/trunk/grgen/doc/grgen.pdf>.
- Martin Bodin, Thomas Jensen, and Alan Schmitt. Certified abstract interpretation with pretty-big-step semantics. In *Proceedings of the 2015 Conference on Certified Programs and Proofs, CPP 2015, Mumbai, India, January 15-17, 2015*, pages 29–40, 2015. doi: 10.1145/2676724.2693174. URL <http://doi.acm.org/10.1145/2676724.2693174>.
- Aaron Bohannon, J. Nathan Foster, Benjamin C. Pierce, Alexandre Pilkiewicz, and Alan Schmitt. Boomerang: resourceful lenses for string data. In *Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2008, San Francisco, California, USA, January 7-12, 2008*, pages 407–419, 2008. doi: 10.1145/1328438.1328487. URL <http://doi.acm.org/10.1145/1328438.1328487>.
- Egon Börger and Robert F Stärk. *Abstract State Machines: A Method for High-level System Design and Analysis*. Springer, 2003.
- Peter Borovanský, Claude Kirchner, H el ene Kirchner, Pierre-Etienne Moreau, and Christophe Ringeissen. An overview of ELAN. *Electr. Notes Theor. Comput. Sci.*, 15:55–70, 1998. doi: 10.1016/S1571-0661(05)82552-6. URL [https://doi.org/10.1016/S1571-0661\(05\)82552-6](https://doi.org/10.1016/S1571-0661(05)82552-6).
- Ahmed Bouajjani, Cezara Dragoi, Constantin Enea, and Mihaela Sighireanu. Abstract domains for automated reasoning about list-manipulating programs with infinite data. In *Verification, Model Checking, and Abstract Interpretation - 13th International Conference, VMCAI 2012, Philadelphia, PA, USA, January 22-24, 2012. Proceedings*, pages 1–22, 2012. doi: 10.1007/978-3-642-27940-9_1. URL https://doi.org/10.1007/978-3-642-27940-9_1.
- Peter Braun and Frank Marschall. BOTL - the bidirectional object oriented transformation language. Technical report, Technische Universit at M unchen, 2003.

- Martin Bravenboer, Karl Trygve Kalleberg, Rob Vermaas, and Eelco Visser. Stratego/xt 0.17. A language and toolset for program transformation. *Sci. Comput. Program.*, 72(1-2):52–70, 2008. doi: 10.1016/j.scico.2007.11.003. URL <https://doi.org/10.1016/j.scico.2007.11.003>.
- Stefan Bunzel. AUTOSAR - the standardized software architecture. *Informatik Spektrum*, 34(1):79–83, 2011. doi: 10.1007/s00287-010-0506-7. URL <http://dx.doi.org/10.1007/s00287-010-0506-7>.
- Fabian Büttner, Marina Egea, and Jordi Cabot. On verifying ATL transformations using 'off-the-shelf' SMT solvers. In Robert B. France, Jürgen Kazmeier, Ruth Breu, and Colin Atkinson, editors, *Model Driven Engineering Languages and Systems - 15th International Conference, MODELS 2012, Innsbruck, Austria, September 30-October 5, 2012. Proceedings*, volume 7590 of *Lecture Notes in Computer Science*, pages 432–448. Springer, 2012. ISBN 978-3-642-33665-2. doi: 10.1007/978-3-642-33666-9_28. URL http://dx.doi.org/10.1007/978-3-642-33666-9_28.
- Cristian Cadar and Alastair F. Donaldson. Analysing the program analyser. In *Proceedings of the 38th International Conference on Software Engineering, ICSE 2016, Austin, TX, USA, May 14-22, 2016 - Companion Volume*, pages 765–768, 2016. doi: 10.1145/2889160.2889206. URL <http://doi.acm.org/10.1145/2889160.2889206>.
- Cristian Cadar, Daniel Dunbar, and Dawson R. Engler. KLEE: unassisted and automatic generation of high-coverage tests for complex systems programs. In Richard Draves and Robbert van Renesse, editors, *8th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2008, December 8-10, 2008, San Diego, California, USA, Proceedings*, pages 209–224. USENIX Association, 2008. ISBN 978-1-931971-65-2. URL http://www.usenix.org/events/osdi08/tech/full_papers/cadar/cadar.pdf.
- Cristiano Calcagno, Dino Distefano, Peter W. O’Hearn, and Hongseok Yang. Compositional shape analysis by means of bi-abduction. *J. ACM*, 58(6):26:1–26:66, 2011. doi: 10.1145/2049697.2049700. URL <http://doi.acm.org/10.1145/2049697.2049700>.
- Daniel Calegari and Nora Szasz. Verification of model transformations: A survey of the state-of-the-art. *Electr. Notes Theor. Comput. Sci.*, 292: 5–25, 2013. doi: 10.1016/j.entcs.2013.02.002. URL <https://doi.org/10.1016/j.entcs.2013.02.002>.

- Giuseppe Castagna and Kim Nguyen. Typed iterators for XML. In *Proceeding of the 13th ACM SIGPLAN international conference on Functional programming, ICFP 2008, Victoria, BC, Canada, September 20-28, 2008*, pages 15–26, 2008. doi: 10.1145/1411204.1411210. URL <http://doi.acm.org/10.1145/1411204.1411210>.
- María Victoria Cengarle and Alexander Knapp. Ocl 1.4/5 vs. 2.0 expressions formal semantics and expressiveness. *Software and Systems Modeling*, 3(1):9–30, 2004. ISSN 1619-1366. doi: 10.1007/s10270-003-0035-9. URL <http://dx.doi.org/10.1007/s10270-003-0035-9>.
- Bor-Yuh Evan Chang and Xavier Rival. Relational inductive shape analysis. In *Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2008, San Francisco, California, USA, January 7-12, 2008*, pages 247–260, 2008. doi: 10.1145/1328438.1328469. URL <http://doi.acm.org/10.1145/1328438.1328469>.
- James Chapman, Pierre-Évariste Dagand, Conor McBride, and Peter Morris. The gentle art of levitation. In *Proceeding of the 15th ACM SIGPLAN international conference on Functional programming, ICFP 2010, Baltimore, Maryland, USA, September 27-29, 2010*, pages 3–14, 2010. doi: 10.1145/1863543.1863547. URL <http://doi.acm.org/10.1145/1863543.1863547>.
- Magnus Christerson, David Frankel, and Todd Schiller. Financial domain-specific language listing, October 2013. URL <http://www.webcitation.org/query?url=http%3A%2F%2Fwww.dslfin.org%2Fresources.html&date=2017-07-24>.
- Horatiu Cirstea and Claude Kirchner. The rewriting calculus - part I. *Logic Journal of the IGPL*, 9(3):339–375, 2001a. doi: 10.1093/jigpal/9.3.339. URL <https://doi.org/10.1093/jigpal/9.3.339>.
- Horatiu Cirstea and Claude Kirchner. The rewriting calculus - part II. *Logic Journal of the IGPL*, 9(3):377–410, 2001b. doi: 10.1093/jigpal/9.3.377. URL <https://doi.org/10.1093/jigpal/9.3.377>.
- Koen Claessen and John Hughes. Quickcheck: a lightweight tool for random testing of haskell programs. In *Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming (ICFP '00), Montreal, Canada, September 18-21, 2000.*, pages 268–279, 2000. doi:

10.1145/351240.351266. URL <http://doi.acm.org/10.1145/351240.351266>.

Robert Clarisó, Jordi Cabot, Esther Guerra, and Juan de Lara. Backwards reasoning for model transformations: Method and applications. *Journal of Systems and Software*, 116:113–132, 2016. doi: 10.1016/j.jss.2015.08.017. URL <https://doi.org/10.1016/j.jss.2015.08.017>.

Manuel Clavel, Francisco Durán, Steven Eker, Patrick Lincoln, Narciso Martí-Oliet, José Meseguer, and Carolyn L. Talcott, editors. *All About Maude - A High-Performance Logical Framework, How to Specify, Program and Verify Systems in Rewriting Logic*, volume 4350 of *Lecture Notes in Computer Science*, 2007. Springer. ISBN 978-3-540-71940-3. doi: 10.1007/978-3-540-71999-1. URL <https://doi.org/10.1007/978-3-540-71999-1>.

B. Jack Copeland. The church-turing thesis. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2015 edition, 2015. URL <https://plato.stanford.edu/archives/sum2015/entries/church-turing/>.

James R. Cordy. The TXL source transformation language. *Sci. Comput. Program.*, 61(3):190–210, 2006. doi: 10.1016/j.scico.2006.04.002. URL <http://dx.doi.org/10.1016/j.scico.2006.04.002>.

Ricardo Corin and Felipe Andrés Manzano. Efficient symbolic execution for analysing cryptographic protocol implementations. In Úlfar Erlingsson, Roel Wieringa, and Nicola Zannone, editors, *Engineering Secure Software and Systems - Third International Symposium, ESSoS 2011, Madrid, Spain, February 9-10, 2011. Proceedings*, volume 6542 of *Lecture Notes in Computer Science*, pages 58–72. Springer, 2011. ISBN 978-3-642-19124-4. doi: 10.1007/978-3-642-19125-1_5. URL http://dx.doi.org/10.1007/978-3-642-19125-1_5.

Patrick Cousot. Verification by abstract interpretation. In *Verification: Theory and Practice, Essays Dedicated to Zohar Manna on the Occasion of His 64th Birthday*, pages 243–268, 2003. doi: 10.1007/978-3-540-39910-0_11. URL https://doi.org/10.1007/978-3-540-39910-0_11.

Patrick Cousot and Radhia Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Conference Record of the Fourth ACM Symposium*

on *Principles of Programming Languages*, Los Angeles, California, USA, January 1977, pages 238–252, 1977. doi: 10.1145/512950.512973. URL <http://doi.acm.org/10.1145/512950.512973>.

Patrick Cousot and Radhia Cousot. Formal language, grammar and set-constraint-based program analysis by abstract interpretation. In *Proceedings of the seventh international conference on Functional programming languages and computer architecture, FPCA 1995, La Jolla, California, USA, June 25-28, 1995*, pages 170–181, 1995. URL <http://doi.acm.org/10.1145/224164.224199>.

Patrick Cousot and Radhia Cousot. Modular static program analysis. In *Compiler Construction, 11th International Conference, CC 2002, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2002, Grenoble, France, April 8-12, 2002, Proceedings*, pages 159–178, 2002. doi: 10.1007/3-540-45937-5_13. URL https://doi.org/10.1007/3-540-45937-5_13.

Arlen Cox, Bor-Yuh Evan Chang, Huisong Li, and Xavier Rival. Abstract domains and solvers for sets reasoning. In *Logic for Programming, Artificial Intelligence, and Reasoning - 20th International Conference, LPAR-20 2015, Suva, Fiji, November 24-28, 2015, Proceedings*, pages 356–371, 2015. doi: 10.1007/978-3-662-48899-7_25. URL https://doi.org/10.1007/978-3-662-48899-7_25.

J. Sanchez Cuadrado, E. Guerra, and J. de Lara. Static analysis of model transformations. *IEEE Transactions on Software Engineering*, PP(99):1–1, 2017. ISSN 0098-5589. doi: 10.1109/TSE.2016.2635137.

CWI Amsterdam. Abstract syntax tree of muRascal. online, February 2017. URL <https://web.archive.org/web/20170726155735/https://raw.githubusercontent.com/usetheSource/rascal/master/src/org/rascalmpl/library/experiments/Compiler/muRascal/AST.rsc>.

Krzysztof Czarnecki and Simon Helsen. Feature-based survey of model transformation approaches. *IBM Systems Journal*, 45(3):621–646, 2006. doi: 10.1147/sj.453.0621. URL <https://doi.org/10.1147/sj.453.0621>.

Juan de Lara, Hans Vangheluwe, and Manuel Alfonseca. Meta-modelling and graph grammars for multi-paradigm modelling in

- atom3. *Software and Systems Modeling*, 3(3):194–209, 2004. ISSN 1619-1366. doi: 10.1007/s10270-003-0047-5. URL <http://dx.doi.org/10.1007/s10270-003-0047-5>.
- Xianghua Deng, Jooyong Lee, and Robby. Efficient and formal generalized symbolic execution. *Autom. Softw. Eng.*, 19(3):233–301, 2012. doi: 10.1007/s10515-011-0089-9. URL <http://dx.doi.org/10.1007/s10515-011-0089-9>.
- Nachum Dershowitz and Zohar Manna. Proving termination with multiset orderings. *Commun. ACM*, 22(8):465–476, 1979. doi: 10.1145/359138.359142. URL <http://doi.acm.org/10.1145/359138.359142>.
- Aleksandar S. Dimovski, Ahmad Salim Al-Sibahi, Claus Brabrand, and Andrzej Wasowski. Family-based model checking without a family-based model checker. In *Model Checking Software - 22nd International Symposium, SPIN 2015, Stellenbosch, South Africa, August 24-26, 2015, Proceedings*, pages 282–299, 2015a. doi: 10.1007/978-3-319-23404-5_18.
- Aleksandar S. Dimovski, Ahmad Salim Al-Sibahi, Claus Brabrand, and Andrzej Wasowski. Family-based model checking using off-the-shelf model checkers: extended abstract. In *Proceedings of the 19th International Conference on Software Product Line, SPLC 2015, Nashville, TN, USA, July 20-24, 2015*, page 397, 2015b. doi: 10.1145/2791060.2791119. URL <http://doi.acm.org/10.1145/2791060.2791119>.
- Aleksandar S. Dimovski, Ahmad Salim Al-Sibahi, Claus Brabrand, and Andrzej Wasowski. Efficient family-based model checking via variability abstractions. *International Journal on Software Tools for Technology Transfer*, pages 1–19, 2016. ISSN 1433-2787. doi: 10.1007/s10009-016-0425-2. URL <http://dx.doi.org/10.1007/s10009-016-0425-2>.
- Keith Duddy, Michael Lawley, and Sridhar Iyengar. Mof query/views/-transformations: Second revised submission. Technical report, 2004. URL <http://tefkat.sourceforge.net/publications/ad-04-01-06.pdf>.
- Levent Erkök and John Matthews. High assurance programming in cryptol. In *Fifth Cyber Security and Information Intelligence Research Workshop, CSIIIRW '09, Knoxville, TN, USA, April 13-15, 2009*, page 60, 2009. doi: 10.1145/1558607.1558676. URL <http://doi.acm.org/10.1145/1558607.1558676>.

- Antonio Estévez, Javier Padrón, E. Victor Sánchez Rebull, and José L. Roda. ATC: A low-level model transformation language. In *Model-Driven Enterprise Information Systems, Proceedings of the 2nd International Workshop on Model-Driven Enterprise Information Systems, MDEIS 2006, In conjunction with ICEIS 2006, Paphos, Cyprus, May 2006*, pages 64–74, 2006.
- Alexander Felfernig, Lothar Hotz, Claire Bagley, and Juha Tiihonen. *Knowledge-based Configuration: From Research to Business Cases*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1 edition, 2014. ISBN 012415817X, 9780124158177.
- Matthias Felleisen, Robert Bruce Findler, and Matthew Flatt. *Semantics Engineering with PLT Redex*. MIT Press, 2009. ISBN 978-0-262-06275-6. URL <http://mitpress.mit.edu/catalog/item/default.asp?ttype=2&tid=11885>.
- Matthias Felleisen, Robert Bruce Findler, Matthew Flatt, Shriram Krishnamurthi, Eli Barzilay, Jay A. McCarthy, and Sam Tobin-Hochstadt. The racket manifesto. In *1st Summit on Advances in Programming Languages, SNAPL 2015, May 3-6, 2015, Asilomar, California, USA*, pages 113–128, 2015. doi: 10.4230/LIPIcs.SNAPL.2015.113. URL <https://doi.org/10.4230/LIPIcs.SNAPL.2015.113>.
- Olivier Finot, Jean-Marie Mottu, Gerson Sunyé, and Thomas Degueule. Using meta-model coverage to qualify test oracles. In *Proceedings of the Second Workshop on the Analysis of Model Transformations (AMT 2013), Miami, FL, USA, September 29, 2013*, 2013. URL http://ceur-ws.org/Vol-1077/amt13_submission_3.pdf.
- Thorsten Fischer, Jörg Niere, Lars Torunski, and Albert Zündorf. Story diagrams: A new graph rewrite language based on the unified modeling language and java. In Hartmut Ehrig, Gregor Engels, Hans-Jörg Kreowski, and Grzegorz Rozenberg, editors, *Theory and Application of Graph Transformations*, volume 1764 of *Lecture Notes in Computer Science*, pages 296–309. Springer Berlin Heidelberg, 2000. ISBN 978-3-540-67203-6. doi: 10.1007/978-3-540-46464-8_21. URL http://dx.doi.org/10.1007/978-3-540-46464-8_21.
- Martin Fowler. *Refactoring - Improving the Design of Existing Code*. Addison Wesley object technology series. Addison-Wesley, 1999. ISBN 978-

0-201-48567-7. URL <http://martinfowler.com/books/refactoring.html>.

Martin Fowler. *Domain-Specific Languages*. The Addison-Wesley signature series. Addison-Wesley, 2011. ISBN 978-0-321-71294-3. URL http://vig.pearsoned.com/store/product/1,1207,store-12521_isbn-0321712943,00.html.

Timothy S. Freeman and Frank Pfenning. Refinement types for ML. In *Proceedings of the ACM SIGPLAN'91 Conference on Programming Language Design and Implementation (PLDI), Toronto, Ontario, Canada, June 26-28, 1991*, pages 268–277, 1991. URL <http://doi.acm.org/10.1145/113445.113468>.

Patrice Godefroid, Nils Klarlund, and Koushik Sen. DART: directed automated random testing. In *Proceedings of the ACM SIGPLAN 2005 Conference on Programming Language Design and Implementation, Chicago, IL, USA, June 12-15, 2005*, pages 213–223, 2005. doi: 10.1145/1065010.1065036. URL <http://doi.acm.org/10.1145/1065010.1065036>.

C Gomes, Bruno Barroca, and Vasco Amaral. Classification of Model Transformation Tools: Pattern Matching Techniques. *Model-Driven Engineering Languages and Systems*, pages 619–635, 2014. URL http://link.springer.com/chapter/10.1007/978-3-319-11653-2_38.

Carlos A. González and Jordi Cabot. Atltest: A white-box test generation approach for ATL transformations. In *Model Driven Engineering Languages and Systems - 15th International Conference, MODELS 2012, Innsbruck, Austria, September 30-October 5, 2012. Proceedings*, pages 449–464, 2012. doi: 10.1007/978-3-642-33666-9_29. URL http://dx.doi.org/10.1007/978-3-642-33666-9_29.

James Gosling, Bill Joy, Guy L. Steele Jr., Gilad Bracha, and Alex Buckley. *The Java Language Specification*. Addison-Wesley, 8th edition, 2014. ISBN 978-0133900699.

Jim Gray. Why do computers stop and what can be done about it? In *Fifth Symposium on Reliability in Distributed Software and Database Systems, SRDS 1986, Los Angeles, California, USA, January 13-15, 1986, Proceedings*, pages 3–12. IEEE Computer Society, 1986. ISBN 0-8186-0690-8.

- Yuri Gurevich. Evolving algebras 1993: Lipari guide. *Specification and validation methods*, pages 9–36, 1995.
- Nicolas Halbwachs and Mathias Péron. Discovering properties about arrays in simple programs. In *Proceedings of the ACM SIGPLAN 2008 Conference on Programming Language Design and Implementation, Tucson, AZ, USA, June 7-13, 2008*, pages 339–348, 2008. doi: 10.1145/1375581.1375623. URL <http://doi.acm.org/10.1145/1375581.1375623>.
- John Harrison. *Handbook of Practical Logic and Automated Reasoning*. Cambridge University Press, 2009. ISBN 978-0-521-89957-4.
- Anders Hejlsberg, Mads Torgersen, Scott Wiltamuth, and Peter Golde. *The C# Programming Language*. Addison-Wesley, 4th edition, 2011. ISBN 978-0321741769.
- Soichiro Hidaka, Massimo Tisi, Jordi Cabot, and Zhenjiang Hu. Feature-based classification of bidirectional transformation approaches. *Software and System Modeling*, 15(3):907–928, 2016. doi: 10.1007/s10270-014-0450-0. URL <https://doi.org/10.1007/s10270-014-0450-0>.
- C. A. R. Hoare. An axiomatic basis for computer programming. *Commun. ACM*, 12(10):576–580, 1969. doi: 10.1145/363235.363259. URL <http://doi.acm.org/10.1145/363235.363259>.
- C. A. R. Hoare. The verifying compiler, a grand challenge for computing research. In *Verification, Model Checking, and Abstract Interpretation, 6th International Conference, VMCAI 2005, Paris, France, January 17-19, 2005, Proceedings*, pages 78–78, 2005. doi: 10.1007/978-3-540-30579-8_5. URL http://dx.doi.org/10.1007/978-3-540-30579-8_5.
- J.E. Hopcroft, R. Motwani, and J.D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Pearson/Addison Wesley, 2007. ISBN 9780321455369. URL <http://books.google.dk/books?id=6ytHAQAAIAAJ>.
- Zhenjiang Hu, Shin-Cheng Mu, and Masato Takeichi. A programmable editor for developing structured documents based on bidirectional transformations. *Higher-Order and Symbolic Computation*, 21(1-2):89–118, 2008. doi: 10.1007/s10990-008-9025-5. URL <https://doi.org/10.1007/s10990-008-9025-5>.

- Paul Hudak. Building domain-specific embedded languages. *ACM Comput. Surv.*, 28(4es):196, 1996. doi: 10.1145/242224.242477. URL <http://doi.acm.org/10.1145/242224.242477>.
- Laura Inozemtseva and Reid Holmes. Coverage is not strongly correlated with test suite effectiveness. In *36th International Conference on Software Engineering, ICSE '14, Hyderabad, India - May 31 - June 07, 2014*, pages 435–445, 2014. doi: 10.1145/2568225.2568271. URL <http://doi.acm.org/10.1145/2568225.2568271>.
- Alexandru Florin Iosif-Lazăr, Ahmad Salim Al-Sibahi, Aleksandar S. Dimovski, Juha Erik Savolainen, Krzysztof Sierszecki, and Andrzej Wasowski. Experiences from designing and validating a software modernization transformation (E). In *30th IEEE/ACM International Conference on Automated Software Engineering, ASE 2015, Lincoln, NE, USA*, pages 597–607, 2015. doi: 10.1109/ASE.2015.84.
- Ethan K. Jackson, Tihamer Levendovszky, and Daniel Balasubramanian. Reasoning about metamodeling with formal specifications and automatic proofs. In *Model Driven Engineering Languages and Systems, 14th International Conference, MODELS 2011, Wellington, New Zealand, October 16-21, 2011. Proceedings*, pages 653–667, 2011. doi: 10.1007/978-3-642-24485-8_48. URL https://doi.org/10.1007/978-3-642-24485-8_48.
- Edgar Jakumeit. Erweiterung der regelsprache eines graphersetzungswerkzeugs um rekursive regeln mittels sterngraphgrammatiken und paragraphgrammatiken. Diploma thesis, Universität Karlsruhe, jul 2008.
- Edgar Jakumeit, Sebastian Buchwald, and Moritz Kroll. Grgen.net. *International Journal on Software Tools for Technology Transfer*, 12(3/4):263 – 271, 2010. ISSN 14332779. URL <http://search.ebscohost.com/login.aspx?direct=true&db=aph&AN=51523140&site=ehost-live>.
- Neil D. Jones, Carsten K. Gomard, and Peter Sestoft. *Partial evaluation and automatic program generation*. Prentice Hall international series in computer science. Prentice Hall, 1993. ISBN 978-0-13-020249-9.
- Frédéric Jouault, Freddy Allilaire, Jean Bézivin, and Ivan Kurtev. ATL: A model transformation tool. *Sci. Comput. Program.*, 72(1-2):31–39, 2008.

doi: 10.1016/j.scico.2007.08.002. URL <http://dx.doi.org/10.1016/j.scico.2007.08.002>.

Gilles Kahn. Natural semantics. In *STACS 87, 4th Annual Symposium on Theoretical Aspects of Computer Science, Passau, Germany, February 19-21, 1987, Proceedings*, pages 22–39, 1987. doi: 10.1007/BFb0039592. URL <https://doi.org/10.1007/BFb0039592>.

Lennart C. L. Kats and Eelco Visser. The spoofax language workbench: rules for declarative specification of languages and ides. In William R. Cook, Siobhán Clarke, and Martin C. Rinard, editors, *Proceedings of the 25th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2010, October 17-21, 2010, Reno/Tahoe, Nevada, USA*, pages 444–463. ACM, 2010. ISBN 978-1-4503-0203-6. doi: 10.1145/1869459.1869497. URL <http://doi.acm.org/10.1145/1869459.1869497>.

Michael Kay. XSL transformations (XSLT) version 3.0. W3C recommendation, W3C, June 2017. <https://www.w3.org/TR/2017/REC-xslt-30-20170608/>.

Sarfraz Khurshid, Corina S. Pasareanu, and Willem Visser. Generalized symbolic execution for model checking and testing. In *Tools and Algorithms for the Construction and Analysis of Systems, 9th International Conference, TACAS 2003, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2003, Warsaw, Poland, April 7-11, 2003, Proceedings*, pages 553–568, 2003. doi: 10.1007/3-540-36577-X_40. URL http://dx.doi.org/10.1007/3-540-36577-X_40.

James C. King. Symbolic execution and program testing. *Commun. ACM*, 19(7):385–394, 1976. doi: 10.1145/360248.360252. URL <http://doi.acm.org/10.1145/360248.360252>.

Stephen Cole Kleene. *Mathematical Logic*. John Wiley & Sons, 1 edition, 1967. ISBN 978-0471490333.

Paul Klint, Tijs van der Storm, and Jurgen J. Vinju. RASCAL: A domain specific language for source code analysis and manipulation. In *Ninth IEEE International Working Conference on Source Code Analysis and Manipulation, SCAM 2009, Edmonton, Alberta, Canada, September 20-21, 2009*, pages 168–177, 2009. doi: 10.1109/SCAM.2009.28. URL <https://doi.org/10.1109/SCAM.2009.28>.

- Paul Klint, Tijs van der Storm, and Jurgen Vinju. EASY Meta-programming with Rascal. In JoãoM. Fernandes, Ralf Lämmel, Joost Visser, and João Saraiva, editors, *Generative and Transformational Techniques in Software Engineering III*, volume 6491 of *Lecture Notes in Computer Science*, pages 222–289. Springer Berlin Heidelberg, 2011. ISBN 978-3-642-18022-4. doi: 10.1007/978-3-642-18023-1_6. URL http://dx.doi.org/10.1007/978-3-642-18023-1_6.
- Dimitrios S. Kolovos, Richard F. Paige, and Fiona Polack. The epsilon transformation language. In *Theory and Practice of Model Transformations, First International Conference, ICMT 2008, Zürich, Switzerland, July 1-2, 2008, Proceedings*, pages 46–60, 2008. doi: 10.1007/978-3-540-69927-9_4. URL https://doi.org/10.1007/978-3-540-69927-9_4.
- Dimitris Kolovos, Louis Rose, Antonio García-Domínguez, and Richard Paige. The epsilon book. September 2014. URL <http://eclipse.org/epsilon/doc/book/>.
- B. Korel. A dynamic approach of test data generation. In *Software Maintenance, 1990, Proceedings., Conference on*, pages 311–317, Nov 1990. doi: 10.1109/ICSM.1990.131379.
- Alexander Kraas. Realizing model simplifications with QVT operational mappings. In *Proceedings of the 14th International Workshop on OCL and Textual Modelling co-located with 17th International Conference on Model Driven Engineering Languages and Systems (MODELS 2014), Valencia, Spain, September 30, 2014.*, pages 53–62, 2014. URL <http://ceur-ws.org/Vol-1285/paper06.pdf>.
- Daniel Kröning, Philipp Rümmer, and Georg Weissenbacher. A proposal for a theory of finite sets, lists, and maps for the smt-lib standard. In *Informal proceedings, 7th International Workshop on Satisfiability Modulo Theories at CADE. Vol. 22. 2009*, 2009.
- Peter J Landin. The mechanical evaluation of expressions. *The Computer Journal*, 6(4):308–320, 1964.
- Chris Lattner and Vikram S. Adve. LLVM: A compilation framework for lifelong program analysis & transformation. In *2nd IEEE / ACM International Symposium on Code Generation and Optimization (CGO 2004), 20-24 March 2004, San Jose, CA, USA*, pages 75–88. IEEE Computer Society, 2004. ISBN 0-7695-2102-9. doi: 10.1109/CGO.2004.1281665. URL <http://dx.doi.org/10.1109/CGO.2004.1281665>.

- Marius Lauder, Anthony Anjorin, Gergely Varró, and Andy Schürr. Bidirectional model transformation with precedence triple graph grammars. In *Modelling Foundations and Applications - 8th European Conference, ECMFA 2012, Kgs. Lyngby, Denmark, July 2-5, 2012. Proceedings*, pages 287–302, 2012. doi: 10.1007/978-3-642-31491-9_22. URL https://doi.org/10.1007/978-3-642-31491-9_22.
- Michael Lawley and Jim Steel. Practical declarative model transformation with tefkat. In Jean-Michel Bruel, editor, *Satellite Events at the MoDELS 2005 Conference*, volume 3844 of *Lecture Notes in Computer Science*, pages 139–150. Springer Berlin Heidelberg, 2006. ISBN 978-3-540-31780-7. doi: 10.1007/11663430_15. URL http://dx.doi.org/10.1007/11663430_15.
- László Lengyel, Tihamér Levendovszky, and Hassan Charaf. A visual control flow language and its termination properties. *International Journal of Computer, Information, Systems and Control Engineering*, 1(8):2505 – 2510, 2007. ISSN 1307-6892. URL <http://waset.org/Publications?p=8>.
- Xavier Leroy. The Caml Light system, documentation and user’s guide 0.74. <https://caml.inria.fr/pub/docs/manual-caml-light/>, December 1997.
- Tihamér Levendovszky, László Lengyel, Gergely Mezei, and Hassan Charaf. A systematic approach to metamodeling environments and model transformation systems in VMTS. *Electronic Notes in Theoretical Computer Science*, 127(1):65 – 75, 2005. ISSN 1571-0661. doi: <http://dx.doi.org/10.1016/j.entcs.2004.12.040>. URL <http://www.sciencedirect.com/science/article/pii/S1571066105001155>. Proceedings of the International Workshop on Graph-Based Tools (GraBaTs 2004) Graph-Based Tools 2004.
- Andriy Levytsky and Eugene J. H. Kerckhoffs. From class diagrams to zope products with the meta-modelling tool atom. *Summer Computer Simulation Conference*, pages 295–300, 2003.
- Jiangchao Liu and Xavier Rival. An array content static analysis based on non-contiguous partitions. *Computer Languages, Systems & Structures*, 47:104–129, 2017. doi: 10.1016/j.cl.2016.01.005. URL <https://doi.org/10.1016/j.cl.2016.01.005>.

- Benjamin Livshits, Manu Sridharan, Yannis Smaragdakis, Ondrej Lhoták, José Nelson Amaral, Bor-Yuh Evan Chang, Samuel Z. Guyer, Uday P. Khedker, Anders Møller, and Dimitrios Vardoulakis. In defense of soundness: a manifesto. *Commun. ACM*, 58(2):44–46, 2015. doi: 10.1145/2644805. URL <http://doi.acm.org/10.1145/2644805>.
- Andres Löb and José Pedro Magalhães. Generic programming with indexed functors. In *Proceedings of the seventh ACM SIGPLAN workshop on Generic programming, WGP@ICFP 2011, Tokyo, Japan, September 19–21, 2011*, pages 1–12, 2011. doi: 10.1145/2036918.2036920. URL <http://doi.acm.org/10.1145/2036918.2036920>.
- Levi Lucio and Hans Vangheluwe. Symbolic Execution for the Verification of Model Transformations. *VOLT@STAF*, pages 1–29, April 2013.
- Simon Marlow. *The Haskell 2010 Language Report*. 2010. URL <https://www.haskell.org/definition/haskell2010.pdf>.
- Tom Mens and P Van Gorp. A taxonomy of model transformation. *Electronic Notes in Theoretical Computer Science*, 152(1-2):125–142, 2006. URL <http://www.sciencedirect.com/science/article/pii/S1571066106001435>.
- Marjan Mernik, Jan Heering, and Anthony M. Sloane. When and how to develop domain-specific languages. *ACM Comput. Surv.*, 37(4):316–344, 2005. doi: 10.1145/1118890.1118892. URL <http://doi.acm.org/10.1145/1118890.1118892>.
- Robin Milner, Mads Tofte, and Robert Harper. *Definition of standard ML*. MIT Press, 1990. ISBN 978-0-262-63132-7.
- Neil Mitchell and Colin Runciman. Uniform boilerplate and list processing. In *Proceedings of the ACM SIGPLAN Workshop on Haskell, Haskell 2007, Freiburg, Germany, September 30, 2007*, pages 49–60, 2007. doi: 10.1145/1291201.1291208. URL <http://doi.acm.org/10.1145/1291201.1291208>.
- Melina Mongiovi, Gustavo Mendes, Rohit Gheyi, Gustavo Soares, and Márcio Ribeiro. Scaling testing of refactoring engines. In *30th IEEE International Conference on Software Maintenance and Evolution, Victoria, BC, Canada, September 29 - October 3, 2014*, pages 371–380, 2014. doi: 10.1109/ICSME.2014.59. URL <http://dx.doi.org/10.1109/ICSME.2014.59>.

- Alan Mycroft and Neil D. Jones. A relational framework for abstract interpretation. In *Programs as Data Objects, Proceedings of a Workshop, Copenhagen, Denmark, October 17-19, 1985*, pages 156–171, 1985. doi: 10.1007/3-540-16446-4_9. URL https://doi.org/10.1007/3-540-16446-4_9.
- Arne Nordmann, Nico Hochgeschwender, and Sebastian Wrede. A Survey on Domain-Specific Languages in Robotics. *Journal of Software Engineering for Robotics*, 8810(Chapter 17):195–206, October 2014.
- Object Management Group. *Meta Object Facility (MOF) 2.0 Query/View/Transformation Specification*. June 2016. version 1.3.
- Martin Odersky and Tiark Rompf. Unifying functional and object-oriented programming with scala. *Commun. ACM*, 57(4):76–86, 2014. doi: 10.1145/2591013. URL <http://doi.acm.org/10.1145/2591013>.
- Martin Odersky, Philippe Altherr, Vincent Cremet, Burak Emir, Sebastian Maneth, Stéphane Micheloud, Nikolay Mihaylov, Michel Schinz, Erik Stenman, and Matthias Zenger. An overview of the scala programming language. Technical report, 2004.
- Ruhsan Onder and Zeki Bayram. Xslt version 2.0 is turing-complete: A purely transformation based proof. In *Proceedings of the 11th International Conference on Implementation and Application of Automata, CIAA'06*, pages 275–276, Berlin, Heidelberg, 2006. Springer-Verlag. ISBN 3-540-37213-X, 978-3-540-37213-4. doi: 10.1007/11812128_26. URL http://dx.doi.org/10.1007/11812128_26.
- Oded Padon, Neil Immerman, Sharon Shoham, Aleksandr Karbyshev, and Mooly Sagiv. Decidability of inferring inductive invariants. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, pages 217–231, 2016. doi: 10.1145/2837614.2837640. URL <http://doi.acm.org/10.1145/2837614.2837640>.
- Corina S. Pasareanu, Willem Visser, David H. Bushnell, Jaco Geldenhuys, Peter C. Mehlitz, and Neha Rungta. Symbolic pathfinder: integrating symbolic execution with model checking for java bytecode analysis. *Autom. Softw. Eng.*, 20(3):391–425, 2013. doi: 10.1007/s10515-013-0122-2. URL <http://dx.doi.org/10.1007/s10515-013-0122-2>.

- S Patzina and Lars Patzina. A case study based comparison of atl and sdm. *Proceedings of the 4th International Conference on Applications of Graph Transformations with Industrial Relevance*, 7233, 2012. doi: 10.1007/978-3-642-34176-2. URL http://link.springer.com/chapter/10.1007/978-3-642-34176-2_18<http://www.springerlink.com/index/10.1007/978-3-642-34176-2>.
- Valentin Perrelle and Nicolas Halbwachs. An analysis of permutations in arrays. In *Verification, Model Checking, and Abstract Interpretation, 11th International Conference, VMCAI 2010, Madrid, Spain, January 17-19, 2010. Proceedings*, pages 279–294, 2010. doi: 10.1007/978-3-642-11319-2_21. URL https://doi.org/10.1007/978-3-642-11319-2_21.
- Tuan-Hung Pham and Michael W. Whalen. An improved unrolling-based decision procedure for algebraic data types. In *Verified Software: Theories, Tools, Experiments - 5th International Conference, VSTTE 2013, Menlo Park, CA, USA, May 17-19, 2013, Revised Selected Papers*, pages 129–148, 2013. doi: 10.1007/978-3-642-54108-7_7. URL https://doi.org/10.1007/978-3-642-54108-7_7.
- Gordon D. Plotkin. A structural approach to operational semantics. *J. Log. Algebr. Program.*, 60-61:17–139, 2004.
- Detlef Plump. On Termination of Graph Rewriting. *Graph-Theoretic Concepts in Computer Science*, 1995. URL http://link.springer.com/chapter/10.1007/3-540-60618-1_68.
- Detlef Plump. Termination of graph rewriting is undecidable. *Fundamenta Informaticae*, 33(2), 1998.
- Amir Pnueli, Michael Siegel, and Eli Singerman. Translation validation. In Bernhard Steffen, editor, *Tools and Algorithms for Construction and Analysis of Systems, 4th International Conference, TACAS '98, Held as Part of the European Joint Conferences on the Theory and Practice of Software, ETAPS'98, Lisbon, Portugal, March 28 - April 4, 1998, Proceedings*, volume 1384 of *Lecture Notes in Computer Science*, pages 151–166. Springer, 1998. ISBN 3-540-64356-7. doi: 10.1007/BFb0054170. URL <http://dx.doi.org/10.1007/BFb0054170>.
- Emil L. Post. A variant of a recursively unsolvable problem. *Bulletin of the American Mathematical Society*, 52(4):264–269, April 1946. ISSN

- 0002-9904. doi: 10.1090/S0002-9904-1946-08555-9. URL <http://www.ams.org/journal-getitem?pii=S0002-9904-1946-08555-9>.
- Emil L. Post. Recursive unsolvability of a problem of Thue. *The Journal of Symbolic Logic*, 12(1):1–11, 1947. URL http://journals.cambridge.org/abstract_S0022481200076295.
- Lukman Ab. Rahim and Jon Whittle. A survey of approaches for verifying model transformations. *Software and System Modeling*, 14(2):1003–1028, 2015. doi: 10.1007/s10270-013-0358-0. URL <http://dx.doi.org/10.1007/s10270-013-0358-0>.
- Arend Rensink. The GROOVE simulator: A tool for state space generation. In *Applications of Graph Transformations with Industrial Relevance, Second International Workshop, AGTIVE 2003, Charlottesville, VA, USA, September 27 - October 1, 2003, Revised Selected and Invited Papers*, pages 479–485, 2003. doi: 10.1007/978-3-540-25959-6_40. URL https://doi.org/10.1007/978-3-540-25959-6_40.
- Henry Gordon Rice. Classes of recursively enumerable sets and their decision problems. *Transactions of the American Mathematical Society*, 74(2):358–366, 1953.
- Adrián Riesco. Test-case generation for maude functional modules. In *WADT 2010, Etelsen, Germany, July 1-4, 2010*, pages 287–301, 2010. doi: 10.1007/978-3-642-28412-0_18. URL http://dx.doi.org/10.1007/978-3-642-28412-0_18.
- Adrián Riesco. Using narrowing to test maude specifications. In *WRLA'12, Tallinn, Estonia, March 24-25, 2012*, pages 201–220, Berlin, Heidelberg, 2012. Springer-Verlag. ISBN 978-3-642-34004-8. doi: 10.1007/978-3-642-34005-5_11. URL http://dx.doi.org/10.1007/978-3-642-34005-5_11.
- Xavier Rival. Symbolic transfer function-based approaches to certified compilation. In *Proceedings of the 31st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2004, Venice, Italy, January 14-16, 2004*, pages 1–13, 2004. doi: 10.1145/964001.964002. URL <http://doi.acm.org/10.1145/964001.964002>.

- Xavier Rival, Antoine Toubhans, and Bor-Yuh Evan Chang. Construction of abstract domains for heterogeneous properties (position paper). In *Leveraging Applications of Formal Methods, Verification and Validation. Specialized Techniques and Applications - 6th International Symposium, ISoLA 2014, Imperial, Corfu, Greece, October 8-11, 2014, Proceedings, Part II*, pages 489–492, 2014. doi: 10.1007/978-3-662-45231-8_40. URL https://doi.org/10.1007/978-3-662-45231-8_40.
- Mads Rosendahl. Abstract interpretation as a programming language. In *Semantics, Abstract Interpretation, and Reasoning about Programs: Essays Dedicated to David A. Schmidt on the Occasion of his Sixtieth Birthday, Manhattan, Kansas, USA, 19-20th September 2013.*, pages 84–104, 2013. doi: 10.4204/EPTCS.129.7. URL <https://doi.org/10.4204/EPTCS.129.7>.
- Grzegorz Rozenberg, editor. *Handbook of Graph Grammars and Computing by Graph Transformation: Volume I. Foundations*. World Scientific Publishing Co., Inc., River Edge, NJ, USA, 1997. ISBN 98-102288-48.
- Olga Runge. *The AGG 1.5.0 Development Environment: The User Manual*, 2006. URL <http://user.cs.tu-berlin.de/~gragra/agg/AGG-ShortManual/AGG-ShortManual.html>.
- John M. Rushby, Sam Owre, and Natarajan Shankar. Subtypes for specifications: Predicate subtyping in PVS. *IEEE Trans. Software Eng.*, 24(9):709–720, 1998. doi: 10.1109/32.713327. URL <https://doi.org/10.1109/32.713327>.
- Hanan Samet. Compiler testing via symbolic interpretation. In *Proceedings of the 1976 Annual Conference, Houston, Texas, USA, October 20-22, 1976*, pages 492–497, 1976. doi: 10.1145/800191.805648. URL <http://doi.acm.org/10.1145/800191.805648>.
- Max Schäfer, Torbjörn Ekman, and Oege de Moor. Challenge proposal: verification of refactorings. In *Proceedings of the 3rd ACM Workshop Programming Languages meets Program Verification, PLPV 2009, Savannah, GA, USA, January 20, 2009*, pages 67–72, 2009. doi: 10.1145/1481848.1481859. URL <http://doi.acm.org/10.1145/1481848.1481859>.
- David A. Schmidt. Trace-based abstract interpretation of operational semantics. *Lisp and Symbolic Computation*, 10(3):237–271, 1998.

- Johannes Schönböck, Gerti Kappel, Manuel Wimmer, Angelika Kusel, Werner Retschitzegger, and Wieland Schwinger. Tetrabox - A generic white-box testing framework for model transformations. In *20th Asia-Pacific Software Engineering Conference, APSEC 2013, Ratchathewi, Bangkok, Thailand, December 2-5, 2013 - Volume 1*, pages 75–82, 2013. doi: 10.1109/APSEC.2013.21. URL <http://dx.doi.org/10.1109/APSEC.2013.21>.
- Andy Schürr. PROGRES, a visual language and environment for programming with graph rewriting systems. Technical report, RWTH Aachen, 1994.
- Dana S. Scott. Data types as lattices. *SIAM J. Comput.*, 5(3):522–587, 1976. URL <http://dx.doi.org/10.1137/0205037>.
- Robert C. Seacord, Daniel Plakosh, and Grace A. Lewis. *Modernizing Legacy Systems - Software Technologies, Engineering Processes, and Business Practices*. SEI series in software engineering. Addison-Wesley, 2003. ISBN 978-0-321-11884-4. URL <http://www.informit.com/store/modernizing-legacy-systems-software-technologies-engineering-9780321118844>.
- Gehan M. K. Selim, Fabian Büttner, James R. Cordy, Jürgen Dingel, and Shige Wang. Automated verification of model transformations in the automotive industry. In Ana Moreira, Bernhard Schätz, Jeff Gray, Antonio Vallecillo, and Peter J. Clarke, editors, *Model-Driven Engineering Languages and Systems - 16th International Conference, MODELS 2013, Miami, FL, USA, September 29 - October 4, 2013. Proceedings*, volume 8107 of *Lecture Notes in Computer Science*, pages 690–706. Springer, 2013. ISBN 978-3-642-41532-6. doi: 10.1007/978-3-642-41533-3_42. URL http://dx.doi.org/10.1007/978-3-642-41533-3_42.
- Gehan M. K. Selim, Levi Lucio, James R. Cordy, Jürgen Dingel, and Bentley J. Oakes. Specification and verification of graph-based model transformation properties. In Holger Giese and Barbara König, editors, *Graph Transformation - 7th International Conference, ICGT 2014, Held as Part of STAF 2014, York, UK, July 22-24, 2014. Proceedings*, volume 8571 of *Lecture Notes in Computer Science*, pages 113–129. Springer, 2014. ISBN 978-3-319-09107-5. doi: 10.1007/978-3-319-09108-2_8. URL http://dx.doi.org/10.1007/978-3-319-09108-2_8.

- Gehan M. K. Selim, Shige Wang, James R. Cordy, and Juergen Dingel. Model transformations for migrating legacy deployment models in the automotive industry. *Software and System Modeling*, 14(1):365–381, 2015. doi: 10.1007/s10270-013-0365-1. URL <http://dx.doi.org/10.1007/s10270-013-0365-1>.
- Jim Shore. Fail fast. *IEEE Software*, 21(5):21–25, 2004. doi: 10.1109/MS.2004.1331296. URL <http://doi.ieeecomputersociety.org/10.1109/MS.2004.1331296>.
- Jon Siegel, editor. *CORBA 3 fundamentals and programming*. OMG Press, John Wiley & Sons, New York, 2nd edition, 2000.
- Anthony M. Sloane. Lightweight language processing in kiama. In *Generative and Transformational Techniques in Software Engineering III - International Summer School, GTTSE 2009, Braga, Portugal, July 6-11, 2009. Revised Papers*, pages 408–425, 2009. doi: 10.1007/978-3-642-18023-1_12. URL https://doi.org/10.1007/978-3-642-18023-1_12.
- Michael B. Smyth and Gordon D. Plotkin. The category-theoretic solution of recursive domain equations. *SIAM J. Comput.*, 11(4):761–783, 1982. URL <http://dx.doi.org/10.1137/0211062>.
- Eugene Syriani and Hans Vangheluwe. A modular timed graph transformation language for simulation-based design. *Software & Systems Modeling*, 12(2):387–414, 2013. ISSN 1619-1366. doi: 10.1007/s10270-011-0205-0. URL <http://dx.doi.org/10.1007/s10270-011-0205-0>.
- Gabriele Taentzer. AGG: A graph transformation environment for modeling and validation of software. In JohnL. Pfaltz, Manfred Nagl, and Boris Böhlen, editors, *Applications of Graph Transformations with Industrial Relevance*, volume 3062 of *Lecture Notes in Computer Science*, pages 446–453. Springer Berlin Heidelberg, 2004. ISBN 978-3-540-22120-3. doi: 10.1007/978-3-540-25959-6_35. URL http://dx.doi.org/10.1007/978-3-540-25959-6_35.
- Gabriele Taentzer. What Algebraic Graph Transformations Can Do For Model Transformations. *Electronic Communications of the EASST*, 30, 2010.
- R. Gregory Taylor. *Models of Computation and Formal Languages*. Oxford University Press, Oxford, UK, 1998. ISBN 0-19-510983-X.

- The Eclipse Foundation. *Xtend User Guide*, September 2014. URL <http://www.eclipse.org/xtend/documentation/2.7.0/Xtend%20User%20Guide.pdf>.
- The PROGRES Developer Team. *The PROGRES Language Manual Version 9.x*, 1999. URL <http://www-i3.informatik.rwth-aachen.de/files/progres/PROGRESLanguageManual.pdf>.
- Nikolai Tillmann and Jonathan de Halleux. Pex-white box test generation for .net. In *Tests and Proofs, Second International Conference, TAP 2008, Prato, Italy, April 9-11, 2008. Proceedings*, pages 134–153, 2008. doi: 10.1007/978-3-540-79124-9_10. URL http://dx.doi.org/10.1007/978-3-540-79124-9_10.
- Massimo Tisi, Salvador Martínez, Frédéric Jouault, and Jordi Cabot. Lazy execution of model-to-model transformations. In Jon Whittle, Tony Clark, and Thomas Kühne, editors, *Model Driven Engineering Languages and Systems*, volume 6981 of *Lecture Notes in Computer Science*, pages 32–46. Springer Berlin Heidelberg, 2011. ISBN 978-3-642-24484-1. doi: 10.1007/978-3-642-24485-8_4. URL http://dx.doi.org/10.1007/978-3-642-24485-8_4.
- Emina Torlak and Daniel Jackson. Kodkod: A relational model finder. In *Tools and Algorithms for the Construction and Analysis of Systems, 13th International Conference, TACAS 2007, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2007 Braga, Portugal, March 24 - April 1, 2007, Proceedings*, pages 632–647, 2007. doi: 10.1007/978-3-540-71209-1_49. URL http://dx.doi.org/10.1007/978-3-540-71209-1_49.
- Antoine Toubhans, Bor-Yuh Evan Chang, and Xavier Rival. Reduced product combination of abstract domains for shapes. In *Verification, Model Checking, and Abstract Interpretation, 14th International Conference, VMCAI 2013, Rome, Italy, January 20-22, 2013. Proceedings*, pages 375–395, 2013. doi: 10.1007/978-3-642-35873-9_23. URL https://doi.org/10.1007/978-3-642-35873-9_23.
- Javier Troya and Antonio Vallecillo. A rewriting logic semantics for ATL. *Journal of Object Technology*, 10:5: 1–29, 2011. doi: 10.5381/jot.2011.10.1.a5. URL <http://dx.doi.org/10.5381/jot.2011.10.1.a5>.

- Arie van Deursen, Paul Klint, and Joost Visser. Domain-specific languages: An annotated bibliography. *SIGPLAN Notices*, 35(6):26–36, 2000. doi: 10.1145/352029.352035. URL <http://doi.acm.org/10.1145/352029.352035>.
- Hans Vangheluwe, Juan De Lara, and Pieter J. Mosterman. An introduction to multi-paradigm modelling and simulation. *Proceedings of the AIS '2002 conference (AI, Simulation and Planning in High Autonomy Systems)*, 2002.
- Dániel Varró and Andras Balogh. The model transformation language of the VIATRA2 framework. *Science of Computer Programming*, 68(3): 214 – 234, 2007. ISSN 0167-6423. doi: <http://dx.doi.org/10.1016/j.scico.2007.05.004>. URL <http://www.sciencedirect.com/science/article/pii/S016764230700127X>. Special Issue on Model Transformation.
- Niki Vazou, Patrick Maxim Rondon, and Ranjit Jhala. Abstract refinement types. In *Programming Languages and Systems - 22nd European Symposium on Programming, ESOP 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16-24, 2013. Proceedings*, pages 209–228, 2013. doi: 10.1007/978-3-642-37036-6_13. URL https://doi.org/10.1007/978-3-642-37036-6_13.
- Markus von Detten, Christian Heinzemann, Marie Christian Platenius, Jan Rieke, and Dietrich Travkin. Story diagrams - syntax and semantics. Technical report, University of Paderborn, 2012.
- Dennis Wagelaar, Jörn Guy Süss, Hugo Bruneliere, Ronan, William Piers, Salva Martínez, and Thierry Fortin. Atl/user guide - the atl language, September 2014. URL http://wiki.eclipse.org/ATL/User_Guide_-_The_ATL_Language.
- Junhua Wang, Soon-Kyeong Kim, and David A. Carrington. Verifying metamodel coverage of model transformations. In *17th Australian Software Engineering Conference (ASWEC 2006), 18-21 April 2006, Sydney, Australia*, pages 270–282, 2006. doi: 10.1109/ASWEC.2006.55. URL <http://dx.doi.org/10.1109/ASWEC.2006.55>.
- Xiaoliang Wang, Fabian Büttner, and Yngve Lamo. Verification of graph-based model transformations using alloy. *ECEASST*, 67, 2014. URL <http://journal.ub.tu-berlin.de/eceasst/article/view/943>.

- Jos B. Warmer and Anneke G. Kleppe. *The Object Constraint Language: Precise Modeling With Uml (Addison-Wesley Object Technology Series)*. Addison-Wesley Professional, October 1998. ISBN 0201379406. URL <http://www.worldcat.org/isbn/0201379406>.
- Glynn Winskel. *The formal semantics of programming languages - an introduction*. MIT Press, 1993a. ISBN 978-0-262-23169-5.
- Glynn Winskel. *Information Systems*, chapter 12. In Winskel [1993a], 1993b. ISBN 978-0-262-23169-5.
- Niklaus Wirth. *Compiler Construction*. International computer science series. Addison-Wesley, 1996. ISBN 9780201403534.
- Claes Wohlin, Per Runeson, Martin Höst, Magnus C. Ohlsson, and Björn Regnell. *Experimentation in Software Engineering*. Springer, 2012. ISBN 978-3-642-29043-5. doi: 10.1007/978-3-642-29044-2. URL <http://dx.doi.org/10.1007/978-3-642-29044-2>.
- Hongwei Xi and Frank Pfenning. Eliminating array bound checking through dependent types. In *Proceedings of the ACM SIGPLAN '98 Conference on Programming Language Design and Implementation (PLDI), Montreal, Canada, June 17-19, 1998*, pages 249–257, 1998. doi: 10.1145/277650.277732. URL <http://doi.acm.org/10.1145/277650.277732>.
- Yingfei Xiong, Dongxi Liu, Zhenjiang Hu, Haiyan Zhao, Masato Takeichi, and Hong Mei. Towards automatic model synchronization from model transformations. In *22nd IEEE/ACM International Conference on Automated Software Engineering (ASE 2007), November 5-9, 2007, Atlanta, Georgia, USA*, pages 164–173, 2007. doi: 10.1145/1321631.1321657. URL <http://doi.acm.org/10.1145/1321631.1321657>.
- Yingfei Xiong, Zhenjiang Hu, Haiyan Zhao, Hui Song, Masato Takeichi, and Hong Mei. Supporting automatic model inconsistency fixing. In *Proceedings of the the 7th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on The Foundations of Software Engineering, ESEC/FSE '09*, pages 315–324, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-001-2. doi: 10.1145/1595696.1595757. URL <http://doi.acm.org/10.1145/1595696.1595757>.
- Xuejun Yang, Yang Chen, Eric Eide, and John Regehr. Finding and understanding bugs in C compilers. In *Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation*,

PLDI 2011, San Jose, CA, USA, June 4-8, 2011, pages 283–294, 2011.
doi: 10.1145/1993498.1993532. URL <http://doi.acm.org/10.1145/1993498.1993532>.

Appendix A

Modernization Transformation Proofs

Proposition 4.1. *For each concrete execution path $\pi = (\sigma_{\text{in}}, \sigma_1, \dots, \sigma_{\text{out}})$ of the program P , there exists the corresponding symbolic execution path $\hat{\pi} = ((\hat{\sigma}_{\text{in}}, \text{true}), \dots, (\hat{\sigma}_{\text{out}}, \hat{b}_{\text{out}}))$, such that the following equations are satisfied:*

$$\begin{aligned}\hat{\sigma}_{\text{in}} &= [x_i \mapsto x_i^?, \dots, x_n \mapsto x_n^?] \\ \theta &= [x_0^? \mapsto \sigma_{\text{in}}(x_0), \dots, x_n^? \mapsto \sigma_{\text{in}}(x_n)] \\ ((\hat{\sigma}_{\text{out}})\theta)(\text{ret}) &= \sigma_{\text{out}}(\text{ret}) \\ (\hat{b})\theta &= \text{true}\end{aligned}$$

Proof. By induction on the length of the path π , and by showing that each concrete transition $\sigma_i \xrightarrow{s} \sigma_{i+1}$ from π can be simulated symbolically with some transition $(\hat{\sigma}_i, \hat{b}_i) \xrightarrow{s} (\hat{\sigma}_{i+1}, \hat{b}_{i+1})$ in a corresponding $\hat{\pi}$. Let θ be a substitution assigning each symbol \hat{x}_i the value $\sigma_0(\hat{x}_i)$. We say that a constrained symbolic store $(\hat{\sigma}, \hat{b})$ simulates a concrete store σ if $\sigma = (\hat{\sigma})\theta$ and $(\hat{b})\theta$ is true. Now, let $\sigma_i \xrightarrow{s} \sigma_{i+1}$ be a transition in π and σ_i be simulated by $(\hat{\sigma}_i, \hat{b}_i)$; this transition can be symbolically executed resulting in $(\hat{\sigma}_i, \hat{b}_i) \xrightarrow{s} (\hat{\sigma}_{i+1}, \hat{b}_{i+1})$ so σ_{i+1} is simulated by $(\hat{\sigma}_{i+1}, \hat{b}_{i+1})$. This way, by simulating each transition from π symbolically, we can generate the corresponding symbolic path $\hat{\pi}$ fulfilling the stated requirements. \square

Theorem 4.1. *Two programs P and P' are semantically equivalent $P \sim P'$ iff for each value $v \in \text{Val}$ it holds:*

$$\bigvee_{(\hat{\sigma}_{\text{out}}, \hat{b}_{\text{out}}) \in \Sigma} (\hat{\sigma}_{\text{out}}(\text{ret}) = v \wedge \hat{b}_{\text{out}}) \iff \bigvee_{(\hat{\sigma}'_{\text{out}}, \hat{b}'_{\text{out}}) \in \Sigma'} (\hat{\sigma}'_{\text{out}}(\text{ret}) = v \wedge \hat{b}'_{\text{out}})$$

where

$$\Sigma = \left\{ (\hat{\sigma}_{\text{out}}, \hat{b}_{\text{out}}) \mid ((\hat{\sigma}_{\text{in}}, \text{true}) \dots (\hat{\sigma}_{\text{out}}, \hat{b}_{\text{out}})) \text{ is a symbolic path of } P \right\}$$

$$\Sigma' = \left\{ (\hat{\sigma}'_{\text{out}}, \hat{b}'_{\text{out}}) \mid ((\hat{\sigma}_{\text{in}}, \text{true}) \dots (\hat{\sigma}'_{\text{out}}, \hat{b}'_{\text{out}})) \text{ is a symbolic path of } P' \right\}$$

Proof. Follows directly from Proposition 4.1 and Definition 4.6. □

Appendix B

SymexTRON Auxiliary Functions

$$\widehat{\text{singleton}}(\widehat{e}, \widehat{h}) = \begin{cases} \{(x^?, \widehat{h})\} & \text{if } \widehat{e} = \{x^?\} \\ \widehat{\text{mk-singleton}}(\widehat{e}, \widehat{h}) & \text{otherwise} \end{cases}$$

$$\widehat{\text{mk-singleton}}(\widehat{e}, (\widehat{z}, \widehat{\ell}, \widehat{d}, \widehat{\Gamma}, \widehat{b})) = \{(x^?, \widehat{h}') \mid x^? \text{ fresh} \wedge \widehat{h}' \text{ sat}\}$$

where $x^?$ fresh

$$\widehat{\Gamma}' = \widehat{\Gamma}[x^? \mapsto \widehat{\text{types}}(\widehat{e}, \widehat{\Gamma}, \widehat{z})]$$

$$\widehat{h}' = (\widehat{z}, \widehat{\ell}, \widehat{d}, \widehat{\Gamma}', \widehat{b} \wedge \widehat{e} = \{x^?\})$$

$$\widehat{\text{inst}}(x^?, \widehat{h}) = \begin{cases} \{(\widehat{z}(x^?), \widehat{h})\} & \text{if } x^? \in \text{dom } \widehat{z} \\ \widehat{\text{mk-inst}}(x^?, \widehat{h}) \cup \widehat{\text{alias-inst}}(x^?, \widehat{h}) & \text{otherwise} \end{cases}$$

where $\widehat{h} = (\widehat{z}, \widehat{\ell}, \widehat{d}, \widehat{\Gamma}, \widehat{b})$

$$\widehat{\text{mk-inst}}(x^?, \widehat{h}) = \left\{ (\widehat{o}, \widehat{h}') \mid \begin{array}{l} c \in \text{cs}_{\text{in}} \wedge \\ \widehat{h}' = \widehat{\text{mk-inst-typ}}(x^?, \widehat{o}, c, \text{cs}_{\text{ex}}, \widehat{h}) \wedge \widehat{h}' \text{ sat} \end{array} \right\}$$

where \widehat{o} fresh

$$\widehat{h} = (\widehat{z}, \widehat{\ell}, \widehat{d}, \widehat{\Gamma}, \widehat{b})$$

$$(\text{cs}_{\text{in}}, \text{cs}_{\text{ex}}) = \widehat{\Gamma}(x^?)$$

$$\widehat{\text{mk-inst-typ}}(x^?, \widehat{o}, c, \text{cs}_{\text{ex}}, \widehat{h}) = (\widehat{z}[x^? \mapsto \widehat{o}], \widehat{\ell}', \widehat{d}, \widehat{\Gamma}', \widehat{b})$$

where $(\widehat{fs}, \Gamma'') = \widehat{\text{mk-fields}}(\text{fields}(c), \widehat{\Gamma})$

$$\widehat{\ell}' = \widehat{\ell}[(\widehat{o}, f) \mapsto \widehat{e} \mid (f, \widehat{e}) \in \widehat{fs}]$$

$$\widehat{\Gamma}' = \widehat{\Gamma}'' \left[\widehat{o} \mapsto (\{c\}, \left\{ c' \mid \begin{array}{l} c' \text{ gen}^* c \wedge \\ c' \in \text{cs}_{\text{ex}} \end{array} \right\}) \right]$$

$$\widehat{\text{mk-fields}}(fs_c, \widehat{\Gamma}) = \widehat{\text{mk-fields}}'(\emptyset, fs_c, \widehat{\Gamma})$$

$$\widehat{\text{mk-fields}}'(fs_e, fs_c, \widehat{\Gamma}) = \begin{cases} \widehat{\text{mk-fields}}'(fs_e \cup \{(f, X^?)\}, fs'_c, \widehat{\Gamma}') & \text{if } fs_c = (f, c) \uplus fs'_c \\ (fs_e, \widehat{\Gamma}) & \text{otherwise} \end{cases}$$

where $X^?$ fresh

$$\widehat{\Gamma}' = \widehat{\Gamma}[X^? \mapsto c]$$

$$\widehat{\text{alias-inst}}(x^?, \widehat{h}) = \left\{ (\widehat{o}, \widehat{h}') \mid \begin{array}{l} \widehat{o} \in \widehat{\text{type-insts}}(\widehat{cs}_{\text{in}}, \widehat{\Gamma}) \wedge \\ \widehat{h}' = \widehat{\text{alias-inst-agn}}(x^?, \widehat{o}, cs_{\text{out}}, \widehat{h}) \wedge \widehat{h}' \text{ sat} \end{array} \right\}$$

where $(cs_{\text{in}}, cs_{\text{out}}) = \Gamma(x^?)$

$$\widehat{\text{alias-inst-agn}}(x^?, \widehat{o}, cs_{\text{out}}, \widehat{h}) = (\widehat{z}[x^? \mapsto \widehat{o}], \widehat{\ell}, \widehat{d}, \widehat{\Gamma}', \widehat{b})$$

where $(\{c'\}, cs''_{\text{out}}) = \widehat{\Gamma}(o^?)$

$$cs'_{\text{out}} = cs''_{\text{out}} \cup \{c \mid c \in cs_{\text{out}} \wedge c \text{ gen}^* c'\}$$

$$\widehat{\Gamma}' = \widehat{\Gamma}[\widehat{o} \mapsto (\{c'\}, cs'_{\text{out}})] \setminus \widehat{x}$$

$$\widehat{\text{type-insts}}(cs, \widehat{\Gamma}) = \left\{ \widehat{o} \mid \begin{array}{l} (\{c'\}, cs_{\text{out}}) = \widehat{\Gamma}(o) \wedge \\ (\exists c \in cs.c \text{ gen}^* c' \wedge \nexists c'' \in cs_{\text{out}.c \text{ gen}^* c'') \end{array} \right\}$$

$$\widehat{\text{dc-containment}}(\widehat{e}, c, \widehat{z}, \widehat{d}, \widehat{\Gamma}) = \widehat{\text{dc-reown}}(\widehat{e}, c, \widehat{dcs}', \emptyset, \widehat{z}, \widehat{\Gamma}, \mathbf{true})$$

where $\widehat{dcs}' = \{(\widehat{o}, c', \widehat{e}' \setminus \widehat{e}) \mid (\widehat{o}, c', \widehat{e}') \in \text{graph } \widehat{d}\}$

$$\widehat{\text{dc-reown}}(\widehat{e}, c, \widehat{dcs}, \widehat{dcs}'', \widehat{z}, \widehat{\Gamma}, \widehat{b}) = \begin{cases} \widehat{\text{dc-reown}}(\widehat{e}, c, \widehat{dcs}', \widehat{dcs}'' \cup \{\widehat{dc}'\}, \widehat{\Gamma}', \widehat{b} \wedge \widehat{b}') & \text{if } \widehat{dcs} = \widehat{dc} \uplus \widehat{dcs}' \\ [(\widehat{o}, c') \mapsto \widehat{e}'' \mid (\widehat{o}, c', \widehat{e}'') \in \widehat{dcs}''] & \text{otherwise} \end{cases}$$

where $(\widehat{dc}', \widehat{\Gamma}', \widehat{b}') = \widehat{\text{dc-reown-1}}(\widehat{e}, \widehat{c}, \widehat{dc}, \widehat{z}, \widehat{\Gamma})$

$$\widehat{\text{dc-reown-1}}(\widehat{e}, c, \widehat{dc}, \widehat{z}, \widehat{\Gamma}) = \begin{cases} \widehat{\text{dc-reown-1-sub}}(\widehat{e}, c, \widehat{dc}, \widehat{\Gamma}) & \text{if } c <: c' \\ \widehat{\text{dc-reown-1-sup}}(\widehat{e}, c, \widehat{dc}, \widehat{z}, \widehat{\Gamma}) & \text{if } c' <: c \wedge c \neq c' \\ (\widehat{dc}, \widehat{\Gamma}, \mathbf{true}) & \text{otherwise} \end{cases}$$

$$\widehat{\text{dc-reown-1-sub}}(\widehat{e}, c, (\widehat{o}, c', \widehat{e}'), \widehat{\Gamma}) = ((\widehat{o}, c', \widehat{e}' \cup X^?), \widehat{\Gamma}[X^? \mapsto c], X^? = \emptyset \vee X^? = \widehat{e})$$

where $X^?$ fresh

$$\widehat{\text{dc-reown-1-sup}}(\widehat{e}, c, (\widehat{o}, c', \widehat{e}'), \widehat{z}, \widehat{\Gamma}) = ((\widehat{o}, c', \widehat{e}' \cup X^?), \widehat{\Gamma}', \widehat{b})$$

where $X^?, Y^?, Z^?$ fresh

$$(cs_{\text{in}}, cs_{\text{ex}}) = \widehat{\text{types}}(\widehat{e}, \widehat{\Gamma}, \widehat{z})$$

$$\widehat{\Gamma}' = \widehat{\Gamma}[X^? \mapsto c, Y^? \mapsto c', Z^? \mapsto (cs_{\text{in}}, cs_{\text{ex}} \cup c)']$$

$$\widehat{b} = (X^? = \emptyset \vee X^? = Y^?) \wedge \widehat{e} = Y^? \uplus Z^?$$

$$\widehat{\text{match}}(x^?, X^?, c, \widehat{h}) = \widehat{\text{match-sucs}}(x^?, c, \widehat{h}) \cup \widehat{\text{match-fail}}(x^?, X^?, c, \widehat{h})$$

$$\widehat{\text{match-sucs}}(x^?, c, (\widehat{z}, \widehat{\ell}, \widehat{d}, \widehat{\Gamma}, \widehat{b})) = \left\{ (\text{tt}, (\widehat{z}, \widehat{\ell}, \widehat{d}, \widehat{\Gamma}', \widehat{b})) \mid \begin{array}{l} (\exists c' \in cs_{\text{in}}. c' <: c \vee c <: c') \wedge \\ (\nexists c' \in cs_{\text{ex}}. c' <: c) \end{array} \right\}$$

where $(cs_{\text{in}}, cs_{\text{ex}}) = \widehat{\text{typeof}}(x^?, \widehat{\Gamma}, \widehat{z})$

$$\widehat{\Gamma}' = \begin{cases} \widehat{\Gamma}[\widehat{z}(x^?) \mapsto (c, \{c' \mid c' \in cs_{\text{ex}} \wedge c' <: c\})] & \text{if } x^? \in \text{dom } \widehat{z} \\ \widehat{\Gamma}[x^? \mapsto (c, \{c' \mid c' \in cs_{\text{ex}} \wedge c' <: c\})] & \text{otherwise} \end{cases}$$

$$\widehat{\text{match-fail}}(x^?, X^?, c, (\widehat{z}, \widehat{\ell}, \widehat{d}, \widehat{\Gamma}, \widehat{b})) = \left\{ (\text{ff}, (\widehat{z}, \widehat{\ell}, \widehat{d}, \widehat{\Gamma}', \widehat{b})) \mid \begin{array}{l} (\exists c' \in cs_{\text{in}}. c \not<: c') \vee \\ (\exists c' \in cs_{\text{ex}}. c' <: c) \end{array} \right\}$$

where $(cs_{\text{in}}, cs_{\text{ex}}) = \widehat{\text{typeof}}(x^?, \widehat{\Gamma}, \widehat{z})$

$$(cs'_{\text{in}}, cs'_{\text{ex}}) = \Gamma(X^?)$$

$$\widehat{\Gamma}' = \begin{cases} \Gamma[\widehat{z}(x^?) \mapsto (cs_{\text{in}}, cs_{\text{ex}} \cup c), X^? \mapsto (cs'_{\text{in}}, cs'_{\text{ex}} \cup c)] & \text{if } x^? \in \text{dom } \widehat{z} \\ \Gamma[x^? \mapsto (cs_{\text{in}}, cs_{\text{ex}} \cup c), X^? \mapsto (cs'_{\text{in}}, cs'_{\text{ex}} \cup c)] & \text{otherwise} \end{cases}$$

$$\widehat{\text{dcs}}(x^?, c, \widehat{d}, \widehat{h}) = \left\{ (\widehat{e}, \widehat{h}') \mid (\widehat{o}, \widehat{h}'') \in \widehat{\text{inst}}(x^?, \widehat{h}) \wedge (\widehat{e}, \widehat{h}') = \widehat{\text{dcs}}'(\widehat{o}, c, \widehat{d}, \widehat{h}'') \right\}$$

$$\widehat{\text{dcs}}'(\widehat{o}, c, \widehat{d}_o, \widehat{h}) = \begin{cases} (\widehat{d}_o(\widehat{o}, c), \widehat{h}) & \text{if } (\widehat{o}, c) \in \text{dom } \widehat{d}_o \\ (X^?, (\widehat{z}, \widehat{\ell}, \widehat{d}', \widehat{\Gamma}', \widehat{b})) & \text{otherwise} \end{cases}$$

where $X^?$ fresh

$$\widehat{h} = (\widehat{z}, \widehat{\ell}, \widehat{d}, \widehat{\Gamma}, \widehat{b})$$

$$\widehat{d}' = \widehat{d}[(\widehat{o}, c) \mapsto X^?]$$

$$\widehat{\Gamma}' = \widehat{\Gamma}[X^? \mapsto c]$$

$$\widehat{\text{types}}(X^?, \widehat{\Gamma}, \widehat{z}) = \Gamma(X^?)$$

$$\widehat{\text{types}}(\emptyset, \widehat{\Gamma}, \widehat{z}) = (\emptyset, \emptyset)$$

$$\widehat{\text{types}}(\{x_1^?, \dots, x_n^?\}, \widehat{\Gamma}, \widehat{z}) = \bigsqcup_{i=1}^n \widehat{\text{typeof}}(x_i^?, \Gamma, z)$$

$$\widehat{\text{types}}(\widehat{e}_1 \cup \widehat{e}_2, \widehat{\Gamma}, \widehat{z}) = \widehat{\text{types}}(\widehat{e}_1, \widehat{\Gamma}, \widehat{z}) \sqcup \widehat{\text{types}}(\widehat{e}_2, \widehat{\Gamma}, \widehat{z})$$

$$\widehat{\text{types}}(\widehat{e}_1 \cap \widehat{e}_2, \widehat{\Gamma}, \widehat{z}) = \widehat{\text{types}}(\widehat{e}_1, \widehat{\Gamma}, \widehat{z}) \sqcup \widehat{\text{types}}(\widehat{e}_2, \widehat{\Gamma}, \widehat{z})$$

$$\widehat{\text{types}}(\widehat{e}_1 \setminus \widehat{e}_2, \widehat{\Gamma}, \widehat{z}) = \widehat{\text{types}}(\widehat{e}_1, \widehat{\Gamma}, \widehat{z})$$

$$\widehat{\text{typeof}}(x^?, \widehat{\Gamma}, \widehat{z}) = \begin{cases} \widehat{\Gamma}(\widehat{z}(x^?)) & \text{if } x^? \in \text{dom } \widehat{z} \\ \widehat{\Gamma}(x^?) & \text{otherwise} \end{cases}$$

$$(cs_{\text{in}}, cs_{\text{ex}}) \sqcup (cs'_{\text{in}}, cs'_{\text{ex}}) = (\widehat{\text{type-ub}}(cs_{\text{in}}, cs'_{\text{in}}), \widehat{\text{type-lb}}(cs_{\text{ex}}, cs'_{\text{ex}}))$$

where $\widehat{\text{type-ub}}(cs, cs') = \{c \mid c \in cs \cup cs' \wedge \nexists c' \in cs \cup cs'. c <: c'\}$
 $\widehat{\text{type-lb}}(cs, cs') = \{c \mid c \in cs \cap cs' \wedge \nexists c' \in cs \cap cs'. c <: c'\}$

$$\widehat{\text{partition}}(\widehat{e}, (\widehat{z}, \widehat{\ell}, \widehat{d}, \widehat{\Gamma}', \widehat{b})) = (\widehat{z}, \widehat{\ell}, \widehat{d}, \widehat{\Gamma}', \widehat{b} \wedge \widehat{e} = \{x^?\} \uplus X^?)$$

where $x^?, X^?$ fresh

$$(cs_{\text{in}}, cs_{\text{ex}}) = \widehat{\text{types}}(\widehat{e}, \widehat{\Gamma}, \widehat{z})$$

$$\widehat{\Gamma}' = \widehat{\Gamma}[x^? \mapsto (cs_{\text{in}}, cs_{\text{ex}}), X^? \mapsto (cs_{\text{in}}, cs_{\text{ex}})]$$

Appendix C

SymexTRON Proofs

Theorem 5.1 (Soundness). *If $\exists m. \sigma = m(\hat{\sigma}) \wedge \Gamma, h \stackrel{m}{\models} \hat{h}$ and $\widehat{\mathcal{S}}[[s]](\hat{\sigma}, \hat{h}) = \widehat{\mathbf{M}}$ then for all $(\hat{\sigma}, \hat{h}) \in \widehat{\mathbf{M}}$ there exists σ', Γ' , and h' such that we have a concrete execution $s, \sigma, \Gamma, h \implies \sigma', \Gamma', h'$ and exists a model m' such that $\sigma' = m'(\hat{\sigma}')$ and $\Gamma', h' \stackrel{m'}{\models} \hat{h}'$.*

Proof Sketch. The basic rules follow standard techniques from General Symbolic Execution Khurshid et al. [2003] and the Lazier# algorithm Deng et al. [2012]. The only non-standard components of our formal algorithm are type bounds, containment updates, lazy iteration and deep containment constraints. We discuss the soundness of these components:

- Type bounds are precise for symbolic instances, but may over-approximate for symbols and symbolic references. In particular, during partitioning of symbolic expressions, new symbols and symbolic reference sets are created and their assigned types is inferred from the symbolic expression. Despite this local over-approximation soundness is still preserved, since partitioning constraints are never dropped globally and we can thus only assign symbolic instances to symbols—which is an under-approximative operation—if the assignment agrees with these global constraints (otherwise the resulting heap is unsatisfiable and not considered).
- The containment property, means that each instance is at most contained in a single other instance, and that containment is acyclic. To ensure that this property is preserved in concrete field updates, the rule checks acyclicity before rule assignment and then proceeds

to remove the set of newly assigned instances from the other containment fields. Both acyclicity and unique ownership are correspondingly ensured by the symbolic semantics. The acyclicity constraint is encoded in the symbolic constraints of containment, and thus cyclic assignment produces an unsatisfiable heap. The unique ownership constraint is ensured by the $\widehat{\text{disown}}$ function which instead uses the set difference operation to ensure that the assigned symbolic set references do not contain instances symbolically represented by the newly assigned symbolic expression; since containment constraints are also affected, it ensures to update these as well with the $\widehat{\text{dc-ownership}}$ function.

- Lazy iteration should model concrete iteration and so should symbolically represent the set of values matched in a particular state. However, because of laziness of the operation and possibility of state change at each iteration it is required that the $\widehat{\text{next}}$ operation keeps track of the original set of descendant constraints when performing a new iteration, since otherwise the heap can become inconsistent with concrete iteration.
- Containment constraints are additionally to being consistently updated as mentioned above, also encoded directly as a constraint in the model finder for the symbolic heap. This means that it is only possible to make assignments for referenced symbols to symbolic instances in a way consistent with the type and containment constraints; any non-sound assignment will produce an unsatisfiable symbolic heap, which is not further considered in the execution.

□

Appendix D

TRON Subject programs

Data model for RenameField, RenameMethod, ExtractSuper, ReplaceDelegation:

```
1  class Package {
2    classes :◆ Class*
3  }
4  class Class {
5    name   : String
6    super  :↗ Class
7    fields :◆ Field*
8    methods :◆ Method*
9  }
10 class Field {
11   name : String
12   type :↗ Class
13 }
14 class Method {
15   name   : String
16   params :◆ Parameter*
17   body   :◆ Statement
18   type   :↗ Class
19 }
20 class Parameter {
21   name : String
22   type :↗ Class
23 }
24 class Statement {}
25 class IfStatement extends Statement {
26   then :◆ Statement
27   else :◆ Statement
28   cond :◆ Expr
29 }
30 class Return extends Statement {
31   value :◆ Expr
32 }
33 class Assign Statement {
34   left :◆ AssignableExpr
```

```

35   right :◆ Expr
36   }
37   class Expr {
38     type :… Class
39   }
40   class AssignableExpr extends Expr { }
41   class FieldAccessExpr extends AssignableExpr {
42     field_name : String
43     target     :◆ Expr
44   }
45   class MethodCallExpr extends Expr {
46     method_name : String
47     target      :◆ Expr
48     args       :◆ Arg*
49   }
50   class Arg {
51     name : String
52     value :◆ Expr
53   }

```

RenameField:

```

1 class_fields := class.fields;
2 class.fields := (class_fields \ old_field) ∪ new_field;
3 foreach faexpr ∈ package match* FieldAccessExpr do
4   faexpr_field_name := faexpr.field_name;
5   old_field_name := old_field.name;
6   faexpr_target := faexpr.target;
7   faexpr_target_type := faexpr_target.type;
8   if faexpr_field_name = old_field_name ∧
9     class = faexpr_target_type then
10    new_field_name := new_field.name;
11    faexpr.field_name := new_field_name
12  else skip

```

RenameMethod:

```

1 class_methods := class.methods;
2 class.methods := (class_methods \ old_method) ∪ new_method;
3 foreach mcexpr ∈ package match* MethodCallExpr do
4   mcexpr_method_name := mcexpr.method_name;
5   old_method_name := old_method.name;
6   old_method_params := old_method.params;
7   mcexpr_target := mcexpr.target;
8   mcexpr_target_type := mcexpr_target.type;
9   mcexpr_args := mcexpr.args;
10  paramsmatched := new Any;
11  foreach omp ∈ old_method_params do
12    parammatched := ∅;
13    omp_name := omp.name;
14    foreach mcea ∈ mcexpr_args do
15      mcea_name := mcea.name;
16      if omp_name = mcea_name then
17        parammatched := paramsmatched
18      else skip;
19  if parammatched = ∅ then
20    paramsmatched := ∅

```

```

21   else skip;
22   if mcepr_method_name = old_method_name ∧
23       class = mcepr_target_type ∧
24       ((¬(old_method_params = ∅) ∧ mcepr_args = ∅)
25         ∨ paramsmatched = ∅)
26   then
27     new_method_name := new_method.name;
28     mcepr.method_name := new_method_name
29   else skip

```

ExtractSuper:

```

1  sclass := new Class;
2  package_classes := package.classes;
3  package.classes := package_classes ∪ sclass;
4  class1.super := sclass;
5  class2.super := sclass;
6  sclass.name := sc_name;
7  new_sclass_fields := ∅;
8  rem_class1_fields := ∅;
9  rem_class2_fields := ∅;
10 class1_fields := class1.fields;
11 class2_fields := class2.fields;
12 foreach c1f ∈ class1_fields do
13   foreach c2f ∈ class2_fields do
14     c1f_name := c1f.name;
15     c2f_name := c2f.name;
16     c1f_type := c1f.type;
17     c2f_type := c2f.type;
18     if c1f_name = c2f_name ∧ c2f_type = c1f_type then
19       scf := new Field;
20       scf.name := c1f_name;
21       scf.type := c1f_type;
22       new_sclass_fields := new_sclass_fields ∪ scf;
23       rem_class1_fields := rem_class1_fields ∪ c1f;
24       rem_class2_fields := rem_class2_fields ∪ c2f
25     else
26       skip;
27 class1_fields := class1_fields \ rem_class1_fields;
28 class2_fields := class2_fields \ rem_class2_fields;
29 sclass.fields := new_sclass_fields

```

ReplaceDelegation:

```

1  class_fields := class.fields;
2  field_type := field.type;
3  class.super := field_type;
4  field_type_methods := field_type.methods;
5  class_methods := class.methods;
6  class_new_methods := ∅;
7  foreach ftm ∈ field_type_methods do
8   foreach cm ∈ class_methods do
9     ftm_name := ftm.name;
10    cm_name := cm.name;
11    if ¬(ftm_name = cm_name) then
12      class_new_methods := class_new_methods ∪ cm
13    else skip;

```

```

14 class.methods := class_new_methods;
15 foreach mcexpr ∈ class match* MethodCallExpr do
16   mcexpr_target := mcexpr.target;
17   MCEXPRTARGET := ∅;
18   foreach mcx in mcexpr_target match FieldAccessExpr do
19     MCEXPRTARGET := mcx;
20     if ¬(MCEXPRTARGET = ∅) then
21       mcexpr_target_target := mcexpr_target.target;
22       mcexpr_target_target_type := mcexpr_target_target.type;
23       mcexpr_target_field_name := mcexpr_target.field_name;
24       field_name := field.name;
25       if field_name = mcexpr_target_field_name ∧
26         class = mcexpr_target_target_type then
27         mcexpr.target := mcexpr_target_target
28       else skip
29     else skip;
30 class.fields := class_fields \ field

```

Data models for Fam2Pers:

```

1 // Families meta model
2 class Family {
3   lastName : String
4   father :◆ Member opposite familyFather
5   mother :◆ Member opposite familyMother
6   sons :◆ Member* opposite familySon
7   daughters :◆ Member* opposite familyDaughter
8 }
9 class Member {
10  firstName : String
11  familyFather :↗ Family? opposite father
12  familyMother :↗ Family? opposite mother
13  familySon :↗ Family? opposite sons
14  familyDaughter :↗ Family? opposite daughters
15 }
16 // Persons meta model
17 class Person {
18   fullName : String
19 }
20 class Male extends Person { }
21 class Female extends Person { }

```

Fam2Pers:

```

1 persons := ∅;
2 foreach member ∈ families match* Member do
3   self_familyMother := member.familyMother;
4   self_familyDaughter := member.familyDaughter;
5   // Start inlined isFemale helper
6   if ¬(self_familyMother = ∅) then
7     isFemale := new Any
8   else if ¬(self_familyDaughter = ∅) then
9     isFemale := new Any
10  else
11    isFemale := ∅;
12  // End inlined isFemale helper

```



```

13  if ¬(isFemale = ∅) then
14    female := new Female;
15    self_familyFather := member.familyFather;
16    self_familyMother := member.familyMother;
17    self_familySon := member.familySon;
18    self_familyDaughter := member.familyDaughter;
19    // Start inlined familyName helper
20    if ¬(self_familyFather = ∅) then
21      familyName := self_familyFather.lastName
22    else if ¬(self_familyMother = ∅) then
23      familyName := self_familyMother.lastName
24    else if ¬(self_familySon = ∅) then
25      familyName := self_familySon.lastName
26    else
27      familyName := self_familyDaughter.lastName;
28    // End inlined familyName helper
29    member_firstName := member.firstName;
30    fullName := new Concat;
31    fullName.s1 := member_firstName;
32    fullName.s2 := familyName;
33    female.fullName := fullName;
34    persons := persons ∪ female
35  else
36    male := new Male;
37    self_familyFather := member.familyFather;
38    self_familyMother := member.familyMother;
39    self_familySon := member.familySon;
40    self_familyDaughter := member.familyDaughter;
41    // Start inlined familyName helper
42    if ¬(self_familyFather = ∅) then
43      familyName := self_familyFather.lastName
44    else if ¬(self_familyMother = ∅) then
45      familyName := self_familyMother.lastName
46    else if ¬(self_familySon = ∅) then
47      familyName := self_familySon.lastName
48    else
49      familyName := self_familyDaughter.lastName;
50    // End inlined familyName helper
51    member_firstName := member.firstName;
52    fullName := new Concat;
53    fullName.s1 := member_firstName;
54    fullName.s2 := familyName;
55    male.fullName := fullName;
56    persons := persons ∪ male

```

Data model for Path2Petri:

```

1  // Shared
2  class Element {
3    name : String
4  }
5  // Path expression meta model
6  class PathExp extends Element {
7    transitions :◆ PTransition*
8    states      :◆ State*
9  }
10 class State extends Element {

```

```

11   outgoing :↗ PETransition* opposite source
12   incoming :↗ PETransition* opposite target
13 }
14 class PETransition extends Element {
15   source :↗ State opposite outgoing
16   target :↗ State opposite incoming
17 }
18 // Petri net meta model
19 class PetriNet extends Element {
20   transitions :◆ PNTransition*
21   arcs :◆ Arc*
22   place :◆ Place*
23 }
24 class PNTransition extends Element {
25   outgoing :↗ TransToPlaceArc* opposite source
26   incoming :↗ PlaceToTransArc* opposite target
27 }
28 class Place extends Element {
29   outgoing :↗ PlaceToTransArc* opposite source
30   incoming :↗ TransToPlaceArc* opposite target
31 }
32 class Arc extends Element {
33   weight : Integer
34 }
35 class TransToPlaceArc extends Arc {
36   source :↗ PNTransition opposite outgoing
37   target :↗ Place opposite incoming
38 }
39 class PlaceToTransArc extends Arc {
40   source :↗ Place opposite incoming
41   target :↗ PNTransition opposite outgoing
42 }

```

Path2Petri:

```

1  places := ∅;
2  transitions := ∅;
3  eString := new Empty;
4  int1 := new Int;
5  // First pass to create places
6  foreach st ∈ pe match* State do
7    place := new Place;
8    st._Place := place;
9    place.name := eString;
10   places := places ∪ place;
11  foreach tr ∈ pe match* PETransition do
12    pntr := new PNTransition;
13    tr._PNTransition := pntr;
14    tr_name := tr.name;
15    pntr.name := tr_name;
16    pnia := new PlaceToTransArc;
17    tr._PN_IA := pnia;
18    pntr.incoming := pnia;
19    tr_source := tr.source;
20    tr_source_Place := tr_source._Place;
21    pnia.source := tr_source_Place;
22    pnia.target := pntr;

```

```

23  pnia.weight := int1;
24  pnoa := new TransToPlaceArc;
25  tr._PN_OA := pnoa;
26  pnttr.outgoing := pnoa;
27  pnoa.source := pnttr;
28  tr_target := tr.target;
29  tr_target_Place := tr_target._Place;
30  pnoa.target := tr_target_Place;
31  pnia.weight := int1;
32  transitions := transitions ∪ pnttr;
33  // Second pass to link places to arcs
34  foreach st ∈ pe match* State do
35    st_Place := st._Place;
36    pnoas := ∅;
37    st_incoming := st.incoming;
38    foreach inc ∈ st_incoming do
39      inc_PN_OA := inc._PN_OA;
40      pnoas := pnoas ∪ inc_PN_OA;
41    st_Place.incoming := pnoas;
42    pnias := ∅;
43    st_outgoing := st.outgoing;
44    foreach outg ∈ st_outgoing do
45      outg_PN_IA := outg._PN_IA;
46      pnias := pnias ∪ outg_PN_IA;
47    st_Place.outgoing := pnias;
48  pn := new PetriNet;
49  pe_name := pe.name;
50  pn.name := pe_name;
51  pn.places := places;
52  pn.transitions := transitions;
53  arcs := ∅;
54  foreach pnttr ∈ transitions do
55    pnia := pnttr._PN_IA;
56    pnoa := pnttr._PN_OA;
57    arcs := arcs ∪ pnia ∪ pnoa;
58  pn.arcs := arcs

```

Data models for Class2Rel:

```

1  // Class meta model
2  class NamedElt {
3    name : String
4  }
5  class Package {
6    classifiers :◆ Classifier*
7  }
8  class Classifier extends NamedElt { }
9  class DataType extends Classifier { }
10 class Class extends Classifier {
11   isAbstract : Boolean
12   attributes :◆ Attribute* opposite owner
13   super      :… Class?
14 }
15 class Attribute extends NamedElt {
16   isMultivalued : Boolean
17   type          :… Classifier
18   owner         :… Class opposite attribute

```

```

19 }
20 // Relational meta model
21 class Named {
22     name : String
23 }
24 class Schema {
25     tables :◆ Table*
26     types  :◆ Type*
27 }
28 class Table extends Named {
29     columns :◆ Column*
30     key     :↔ Column
31 }
32 class Column extends Named {
33     type    :↔ Type
34 }

```

Class2Rel:

```

1  objectIdType := new Type;
2  objectIdType.name := integer_name;
3  schema := new Schema;
4  foreach dt ∈ package match* DataType do
5      dt_name := dt.name;
6      if dt_name = integer_name then
7          dt._Type := objectIdType
8      else
9          type := new Type;
10         dt._Type := type;
11         type.name := dt_name;
12         schema.types := schema.types;
13         schema.types := schema.types ∪ type;
14  idString := new String;
15  objectIdString := new String;
16  foreach at ∈ package match* Attribute do
17      at_type := at.type;
18      at_isMultivalued := at.isMultivalued;
19      foreach _ ∈ at_type match DataType do
20          if at_isMultivalued = ∅ then
21              at_name := at.name;
22              at_type_Type := at_type._Type;
23              column := new Column;
24              column.name := at_name;
25              column.type := at_type_Type;
26              at._Column := column
27          else
28              at_owner := at.owner;
29              at_owner_name := at_owner.name;
30              at_name := at.name;
31              at_type_Type := at_type._Type;
32              tableName := new Concat;
33              tableName.s1 := at_owner_name;
34              tableName.s2 := at_name;
35              idName := new Concat;
36              idName.s1 := at_owner_name;
37              idName.s2 := idString;
38              id := new Column;

```

```

39     id.name := idName;
40     id.type := objectIdType;
41     value := new Column;
42     value.name := at_name;
43     value.type := at_type_Type;
44     table := new Table;
45     table.name := tableName;
46     table.key := id;
47     table.columns := id ∪ value;
48     schema_tables := schema.tables;
49     schema.tables := schema_tables ∪ table;
50   foreach _ ∈ at_type match Class do
51     if at_isMultivalued = ∅ then
52       at_name := at.name;
53       column_name := new Concat;
54       column_name.s1 := at_name;
55       column_name.s2 := idString;
56       column := new Column;
57       column.name := column_name;
58       column.type := objectIdType;
59       at._Column := column
60     else
61       at_owner := at.owner;
62       at_owner_name := at_owner.name;
63       at_name := at.name;
64       tableName := new Concat;
65       tableName.s1 := at_owner_name;
66       tableName.s2 := at_name;
67       idName := new Concat;
68       idName.s1 := at_owner_name;
69       idName.s2 := idString;
70       id := new Column;
71       id.name := idName;
72       id.type := objectIdType;
73       foreignKey := new Column;
74       foreignKey.name := at_name;
75       foreignKey.type := objectIdType;
76       table := new Table;
77       table.name := tableName;
78       table.key := id;
79       table.columns := id ∪ foreignKey;
80       schema_tables := schema.tables;
81       schema.tables := schema_tables ∪ table;
82   foreach class ∈ package match* Class do
83     class_name := class.name;
84     class_attributes := class.attributes;
85     key := new Column;
86     key.name := objectIdString;
87     key.type := objectIdType;
88     cols := key;
89     foreach at ∈ class_attributes do
90       at_isMultivalued := at.isMultivalued;
91       if at_isMultivalued = ∅ then
92         at_Column := at._Column;
93         cols := cols ∪ at_Column
94       else skip;

```

```

95     table := new Table;
96     table.name := class_name;
97     table.key := key;
98     table.columns := cols;
99     schema_tables := schema.tables;
100    schema_tables := schema_tables ∪ table

```

Toy1:

```

1 containselem := ∅;
2 foreach sublist ∈ list match* IntList do
3   sublist_data := sublist.data;
4   if elem = sublist_data then
5     containselem := new Any
6   else skip

```

Toy2:

```

1 if list = ∅ then
2   res := new Any
3 else
4   head := list.data;
5   list_next := list.next;
6   if list_next = ∅ then
7     res := new Any
8   else
9     fix list_next do
10      list_next_next := list_next.next;
11      if list_next_next = ∅ then
12        tail := list_next.data
13      else
14        list_next := list_next_next;
15      if head = tail then
16        res := new Any
17      else
18        res := ∅

```

Toy3:

```

1 table := new Table;
2 idcol := new IdColumn;
3 table.id := idcol;
4 table.columns := idcol;
5 class_attributes := class.attributes;
6 foreach attr ∈ class_attributes do
7   col := new DataColumn;
8   attrtype := attr.type;
9   col.type := attrtype;
10  tablecolumns := table.columns;
11  table.columns := tablecolumns ∪ col

```

Toy4:

```

1 table := new Table;
2 idcol := new IdColumn;
3 table.id := idcol; table.columns := idcol;
4 foreach attr ∈ class match* Attribute do

```

```

5  col := new DataColumn;
6  attrtype := attr.type;
7  col.type := attrtype;
8  tablecolumns := table.columns;
9  table.columns := tablecolumns ∪ col

```

Toy5:

```

1 timestamps := ∅;
2 foreach ts ∈ post match* Timestamp do
3   timestamps := timestamps ∪ ts

```

Toy6:

```

1 foreach sp ∈ post match* SinglePost do
2   sp_title := sp.title;
3   sp_title_value := sp_title.value;
4   new_sp_title := new CapitalisedTitle;
5   new_sp_title.value := sp_title_value;
6   sp.title := new_sp_title

```

Toy7:

```

1 invitationlist := ∅;
2 foreach person ∈ contactbook match* Person do
3   isadult := ∅;
4   person_age := person.age;
5   person_name := person.name;
6   foreach age ∈ person_age match Adult do
7     isadult := new Any;
8     if ¬(isadult = ∅) then
9       invited := new Invited;
10      invited.name := person_name;
11      invitationlist := invitationlist ∪ invited
12    else skip

```


Appendix E

Rascal Light Semantics Proofs

Theorem 6.1 (Backtracking purity). *If $\underline{cs}; v; \sigma \xrightarrow[\text{cases}]{\mathcal{CS}} \mathbf{fail}; \sigma'$ then $\sigma' = \sigma$*

Proof. By induction on the derivation \mathcal{CS} :

- Case $\mathcal{CS} = \text{ECS-EMP} \frac{}{\varepsilon; v; \sigma \xrightarrow[\text{cases}]{} \mathbf{fail}; \sigma}$, so $\underline{cs} = \varepsilon$ and $\sigma' = \sigma$. Holds by definition.

$$\sigma \vdash p \stackrel{?}{=} v \xrightarrow[\text{match}]{\mathcal{M}} \underline{\rho} \quad \underline{\rho}; e; \sigma \xrightarrow[\text{case}]{\mathcal{C}} \mathbf{fail}; \sigma''$$

- Case $\mathcal{CS} = \text{ECS-MORE-FAIL} \frac{\underline{cs}'; v; \sigma \xrightarrow[\text{cases}]{\mathcal{CS}'} \mathbf{fail}; \sigma'}{\mathbf{case} \ p \Rightarrow e, \underline{cs}'; v; \sigma \xrightarrow[\text{cases}]{} \mathbf{fail}; \sigma'}$,
so $\underline{cs} = (\mathbf{case} \ p \Rightarrow e, \underline{cs}')$ and $vres = \mathbf{fail}$.

– By inductive hypothesis of \mathcal{CS}' we get $\sigma' = \sigma$.

- Note: the rule ECS-MORE-ORD is inapplicable since its premise states that the result value $vres \neq \mathbf{fail}$.

□

In order to prove Theorem 6.2, we need to state some helper lemmas about sub-derivations.

We have a lemma for the auxiliary merge function:

Lemma E.1. *If we have $\underline{\rho}' = \text{merge}(\underline{\rho}_1, \dots, \underline{\rho}_n)$ and for each value $v \in \text{img } \rho_{i,j}$ in an environment in the input sequence of environment sequences $\underline{\rho}_1, \dots, \underline{\rho}_n$ there exists a type t so that $v : t$, then we have that for each value $v' \in \text{img } \rho'_k$ in an environment in the resulting environment sequence $\underline{\rho}'$ we have a type t' so that $v' : t'$*

Proof. Follows directly from the premises by induction of the input sequence of environment sequences $\underline{\rho}_1, \dots, \underline{\rho}_n$ \square

We have a lemma for the auxiliary children function:

Lemma E.2. *For a value v such that we have $\underline{v}' = \text{children}(v)$ and $\frac{\mathcal{T}}{v : t'}$, then there exists a type sequence \underline{t}' such that $\underline{v}' : \underline{t}'$*

Proof. By induction on the syntax of v :

- Cases $v = vb$ and $v = \blacksquare$, so $\underline{v}' = \varepsilon$. Holds trivially.

- Case $v = k(\underline{v}')$

- By inversion we get

$$\text{T-CONSTRUCTOR} \frac{\text{data } at = \dots \mid k(\underline{t}'') \mid \dots \quad \frac{\mathcal{T}'}{v : t'} \quad \frac{S\mathcal{T}'}{t' <: t''}}{k(\underline{v}') : at}$$

so $t = at$

- Here, $\underline{\mathcal{T}'}$ exactly represents our target goal

- Case $v = [\underline{v}']$

- By inversion we get $\mathcal{T} = \text{T-LIST} \frac{\frac{\mathcal{T}'}{v' : t'}}{[\underline{v}'] : \text{list}(\underline{\square} \underline{t}')}$, so $t = \text{list}(\underline{\square} \underline{t}')$

- Here, $\underline{\mathcal{T}'}$ exactly represents our target goal

- Case $v = \{\underline{v}'\}$

- By inversion we get $\mathcal{T} = \text{T-SET} \frac{\frac{\mathcal{T}'}{v' : t'}}{\{\underline{v}'\} : \text{set}(\underline{\square} \underline{t}')}$, so $t = \text{set}(\underline{\square} \underline{t}')$

- Here, $\underline{\mathcal{T}'}$ exactly represents our target goal

- Case $v = (\underline{v''} : \underline{v'''})$, so $\underline{v'} = \underline{v''}, \underline{v'''}$
 - By inversion we get

$$\mathcal{T} = \text{T-MAP} \frac{\frac{\mathcal{T}''}{\underline{v''} : t''} \quad \frac{\mathcal{T}'''}{\underline{v'''} : t'''}}{(\underline{v''} : \underline{v'''}) : \text{map}(\sqcup \underline{t''}, \sqcup \underline{t'''})}$$

- Here we take the concatenation of the premises $\mathcal{T}'', \mathcal{T}'''$ to fulfill our goal. □

We have two mutually recursive lemmas on pattern matching: Lemma E.3 and Lemma E.4.

Lemma E.3. *If $\sigma \vdash p \stackrel{\mathcal{M}}{?} \underset{\text{match}}{\underline{v}} \Longrightarrow \underline{\rho}$ and there exists a type t such that $\frac{\mathcal{T}}{v : t}$ and a type t' such that $v' : t'$ for each value in the input store $v' \in \text{img } \sigma$, then we have a type t'' for each value $v'' \in \text{img } \rho_i$ in an environment in the output sequence $\underline{\rho}$.*

Proof. By induction on the derivation \mathcal{M} :

- Cases $\mathcal{M} = \text{P-VAL-SUCS}$, $\mathcal{M} = \text{P-VAL-FAIL}$, $\mathcal{M} = \text{P-VAR-UNI}$, $\mathcal{M} = \text{P-VAR-FAIL}$, $\mathcal{M} = \text{P-CONS-FAIL}$, $\mathcal{M} = \text{P-TYPE-FAIL}$, $\mathcal{M} = \text{P-LIST-FAIL}$, $\mathcal{M} = \text{P-SET-FAIL}$, $\mathcal{M} = \text{P-NEG-SUCS}$, $\mathcal{M} = \text{P-NEG-FAIL}$ hold trivially since we have that the output environment sequence $\underline{\rho} = []$ or $\underline{\rho} = \varepsilon$, in both cases containing no values.
- Case $\mathcal{M} = \text{P-VAR-BIND}$ holds by the premise derivation \mathcal{T} .

- Case $\mathcal{M} = \text{P-CONS-SUCS}$ $\frac{\sigma \vdash p'_1 \stackrel{\mathcal{M}'_1}{?} \underset{\text{match}}{\underline{v'_1}} \Longrightarrow \underline{\rho'_1} \quad \dots \quad \sigma \vdash p'_n \stackrel{\mathcal{M}'_n}{?} \underset{\text{match}}{\underline{v'_n}} \Longrightarrow \underline{\rho'_n}}{\sigma \vdash k(\underline{p'}) \stackrel{?}{\underset{\text{match}}{\underline{v}}} \Longrightarrow \text{merge}(\underline{\rho'_1}, \dots, \underline{\rho'_n})}$,
so $p = k(\underline{p'})$, $v = k(\underline{v'})$ and $\underline{\rho} = \text{merge}(\underline{\rho'_1}, \dots, \underline{\rho'_n})$

- By inversion we get

$$\mathcal{T} = \text{T-CONSTRUCTOR} \frac{\text{data } at = \dots \mid k(\underline{t''}) \mid \dots \quad \frac{\mathcal{T}'}{v : t'} \quad \frac{S\mathcal{T}'}{t' <: t''}}{k(\underline{v'_1}, \dots, \underline{v'_n}) : at}$$

so $t = at$

– Now by induction hypotheses of \mathcal{M}' using \mathcal{T}' , and then using Lemma E.1 we get that $\underline{\rho}$ has well-typed values.

- Case $\mathcal{M} = \text{P-TYPE-SUCS}$ holds by the premise derivation \mathcal{T} and Lemma E.1.

$$\bullet \text{ Case } \mathcal{M} = \text{P-LIST-SUCS} \frac{\mathcal{MS}' \quad \sigma \vdash \star p' \stackrel{?}{=} \underline{v'} \mid \emptyset \xrightarrow{\text{match}_\star} \underline{\rho}}{\sigma \vdash [\star p'] \stackrel{?}{=} [\underline{v'}] \xrightarrow{\text{match}} \underline{\rho}}, \text{ so } p = [\star p'] \text{ and } v = [\underline{v'}]$$

– By inversion we get $\mathcal{T} = \text{T-LIST} \frac{\mathcal{T}' \quad v' : t'}{[\underline{v'}] : \text{list}(\underline{\square} t')}$, so $t = \text{list}(\underline{\square} t')$

– Partitioning lists by splitting on concatenation, preserves typing since we can for each value pick the corresponding type derivation in the sequence.

– By induction hypothesis (using Lemma E.4) of \mathcal{MS}' using \mathcal{T}' and above fact we get that $\underline{\rho}$ has well-typed values.

$$\bullet \text{ Case } \mathcal{M} = \text{P-SET-SUCS} \frac{\mathcal{MS}' \quad \sigma \vdash \star p' \stackrel{?}{=} \underline{v'} \mid \emptyset \xrightarrow{\text{match}_\star} \underline{\rho}}{\sigma \vdash \{\star p'\} \stackrel{?}{=} \{\underline{v'}\} \xrightarrow{\text{match}} \underline{\rho}}, \text{ so } p = \{\star p'\} \text{ and } v = \{\underline{v'}\}$$

– By inversion we get $\mathcal{T} = \text{T-SET} \frac{\mathcal{T}' \quad v' : t'}{\{\underline{v'}\} : \text{set}(\underline{\square} t')}$, so $t = \text{set}(\underline{\square} t')$

– Partitioning sets using disjoint union \uplus preserves typing of value sub-sequences, since we can for each value v_i in a sub-sequence pick the corresponding typing derivation \mathcal{T}'_i

– By induction hypothesis (using Lemma E.4) of \mathcal{MS}' using \mathcal{T}' and above fact we get that $\underline{\rho}$ has well-typed values.

- $$\frac{\begin{array}{c} \mathcal{M}' \\ \sigma \vdash p' \stackrel{?}{=} v \xrightarrow[\text{match}]{\rho'} \underline{\rho'} \quad v'_1, \dots, v'_n = \text{children}(v) \\ \mathcal{M}''_1 \quad \mathcal{M}''_n \\ \sigma \vdash /p' \stackrel{?}{=} v'_1 \xrightarrow[\text{match}]{\rho''_1} \underline{\rho''_1} \quad \dots \quad \sigma \vdash /p' \stackrel{?}{=} v'_n \xrightarrow[\text{match}]{\rho''_n} \underline{\rho''_n} \end{array}}{\sigma \vdash /p' \stackrel{?}{=} v \xrightarrow[\text{match}]{\rho', \rho''_1, \dots, \rho''_n} \underline{\rho', \rho''_1, \dots, \rho''_n}}$$
- Case $\mathcal{M} = \text{P-DEEP}$
 - By induction hypothesis on \mathcal{M}' , we get that $\underline{\rho'}$ has well-typed values
 - By using Lemma E.2 on $\underline{v'}$, we get $\underline{v'} : t'$
 - By induction hypotheses on $\underline{\mathcal{M}''}$ using \mathcal{T}' we get that $\underline{\rho''_1}, \dots, \underline{\rho''_n}$ is well-typed
 - Now, we can show that $\underline{\rho'}, \underline{\rho''_1}, \dots, \underline{\rho''_n}$ is well-typed using above facts

□

Lemma E.4. If $\sigma \vdash \star p \stackrel{?}{=} \underline{v} \mid \mathbb{V} \xrightarrow[\text{match}\star]{\rho} \underline{\rho}$ and the following properties hold:

1. There exists a type sequence \underline{t} such that $\underline{v} : \underline{t}$
2. There exists a type t' such that $v' : t'$ for each value in the input store $v' \in \text{img } \sigma$
3. There exists a type t' such that $v' : t'$ for each value in the visited value set $v' \in \mathbb{V}$
4. If for all value sequences $\underline{v'}, \underline{v''}, \underline{v'''}$ where we have $\underline{v'} = \underline{v''} \otimes \underline{v'''}$ and there exists a type sequence $\underline{t'}$ so $\underline{v'} : \underline{t'}$, then we have that there exists two type sequences $\underline{t''}$ and $\underline{t'''}$ so that $\underline{v''} : \underline{t''}$ and $\underline{v'''} : \underline{t'''}$

Then we have a type t' so that $v' : t'$ for each value $v' \in \text{img } \rho_i$ in an environment in the output sequence $\underline{\rho}$.

Proof. By induction on the derivation \mathcal{MS} :

- Cases $\mathcal{MS} = \text{PL-EMP-BOTH}$, $\mathcal{MS} = \text{PL-EMP-PAT}$, $\mathcal{MS} = \text{PL-EMP-VAL}$, $\mathcal{MS} = \text{PL-MORE-STAR-PAT-FAIL}$, and $\mathcal{MS} = \text{PL-MORE-STAR-VAL-FAIL}$ hold trivially since we have $\underline{\rho} = []$ or $\underline{\rho} = \varepsilon$, and there are no values in either $[]$ or ε .

- Case $\mathcal{MS} = \text{PL-MORE-PAT}$ holds by induction hypotheses (including Lemma E.3), and Lemma E.1.

$$x \in \text{dom } \sigma \quad \sigma(x) = \underline{v'} \quad \underline{v} = \underline{v'} \otimes \underline{v''}$$

$$\mathcal{MS}'$$

$$\sigma \vdash \underline{\star p'} := \underline{v''} \mid \emptyset \xrightarrow[\text{match}\star]{\otimes} \underline{\rho}$$

$$\bullet \text{ Case } \mathcal{MS} = \text{PL-MORE-STAR-UNI} \xrightarrow{\quad \quad \quad},$$

$$\sigma \vdash \underline{\star x}, \underline{\star p'} := \underline{v} \mid \mathbb{V} \xrightarrow[\text{match}\star]{\otimes} \underline{\rho}$$

so $\underline{\star p} = \underline{\star x}, \underline{\star p'}$.

By property 4 we can derive that there exists a type sequence $\underline{t''}$ such that $\underline{v''} : \underline{t''}$. Then we can apply induction hypothesis on derivation \mathcal{MS}' , and get our target result. □

We have a lemma on the auxiliary function if-fail:

Lemma E.5. *If we have $v' = \text{if-fail}(vfres, v')$, and:*

1. *If $vfres = \text{success } v$, then we have $v : t$ for some type t*
2. *We have $v'' : t''$ for some type t''*

Then there exists a type t' such that $v' : t'$

Proof. Straightforwardly by case analysis on $vfres$ □

Similarly, we have a lemma on the auxiliary function vcombine

Lemma E.6. *If we have $vfres\star'' = \text{vcombine}(vfres, vfres\star', v, v')$, and:*

1. *If $vfres = \text{success } v'''$, then we have $v''' : t'''$ for some type t*
2. *If $vfres\star' = \text{success } \underline{v''''}$, then we $\underline{v''''} : \underline{t''''}$ for some typing sequence $\underline{t''''}$*
3. *We have $v : t$ for some type t*
4. *We have $\underline{v'} : \underline{t'}$ for some type t'*

Then if $vfres\star'' = \text{success } \underline{v''}$ there exists a type sequence $\underline{t''}$ such that $\underline{v''} : \underline{t''}$

Proof. Straightforwardly by case analysis on $vfres$ and $vfres\star'$ and using Lemma E.5. □

We have a lemma on the reconstruct derivation:

Lemma E.7. *If we have \mathcal{RC} $\text{recons } v \text{ using } \underline{v'} \text{ to } \text{rcres}$, $v : t$ for some type t , and $\underline{v'} : t'$ for some type sequence t' , then when $\text{rcres} = \mathbf{success } \underline{v''}$ there exists a type t'' such that if $\underline{v''} : t''$*

Proof. By induction on the derivation \mathcal{RC} :

- Cases $\mathcal{RC} = \text{RC-VAL-ERR}$, $\mathcal{RC} = \text{RC-CONS-ERR}$, $\mathcal{RC} = \text{RC-LIST-ERR}$, $\mathcal{RC} = \text{RC-SET-ERR}$, $\mathcal{RC} = \text{RC-MAP-ERR}$, $\mathcal{RC} = \text{RC-BOT-ERR}$ hold trivially.
- Case $\mathcal{RC} = \text{RC-VAL-SUCS}$ holds using the premises.
- Case $\mathcal{RC} = \text{RC-CONS-SUCS}$ holds using the premises and rule T-CONSTRUCTOR.
- Case $\mathcal{RC} = \text{RC-LIST-SUCS}$ holds using the premises and rule T-LIST.
- Case $\mathcal{RC} = \text{RC-SET-SUCS}$ holds using the premises and rule T-SET.
- Case $\mathcal{RC} = \text{RC-MAP-SUCS}$ holds using the premises and rule T-MAP.
- Case $\mathcal{RC} = \text{RC-BOT-SUCS}$ holds using the premises and rule T-BOT.

□

We now have a series of mutually inductive lemmas with our Theorem 6.2, since the operational semantics rules are mutually inductive themselves. The lemmas are Lemma E.8, Lemma E.9, Lemma E.10, Lemma E.11, Lemma E.12, Lemma E.13, Lemma E.14, Lemma E.15, Lemma E.16, and Lemma E.17.

Lemma E.8. *If \mathcal{ES} $\underline{e}; \sigma \xrightarrow[\text{expr}^*]{\text{vres}^*} \text{vres}^*; \sigma'$ and there exists a type t such that $v : t$ for each value in the input store $v \in \text{img } \sigma$, then*

1. *There exists a type t' such that $\underline{v'} : t'$ for each value in the result store $\underline{v'} \in \text{img } \sigma'$.*
2. *If the result value vres^* is $\mathbf{success } \underline{v''}$, then there exists a type sequence $\underline{t''}$ such that $\underline{v''} : \underline{t''}$.*

3. If the result value $vres^*$ is either **return** v'' , or **throw** v'' , then there exists a type t'' such that $v'' : t''$

Proof. By induction on the derivation \mathcal{ES} :

- Case $\mathcal{ES} = \text{ES-EMP}$ holds directly using premises.
- Cases $\mathcal{ES} = \text{ES-MORE}$, $\mathcal{ES} = \text{ES-Exc1}$ and $\mathcal{ES} = \text{ES-Exc2}$ hold directly from the induction hypotheses (including the one given by Theorem 6.2).

□

Lemma E.9. If $e; \underline{\rho}; \sigma \xrightarrow[\text{each}]{\mathcal{EE}} vres; \sigma'$, there exists a type t such that $v : t$ for each value in the input store $v \in \text{img } \sigma$, and there exists a type t such that $v : t$ for each value in an environment $v \in \text{img } \rho_i$ in the environment sequence $\underline{\rho}$ then

1. There exists a type t' such that $v' : t'$ for each value in the result store $v' \in \text{img } \sigma'$.
2. If the result value $vres$ is either **success** v'' , **return** v'' , or **throw** v'' , then there exists a type t'' such that $v'' : t''$

Proof. By induction on the derivation \mathcal{EE} :

- Case $\mathcal{EE} = \text{EE-EMP}$ holds directly using the premises.
- Cases $\mathcal{EE} = \text{EE-MORE-SUCS}$, $\mathcal{EE} = \text{EE-MORE-BREAK}$ and $\mathcal{EE} = \text{EE-MORE-EXC}$ hold directly from the induction hypotheses (including Theorem 6.2), and the (trivial) facts that if any store σ is well-typed then $\sigma \setminus \mathbb{X}$ is well-typed for any set of variables \mathbb{X} and $\sigma\rho$ is well-typed for any well-typed environment ρ .

□

Lemma E.10. If $g; \sigma \xrightarrow[\text{gexpr}]{\mathcal{G}} envres; \sigma'$, and there exists a type t such that $v : t$ for each value in the input store $v \in \text{img } \sigma$ then

1. There exists a type t' such that $v' : t'$ for each value in the result store $v' \in \text{img } \sigma'$.

2. If the result value envres is **success** $\underline{\rho}$, then there exists a type t'' for each value in an environment $v'' \in \rho_i$ in the environment sequence $\underline{\rho}$ such that $v'' : t''$
3. If the result value envres is either **return** v''' , or **throw** v''' , then there exists a type t''' such that $v''' : t'''$

Proof. By induction on the derivation \mathcal{G} :

- Cases $\mathcal{G} = \text{G-PAT-SUCS}$ and $\mathcal{G} = \text{G-PAT-EXC}$ hold directly by induction hypothesis given by Theorem 6.2, and then using Lemma E.3 if necessary.
- Cases $\mathcal{G} = \text{G-ENUM-LIST}$, $\mathcal{G} = \text{G-ENUM-SET}$ and $\mathcal{G} = \text{G-ENUM-MAP}$ hold by the induction hypothesis given by Theorem 6.2 and then using inversion on the type derivation of the result collection value to extract the type derivations of the contained values in the result environments (from rules T-LIST, T-SET, and T-MAP respectively).
- Cases $\mathcal{G} = \text{G-ENUM-ERR}$ and $\mathcal{G} = \text{G-ENUM-EXC}$ hold directly by the induction hypothesis given by Theorem 6.2.

□

Lemma E.11. If $\underline{\rho}; e; \sigma \xrightarrow[\text{case}]{\mathcal{C}} \text{vres}; \sigma'$, there exists a type t such that $v : t$ for each value in an environment $v \in \text{img } \rho_i$ in the environment sequence $\underline{\rho}$, and there exists a type t such that $v : t$ for each value in the input store $v \in \text{img } \sigma$, then

1. There exists a type t' such that $v' : t'$ for each value in the result store $v' \in \text{img } \sigma'$.
2. If the result value vres is either **success** v'' , **return** v'' , or **throw** v'' , then there exists a type t'' such that $v'' : t''$

Proof. By induction on derivation \mathcal{C} :

- Case $\mathcal{C} = \text{EC-EMP}$ holds directly using the premises.
- Case $\mathcal{C} = \text{EC-MORE-FAIL}$ and $\mathcal{C} = \text{EC-MORE-ORD}$ hold directly using induction hypotheses (including Theorem 6.2) and the facts of well-typedness of store extension by well-typed environments and well-typedness of variable removals from stores.

□

Lemma E.12. If $\frac{\mathcal{CS}}{\underline{cs}; v; \sigma \xrightarrow[\text{cases}]{\Rightarrow} vres; \sigma'}$, there exists a type t such that $v : t$, and there exists a type t' such that $v' : t'$ for each value in the input store $v' \in \text{img } \sigma$, then

1. There exists a type t'' such that $v'' : t''$ for each value in the result store $v'' \in \text{img } \sigma'$.
2. If the result value $vres$ is either **success** v''' , **return** v''' , or **throw** v''' , then there exists a type t''' such that $v''' : t'''$

Proof. By induction on the derivation \mathcal{CS} :

- Case $\mathcal{CS} = \text{ECS-EMP}$ holds directly using the premises.
- Cases $\mathcal{CS} = \text{ECS-MORE-FAIL}$ and $\mathcal{CS} = \text{ECS-MORE-FAIL}$ hold using Lemma E.3 and the induction hypotheses (including Lemma E.11).

□

Lemma E.13. If $\frac{\mathcal{V}}{\underline{cs}; v; \sigma \xrightarrow[\text{visit}]{\xrightarrow{st}} vres; \sigma'}$, there exists a type t such that $v : t$, and there exists a type t' such that $v' : t'$ for each value in the input store $v' \in \text{img } \sigma$, then

1. There exists a type t'' such that $v'' : t''$ for each value in the result store $v'' \in \text{img } \sigma'$.
2. If the result value $vres$ is either **success** v''' , **return** v''' , or **throw** v''' , then there exists a type t''' such that $v''' : t'''$

Proof. By induction on the derivation \mathcal{V} :

- Cases $\mathcal{V} = \text{EV-TD}$, $\mathcal{V} = \text{EV-TDB}$, $\mathcal{V} = \text{EV-OM-EQ}$, $\mathcal{V} = \text{EV-OM-NEQ}$ and $\mathcal{V} = \text{EV-OM-Exc}$ hold by induction hypotheses (including Lemma E.14).
- Cases $\mathcal{V} = \text{EV-BU}$, $\mathcal{V} = \text{EV-BUB}$, $\mathcal{V} = \text{EV-IM-EQ}$, $\mathcal{V} = \text{EV-IM-NEQ}$ and $\mathcal{V} = \text{EV-IM-Exc}$ hold by induction hypotheses (including Lemma E.16).

□

Lemma E.14. If $\frac{\mathcal{VT}}{\underline{cs}; \underline{v}; \sigma \xrightarrow[\text{td-visit}]{br} \underline{vres}; \sigma'}$, there exists a type t such that $v : t$, and there exists a type t' such that $v' : t'$ for each value in the input store $v' \in \text{img } \sigma$, then

1. There exists a type t'' such that $v'' : t''$ for each value in the result store $v'' \in \text{img } \sigma'$.
2. If the result value $vres$ is either **success** v''' , **return** v''' , or **throw** v''' , then there exists a type t''' such that $v''' : t'''$

Proof. By induction on the derivation \mathcal{VT} :

- All cases ($\mathcal{VT} = \text{ETV-BREAK-SUCS}$, $\mathcal{VT} = \text{ETV-ORD-SUCS1}$, $\mathcal{VT} = \text{ETV-ORD-SUCS2}$, $\mathcal{VT} = \text{ETV-Exc1}$ and $\mathcal{VT} = \text{ETV-Exc2}$) hold by induction hypotheses (including Lemma E.11 and Lemma E.15), Lemma E.5, Lemma E.2 and Lemma E.7.

□

Lemma E.15. If $\frac{\mathcal{VTS}}{\underline{cs}; \underline{v}; \sigma \xrightarrow[\text{td-visit}]{br} \underline{vres} \star; \sigma'}$, there exists a type sequence \underline{t} such that $v : \underline{t}$, and there exists a type t' such that $v' : t'$ for each value in the input store $v' \in \text{img } \sigma$, then

1. There exists a type t'' such that $v'' : t''$ for each value in the result store $v'' \in \text{img } \sigma'$.
2. If the result value $vres \star$ is **success** $\underline{v'''}$ then there exists a type sequence $\underline{t'''} such that $\underline{v''' : t'''}$$
3. If the result value $vres \star$ is either **return** v''' , or **throw** v''' , then there exists a type t''' such that $v''' : t'''$

Proof. By induction on the derivation \mathcal{VTS}

- Case $\mathcal{VTS} = \text{ETVS-EMP}$ holds using the premises.
- Cases $\mathcal{VTS} = \text{ETVS-BREAK}$, $\mathcal{VTS} = \text{ETVS-MORE}$, $\mathcal{VTS} = \text{ETVS-Exc1}$ and $\mathcal{VTS} = \text{ETVS-Exc2}$ holds using the induction hypotheses (including Lemma E.14) and Lemma E.6.

□

Lemma E.16. If $\frac{\mathcal{VB}}{\underline{cs}; \underline{v}; \sigma \xrightarrow[\text{bu-visit}]{br} \underline{vres}; \sigma'}$, there exists a type t such that $\underline{v} : t$, and there exists a type t' such that $\underline{v}' : t'$ for each value in the input store $\underline{v}' \in \text{img } \sigma$, then

1. There exists a type t'' such that $\underline{v}'' : t''$ for each value in the result store $\underline{v}'' \in \text{img } \sigma'$.
2. If the result value \underline{vres} is either **success** \underline{v}''' , **return** \underline{v}''' , or **throw** \underline{v}''' , then there exists a type t''' such that $\underline{v}''' : t'''$

Proof. By induction on the derivation \mathcal{VB} :

- All cases ($\mathcal{VB} = \text{EBU-NO-BREAK-SUCS}$, $\mathcal{VB} = \text{EBU-BREAK-SUCS}$, $\mathcal{VB} = \text{EBU-FAIL-SUCS}$, $\mathcal{VB} = \text{EBU-NO-BREAK-EXC}$, $\mathcal{VB} = \text{EBU-EXC}$, and $\mathcal{VB} = \text{EBU-NO-BREAKERR}$) hold by induction hypotheses (including Lemma E.11 and Lemma E.17), Lemma E.5, Lemma E.2 and Lemma E.7.

□

Lemma E.17. If $\frac{\mathcal{VBS}}{\underline{cs}; \underline{v}; \sigma \xrightarrow[\text{bu-visit}\star]{br} \underline{vres} \star; \sigma'}$, there exists a type sequence \underline{t} such that $\underline{v} : \underline{t}$, and there exists a type t' such that $\underline{v}' : t'$ for each value in the input store $\underline{v}' \in \text{img } \sigma$, then

1. There exists a type t'' such that $\underline{v}'' : t''$ for each value in the result store $\underline{v}'' \in \text{img } \sigma'$.
2. If the result value $\underline{vres} \star$ is **success** \underline{v}''' then there exists a type sequence \underline{t}''' such that $\underline{v}''' : \underline{t}'''$
3. If the result value $\underline{vres} \star$ is either **return** \underline{v}''' , or **throw** \underline{v}''' , then there exists a type t''' such that $\underline{v}''' : t'''$

Proof. By induction on the derivation \mathcal{VBS}

- Case $\mathcal{VBS} = \text{EBUS-EMP}$ holds using the premises.
- Cases $\mathcal{VBS} = \text{EBUS-BREAK}$, $\mathcal{VBS} = \text{EBUS-MORE}$, $\mathcal{VBS} = \text{EBUS-Exc1}$ and $\mathcal{VBS} = \text{EBUS-Exc2}$ holds using the induction hypotheses (including Lemma E.16) and Lemma E.6.

□

Theorem 6.2 (Strong typing). *Assume that semantic unary $\llbracket \ominus \rrbracket$ and binary operators $\llbracket \oplus \rrbracket$ are strongly typed. If $e; \sigma \xrightarrow[\text{expr}]{\mathcal{E}} vres; \sigma'$ and there exists a type*

t such that $\frac{\mathcal{T}}{v : t}$ for each value in the input store $v \in \text{img } \sigma$, then

1. *There exists a type t' such that $v' : t'$ for each value in the result store $v' \in \text{img } \sigma'$.*
2. *If the result value $vres$ is either **success** v'' , **return** v'' , or **throw** v'' , then there exists a type t'' such that $v'' : t''$.*

Proof. By induction on the derivation \mathcal{E} :

- Cases $\mathcal{E} = \text{E-VAL}$, $\mathcal{E} = \text{E-VAR-SUCS}$, $\mathcal{E} = \text{E-VAR-ERR}$, $\mathcal{E} = \text{E-BREAK}$, $\mathcal{E} = \text{E-CONTINUE}$ and $\mathcal{E} = \text{E-FAIL}$ hold directly using the premises.
- Cases $\mathcal{E} = \text{E-UN-EXC1}$, $\mathcal{E} = \text{E-BIN-EXC1}$, $\mathcal{E} = \text{E-BIN-EXC2}$, $\mathcal{E} = \text{E-CONS-ERR}$, $\mathcal{E} = \text{E-CONS-EXC}$, $\mathcal{E} = \text{E-LIST-ERR}$, $\mathcal{E} = \text{E-LIST-EXC}$, $\mathcal{E} = \text{E-SET-ERR}$, $\mathcal{E} = \text{E-SET-EXC}$, $\mathcal{E} = \text{E-MAP-ERR}$, $\mathcal{E} = \text{E-MAP-EXC}$, $\mathcal{E} = \text{E-LOOKUP-ERR}$, $\mathcal{E} = \text{E-LOOKUP-EXC1}$, $\mathcal{E} = \text{E-LOOKUP-EXC2}$, $\mathcal{E} = \text{E-UPDATE-ERR1}$, $\mathcal{E} = \text{E-UPDATE-ERR2}$, $\mathcal{E} = \text{E-UPDATE-EXC1}$, $\mathcal{E} = \text{E-UPDATE-EXC2}$, $\mathcal{E} = \text{E-UPDATE-EXC3}$, $\mathcal{E} = \text{E-CALL-ARG-ERR}$, $\mathcal{E} = \text{E-CALL-ARG-EXC}$, $\mathcal{E} = \text{E-RET-SUCS}$, $\mathcal{E} = \text{E-RET-EXC}$, $\mathcal{E} = \text{E-ASGN-ERR}$, $\mathcal{E} = \text{E-ASGN-EXC}$, $\mathcal{E} = \text{E-IF-TRUE}$, $\mathcal{E} = \text{E-IF-FALSE}$, $\mathcal{E} = \text{E-IF-ERR}$, $\mathcal{E} = \text{E-IF-EXC}$, $\mathcal{E} = \text{E-SWITCH-SUCS}$, $\mathcal{E} = \text{E-SWITCH-EXC1}$, $\mathcal{E} = \text{E-SWITCH-EXC2}$, $\mathcal{E} = \text{E-VISIT-SUCS}$, $\mathcal{E} = \text{E-VISIT-FAIL}$, $\mathcal{E} = \text{E-VISIT-EXC1}$, and $\mathcal{E} = \text{E-VISIT-EXC2}$, $\mathcal{E} = \text{E-FOR-SUCS}$, $\mathcal{E} = \text{E-FOR-EXC}$, $\mathcal{E} = \text{E-WHILE-TRUE-SUCS}$, $\mathcal{E} = \text{E-WHILE-EXC1}$, $\mathcal{E} = \text{E-WHILE-EXC2}$, $\mathcal{E} = \text{E-WHILE-ERR}$, $\mathcal{E} = \text{E-SOLVE-EQ}$, $\mathcal{E} = \text{E-SOLVE-NEQ}$, $\mathcal{E} = \text{E-SOLVE-EXC}$, $\mathcal{E} = \text{E-SOLVE-ERR}$, $\mathcal{E} = \text{E-THR-SUCS}$, $\mathcal{E} = \text{E-THR-ERR}$, $\mathcal{E} = \text{E-FIN-SUCS}$, $\mathcal{E} = \text{E-FIN-EXC}$, and $\mathcal{E} = \text{E-TRY-ORD}$ hold by induction hypotheses (including Lemma E.8, Lemma E.10, Lemma E.12 and Lemma E.13).
- Case $\mathcal{E} = \text{E-UN-SUCS}$ holds by its induction hypothesis and strong typing of $\llbracket \ominus \rrbracket$.

- Case $\mathcal{E} = \text{E-BIN-SUCS}$ holds by its induction hypothesis and strong typing of $\llbracket \oplus \rrbracket$.
- Case $\mathcal{E} = \text{E-CONS-SUCS}$ holds by its induction hypothesis given by Lemma E.8, and then by using T-CONSTRUCTOR with the provided typing premises.
- Case $\mathcal{E} = \text{E-LIST-SUCS}$ holds by its induction hypothesis given by Lemma E.8, and then by using T-LIST.
- Case $\mathcal{E} = \text{E-SET-SUCS}$ holds by its induction hypothesis given by Lemma E.8, and then by using T-SET.
- Case $\mathcal{E} = \text{E-MAP-SUCS}$ holds by its induction hypothesis given by Lemma E.8, and then by using T-MAP.
- Case $\mathcal{E} = \text{E-LOOKUP-SUCS}$ holds by its induction hypotheses and inversion of \mathcal{T} to T-MAP.
- Case $\mathcal{E} = \text{E-LOOKUP-NOKEY}$ holds by its induction hypotheses and using T-CONS on the definition of NoKey.
- Case $\mathcal{E} = \text{E-UPDATE-SUCS}$ holds by its induction hypotheses and by using T-MAP to reconstruct the type derivation for the result.
- Cases $\mathcal{E} = \text{E-CALL-SUCS}$, $\mathcal{E} = \text{E-CALL-RES-EXC}$, $\mathcal{E} = \text{E-CALL-RES-ERR1}$, and $\mathcal{E} = \text{E-CALL-RES-ERR2}$ hold by induction hypotheses (including Lemma E.8) and by using the fact that extracting variables from and extending well-typed stores produces well-typed stores.
- Cases $\mathcal{E} = \text{E-ASGN-SUCS}$ and $\mathcal{E} = \text{E-TRY-CATCH}$ hold by induction hypotheses and the fact that extending well-typed stores with well-typed environments preserves well-typedness.
- Cases $\mathcal{E} = \text{E-SWITCH-FAIL}$, $\mathcal{E} = \text{E-WHILE-FALSE}$, and $\mathcal{E} = \text{E-WHILE-TRUE-BREAK}$ hold by induction hypotheses and using T-VOID.
- Cases $\mathcal{E} = \text{E-BLOCK-SUCS}$, $\mathcal{E} = \text{E-BLOCK-EXC}$ hold by the induction hypothesis given by Lemma E.8 and the fact that removing variables from a well-typed store produces a well-typed store.

□

We will for Theorem 6.3, also need to state some lemmas. We will for the proof focus mainly on cases where the induction hypothesis does not **timeout**, since if it does it is trivially possible to construct a **timeout** derivation for the result syntax.

First, we need a lemma that specifies that the sequence of values produced by the children is strictly smaller than the input value. Let $v \prec v'$ denote the relation that v is syntactically contained in v' , then our target property is specified in Lemma E.18.

Lemma E.18. *If $\underline{v'} = \text{children}(v)$ then $v'_i \prec v$ for all i .*

Proof. Directly by induction on v . □

We have a lemma for progress on reconstruction:

Lemma E.19. *It is possible to construct a derivation **recons** v using $\underline{v'}$ to **rcres** for any well-typed value v and well-typed value sequence $\underline{v'}$.*

Proof. Straightforwardly by case analysis on v . □

We have two mutually recursive lemmas for pattern matching: Lemma E.20 and Lemma E.21.

Lemma E.20. *It is possible to construct a derivation $\sigma \vdash p \stackrel{?}{:=} v \xrightarrow[\text{match}]{} \underline{\rho}$ for any pattern p , well-typed value v , well-typed store σ .*

Proof. By induction on syntax of p :

- Case $p = vb$: We proceed by testing whether v is equal to vb :
 - Case $v = vb$ then use P-VAL-SUCS.
 - Case $v \neq vb$ then use P-VAL-FAIL.
- Case $p = x$: We proceed by testing whether x is in $\text{dom } \sigma$
 - Case $x \in \text{dom } \sigma$: we proceed to test whether v is equal to $\sigma(x)$
 - * Case $v = \sigma(x)$ then use P-VAR-UNI.
 - * Case $v \neq \sigma(x)$ then use P-VAR-FAIL.
 - Case $x \notin \text{dom } \sigma$ then use P-VAR-BIND.

- Case $p = k(\underline{p}')$: We proceed to test whether v is equal to $k(\underline{v}')$ for some \underline{v}'
 - Case $v = k(\underline{v}')$ then use P-CONS-SUCS using the induction hypotheses of \underline{p}' with \underline{v}' .
 - Case $v \neq k(\underline{v}')$ then use P-CONS-FAIL.
- Case $p = t x : p'$: From our premise we know that v is well-typed, i.e. that there exists a t' such that $v : t'$. We proceed to test whether $t' <: t$.
 - Case $t' <: t$ then use P-TYPE-SUCS using the induction hypothesis on \underline{p}' with v .
 - Case $t' \not<: t$ then use P-TYPE-FAIL.
- Case $p = [\star p']$:
We proceed to test whether $v = [\underline{v}']$ for some value sequence \underline{v}' .
 - Case $v = [\underline{v}']$ then use P-LIST-SUCS induction hypothesis given by Lemma E.21 on $\star p'$ with \underline{v}'
 - Case $v \neq [\underline{v}']$ then use P-LIST-FAIL
- Case $p = \{\star p'\}$:
We proceed to test whether $v = \{\underline{v}'\}$ for some value sequence \underline{v}' .
 - Case $v = \{\underline{v}'\}$ then use P-SET-SUCS induction hypothesis given by Lemma E.21 on $\star p'$ with \underline{v}'
 - Case $v \neq \{\underline{v}'\}$ then use P-SET-FAIL
- Case $p = /p'$:
 - Using the induction hypothesis on p' with v , we get $\sigma \vdash p' \stackrel{?}{=} v \xrightarrow[\text{match}]{\quad} \underline{\rho}$.
 - Now, let $\underline{v}' = \text{children}(v)$. In order to handle the self-recursive calls using $/p'$, we proceed by inner well-founded induction on the relation \prec using value sequence \underline{v}' :
 - * Using the inner induction hypothesis we get derivations $\sigma \vdash /p' \stackrel{?}{=} v'_i \xrightarrow[\text{match}]{\quad} \underline{\rho}_i$ for all i

- Finally, we use P-DEEP on the derivations we got from the outer and inner induction hypotheses.

□

Lemma E.21. *It is possible to construct a derivation $\sigma \vdash \underline{\star p} \stackrel{?}{=} v \mid \mathbb{V} \xrightarrow[\text{match}\star]{() \otimes} \underline{\rho}$ for any pattern p , well-typed value v , well-typed store σ , well-typed visited value set \mathbb{V} , type-preserving partition operator \otimes .*

Proof. By induction on the syntax of $\underline{\star p}$:

- Case $\underline{\star p} = \varepsilon$
By case analysis on \underline{v} :
 - Case $\underline{v} = \varepsilon$ then use PL-EMP-BOTH
 - Case $\underline{v} = \underline{v'}, \underline{v''}$ then use PL-EMP-PAT
- Case $\underline{\star p} = \underline{p'}, \underline{\star p''}$:
By case analysis on \underline{v} :
 - Case $\underline{v} = \varepsilon$ then use PL-EMP-VAL
 - Case $\underline{v} = \underline{v'}, \underline{v''}$ then use PL-MORE-PAT using induction hypotheses (including Lemma E.20).
- Case $\underline{\star p} = \underline{\star x}, \underline{\star p''}$: We proceed to test whether x is in $\text{dom } \sigma$:
 - Case $x \in \text{dom } \sigma$:
 - Case $x \notin \text{dom } \sigma$:
Because of backtracking we need to do an inner induction to handle the cases which recurse to the same star patten sequence. Let $\mathbb{V}_{\text{all}} = \{ \underline{v'} \mid \exists \underline{v''}. \underline{v} = \underline{v'} \otimes \underline{v''} \}$, then we proceed by inner induction on $|\mathbb{V}_{\text{all}} - \mathbb{V}|$
 - * Case $|\mathbb{V}_{\text{all}} - \mathbb{V}| = 0$:
Then we have $\mathbb{V} = \mathbb{V}_{\text{all}}$ and so the only applicable rule is PL-MORE-STAR-EXH since \mathbb{V}_{all} covers all partitions and thus $\nexists \underline{v'}, \underline{v''}. \underline{v} = \underline{v'} \otimes \underline{v''} \wedge \underline{v'} \notin \mathbb{V}$ holds.

* Case $|\mathbb{V}_{\text{all}} - \mathbb{V}| > 0$:

Then there exists $\mathbb{V}' \neq \emptyset$ such that $\mathbb{V}_{\text{all}} = \mathbb{V}' \uplus \mathbb{V}$ and we can pick $\underline{v}' \in \mathbb{V}'$ such that $\underline{v} = \underline{v}' \uplus \underline{v}''$ for some \underline{v}'' .

We can now use PL-MORE-STAR-RE by using the outer hypothesis on $\star p$ with \underline{v}'' and \emptyset , and the inner hypothesis on $\mathbb{V} \cup \underline{v}'$ (since the size decreases by 1)

□

We now have a series of lemmas that are mutually recursive with Theorem 6.3: Lemma E.22, Lemma E.23, Lemma E.24, Lemma E.25, Lemma E.26, Lemma E.27, Lemma E.28, Lemma E.29, Lemma E.30, and Lemma E.31.

Lemma E.22. *It is possible to construct a derivation $\underline{e}; \sigma \xrightarrow[\text{expr}\star]{n} \underline{vres} \star; \sigma'$, for any expression sequence \underline{e} , well-typed store σ and fuel n .*

Proof. By induction on n :

- Case $n = 0$ then use corresponding **timeout**-rule.
- Case $n > 0$, then by case analysis on the expression sequence \underline{e} :
 - Case $\underline{e} = \varepsilon$ then use ES-EMP.
 - Case $\underline{e} = \underline{e}', \underline{e}''$:

* Using induction hypothesis given by Theorem 6.3 on $n - 1$ with \underline{e}' and σ we get a derivation $\underline{e}'; \sigma \xrightarrow[\text{expr}]{n-1} \underline{vres}'; \sigma''$

By case analysis on \underline{vres}' :

- Case $\underline{vres}' = \mathbf{success} \underline{v}'$:
By Theorem 6.2 we get that σ'' is well-typed.
By induction hypothesis on $n - 1$ with \underline{e}'' and σ'' we get a derivation $\underline{e}''; \sigma'' \xrightarrow[\text{expr}\star]{n-1} \underline{vres} \star''; \sigma'$
By case analysis on $\underline{vres} \star''$:
 - Case $\underline{vres} \star'' = \mathbf{success} \underline{v}''$ then use ES-MORE
 - Case $\underline{vres} \star'' = \mathbf{exres}$ then use ES-Exc2
- Case $\underline{vres}' = \mathbf{exres}$ then use ES-Exc1

□

Lemma E.23. *It is possible to construct a derivation $e; \underline{\rho}; \sigma \xrightarrow[\text{each}]{n} vres; \sigma'$ for any expression e , well-typed environment sequence $\underline{\rho}$, well-typed store σ and fuel n .*

Proof. By induction on n :

- Case $n = 0$ then use the corresponding **timeout**-derivation.
- Case $n > 0$:

By case analysis on the environment sequence $\underline{\rho}$:

- Case $\underline{\rho} = \varepsilon$ then use EE-EMP
- Case $\underline{\rho} = \rho', \underline{\rho}''$:

By induction hypothesis given by Theorem 6.3 on $n - 1$ with e and $\sigma\rho$ we get a derivation $e; \sigma\rho \xrightarrow[\text{expr}]{n-1} vres'; \sigma''$.

By case analysis on $vres'$:

- * Cases $vres' = \mathbf{success} v'$ and $vres' = \mathbf{continue}$ then use EE-MORE-SUCS with above derivation and the induction hypothesis on $n - 1$ with $\underline{\rho}''$, e , and $\sigma''' = \sigma'' \setminus \text{dom } \rho'$.
- * Cases $vres' = \mathbf{break}$ then use EE-MORE-BREAK
- * Cases $vres' = \mathbf{throw} v'$, $vres' = \mathbf{return} v'$, $vres' = \mathbf{fail}$, and $vres' = \mathbf{error}$ then use EE-MORE-EXC

□

Lemma E.24. *It is possible to construct a derivation $g; \sigma \xrightarrow[\text{gexpr}]{n} envres; \sigma'$ for any generator expression g , well-typed store σ and fuel n .*

Proof. By induction on n :

- Case $n = 0$ then use the corresponding **timeout**-derivation.
- Case $n > 0$:

By case analysis on g :

- Case $g = p := e$:

By induction hypothesis given by Theorem 6.3 on $n - 1$ with e and σ we get a derivation $e; \sigma \xrightarrow[\text{expr}]{n-1} vres'; \sigma'$

By case analysis on $vres'$:

- * Case $vres' = \mathbf{success} \ v'$ then use G-PAT-SUCS with above derivation and the derivation from applying Lemma E.20 on p with σ' and v' .
- * Case $vres' = \mathbf{exres}$ then use G-PAT-EXC:
- Case $g = x \leftarrow e$:
By induction hypothesis given by Theorem 6.3 on $n - 1$ with e and σ we get a derivation $e; \sigma \xRightarrow[\text{expr}]{n-1} vres'; \sigma'$
By case analysis on $vres'$
 - * Case $vres' = \mathbf{success} \ v'$
By case analysis on v' :
 - Case $v' = \underline{v''}$ then use *G – Enum – List*
 - Case $v' = \{\underline{v''}\}$ then use *G – Enum – Set*
 - Case $v' = (\underline{v''} : v''')$ then use *G – Enum – Map*
 - Cases $v' = vb$, $v' = k(\underline{v''})$ and $v' = \blacksquare$ then use *G-ENUM-ERR*
 - * Case $vres' = \mathbf{exres}$ then use G-ENUM-EXC

□

Lemma E.25. *It is possible to construct a derivation $\underline{\rho}; e; \sigma \xRightarrow[\text{case}]{n} vres'; \sigma'$ for any well-typed environment sequence $\underline{\rho}$, expression e , well-typed store σ and fuel n .*

Proof. By induction on n :

- Case $n = 0$ then use the corresponding **timeout**-derivation.
- Case $n > 0$:

By case analysis on the environment sequence $\underline{\rho}$

- Case $\underline{\rho} = 0$ then use EC-EMP.
- Case $\underline{\rho} = \rho', \underline{\rho''}$:

Using the induction hypothesis given by Theorem 6.3 on $n - 1$ with e and σ we get a derivation $e; \sigma \xRightarrow[\text{expr}]{n-1} vres'; \sigma'$.

By case analysis on $vres'$

- * Case $vres' = \mathbf{fail}$ then use EC-MORE-FAIL on above derivation and the derivation from applying the induction hypothesis on $n - 1$ with $\underline{\rho''}$, e and σ .

* Case $vres' \neq \mathbf{fail}$ then use EC-MORE-ORD.

□

Lemma E.26. *It is possible to construct a derivation $\underline{cs}; v; \sigma \xrightarrow[\text{cases}]^n vres; \sigma'$ for any case sequence \underline{cs} , well-typed value v , well-typed store σ and fuel n .*

Proof. By induction on n :

- Case $n = 0$ then use the corresponding **timeout**-derivation.
- Case $n > 0$:

By case analysis on the case sequence \underline{cs}

– Case $\underline{cs} = \varepsilon$ then use ECS-EMP

– Case $\underline{cs} = \mathbf{case} p \Rightarrow e, \underline{cs}'$:

Using Lemma E.20 on p with v and σ gives us a derivation $\sigma \vdash p \stackrel{?}{:=} v \xrightarrow[\text{match}]{} \underline{\rho}$. By Lemma E.3 we know that $\underline{\rho}$ is well-typed.

Using induction hypothesis given by Lemma E.25 on $n - 1$ with $\underline{\rho}$, e and σ we get a derivation $\underline{\rho}; e; \sigma \xrightarrow[\text{case}]^{n-1} vres'; \sigma'$.

By case analysis on $vres'$:

- * Case $vres' = \mathbf{fail}$ then use ECS-MORE-FAIL using above derivation and the derivation given by the induction hypothesis on $n - 1$ with \underline{cs}' , v and σ .
- * Case $vres' \neq \mathbf{fail}$ then use ECS-MORE-ORD

□

Lemma E.27. *It is possible to construct a derivation $\underline{cs}; v; \sigma \xrightarrow[\text{visit}]^{st} vres; \sigma'$ for any case sequence \underline{cs} , well-typed value v , well-typed store σ , traversal strategy st and fuel n .*

Proof. By induction on n :

- Case $n = 0$ then use the corresponding **timeout**-derivation.
- Case $n > 0$:

By case analysis on syntax of st

- Case $st = \mathbf{top-down}$ then use EV-TD with the derivation from the induction hypothesis given by Lemma E.28 on $n - 1$ with $\underline{cs}, v, \sigma$ and **no-break**.
- Case $st = \mathbf{top-down-break}$ then use EV-TDB with the derivation from the induction hypothesis given by Lemma E.28 on $n - 1$ with $\underline{cs}, v, \sigma$ and **break**.
- Case $st = \mathbf{bottom-up}$ then use EV-BU with the derivation from the induction hypothesis given by Lemma E.30 on $n - 1$ with $\underline{cs}, v, \sigma$ and **no-break**.
- Case $st = \mathbf{bottom-up-break}$ then use EV-BUB with the derivation from the induction hypothesis given by Lemma E.30 on $n - 1$ with $\underline{cs}, v, \sigma$ and **break**.

- Case $st = \mathbf{outermost}$:

Using the induction hypothesis by Lemma E.28 on $n - 1$ with $\underline{cs}, v, \sigma$ and **no-break** we get a derivation $\underline{cs}; v; \sigma \xrightarrow[\text{td-visit}]{\mathbf{no-break}^{n-1}} vres'; \sigma''$. We know the well-typedness of the output components from Lemma E.14.

By case analysis on $vres'$:

- * Case $vres' = \mathbf{success} v'$:

We proceed by checking whether $v = v'$:

- Case $v = v'$ then use EV-OM-EQ using above derivation.
- Case $v \neq v'$ then use EV-OM-NEQ using above derivation and the derivation from the induction hypothesis given by Lemma E.28 on $n - 1$ with $\underline{cs}, v', \sigma''$ and **no-break**.

- * Case $vres' = \mathbf{exres}$ then use EV-OM-Exc using above derivation.

- Case $st = \mathbf{innermost}$: Using the induction hypothesis by Lemma E.30 on $n - 1$ with $\underline{cs}, v, \sigma$ and **no-break** we get a derivation $\underline{cs}; v; \sigma \xrightarrow[\text{bu-visit}]{\mathbf{no-break}^{n-1}} vres'; \sigma''$. We know the well-typedness of the output components from Lemma E.16.

By case analysis on $vres'$:

- * Case $vres' = \mathbf{success} v'$:

We proceed by checking whether $v = v'$:

- Case $v = v'$ then use EV-IM-EQ using above derivation.
- Case $v \neq v'$ then use EV-IM-NEQ using above derivation and the derivation from the induction hypothesis given by Lemma E.30 on $n - 1$ with \underline{cs} , v' , σ'' and **no-break**.
- * Case $vres' = exres$ then use EV-IM-Exc using above derivation.

□

Lemma E.28. *It is possible to construct a derivation $\underline{cs}; v; \sigma \xrightarrow[\text{td-visit}]{br}^n vres; \sigma'$ for any \underline{cs} , well-typed value v , well-typed store σ , breaking strategy br and fuel n .*

Proof. By induction on n :

- Case $n = 0$ then use the corresponding **timeout**-derivation.
- Case $n > 0$:

Applying the induction hypothesis given by Lemma E.26 on $n - 1$ with cs , v and σ we get a derivation $\underline{cs}; v; \sigma \xrightarrow[\text{cases}]{\mathcal{CS}}^{n-1} vres'; \sigma''$

We proceed to check whether $vres'$ has syntax $vfres'$ for some $vfres'$:

- Case $vres' = vfres'$:

We proceed to check whether $vfres' = \mathbf{success} v'$ and $br = \mathbf{break}$

- * Case $vfres' = \mathbf{success} v'$ and $br = \mathbf{break}$ then use ETV-BREAK-SUCS with \mathcal{CS} .

- * Case $vfres' \neq \mathbf{success} v'$ or $br \neq \mathbf{break}$:

By Boolean logic, we have that $br \neq \mathbf{break} \Rightarrow vfres' = \mathbf{fail}$ is satisfied.

Let $v'' = \text{if-fail}(vfres', v)$ and $v''' = \text{children}(v'')$ (all well-typed from Lemma E.5 and Lemma E.2).

By induction hypothesis given by Lemma E.15 on $n - 1$ with \underline{cs} , v''' and σ'' , we get a derivation

$$\underline{cs}; \underline{v'''}; \sigma'' \xrightarrow[\text{td-visit}]{br}^{n-1} vres''; \sigma' .$$

We know that the results are well typed from Lemma E.15.

By case analysis on $vres''$:

- Case $vres'' = \mathbf{success} \ \underline{v''''}$ then use ETV-ORD-SUCS2 with \underline{CS} , \underline{CS}' and the result derivation from applying Lemma E.19.
 - Case $vres'' = \mathbf{fail}$ then use ETV-ORD-SUCS1 with \underline{CS} , \underline{CS}' .
 - Case $vres'' = exres$ where $exres \neq \mathbf{fail}$ then use ETV-Exc2 with \underline{CS} , \underline{CS}' .
- Case $vres' \neq vres''$: Necessarily, we then have that $vres' = exres$ where $exres \neq \mathbf{fail}$ and so we use ETV-Exc1 with \underline{CS} .

□

Lemma E.29. *It is possible to construct a derivation $\underline{cs}; \underline{v}; \sigma \xrightarrow[\text{td-visit}]{br}^n vres \star ; \sigma'$ for any \underline{cs} , well-typed value sequence \underline{v} , well-typed store σ , breaking strategy br and fuel n .*

Proof. By induction on n :

- Case $n = 0$ then use the corresponding **timeout**-derivation.
- Case $n > 0$: By case analysis on \underline{v} :
 - Case $\underline{v} = \varepsilon$ then use ETVS-EMP.
 - Case $\underline{v} = v', \underline{v}''$:

By induction hypothesis given by Lemma E.28 on $n - 1$ with

$$\underline{cs}, v \text{ and } \sigma \text{ we get a derivation } \underline{cs}; \underline{v}; \sigma \xrightarrow[\text{td-visit}]{\mathcal{VT} \ br}^{n-1} vres''; \sigma'' .$$

We proceed by checking whether $vres''$ is syntactically a $vres''$ for some $vres''$:

- * Case $vres'' = vres''$:

We proceed by checking whether $vres'' = \mathbf{success} \ v'''$ and $br = \mathbf{break}$:

- Case $vres'' = \mathbf{success} \ v'''$ and $br = \mathbf{break}$ then use ETVS-BREAK with \mathcal{VT} .

- Case $vfres'' \neq \mathbf{success}$ v''' or $br \neq \mathbf{break}$:
By Boolean logic, we have that $br \neq \mathbf{break} \Rightarrow vfres'' = \mathbf{fail}$ is satisfied.
By induction hypothesis on $n - 1$ with \underline{cs} , \underline{v}'' and σ'' we get a derivation

$$\begin{array}{c} \mathcal{VTS} \\ \underline{cs}; \underline{v}''; \sigma'' \xrightarrow[\text{td-visit}^*]{br}^{n-1} vres\star'''; \sigma' \end{array}$$

By case analysis on $vres\star'''$:

- Case $vres\star''' = vfres\star'''$ then use ETVS-MORE with \mathcal{VT} and \mathcal{VTS} .
 - Case $vres\star''' = exres$ where $exres \neq \mathbf{fail}$ then use ETVS-Exc2 with \mathcal{VT} and \mathcal{VTS} .
- * Case $vres'' \neq vfres''$:
Here we then have that $vres'' = exres$, where $exres \neq \mathbf{fail}$ and so we use ETVS-Exc1.

□

Lemma E.30. *It is possible to construct a derivation $\underline{cs}; v; \sigma \xrightarrow[\text{bu-visit}]{br}^n vres; \sigma'$ for any \underline{cs} , well-typed value v , well-typed store σ , breaking strategy br and fuel n .*

Proof. By case analysis on n :

- Case $n = 0$ then use the corresponding **timeout**-derivation.
- Case $n > 0$:

Let $\underline{v}'' = \text{children}(v)$. We know that \underline{v}'' is well-typed from Lemma E.2.

By induction hypothesis given by Lemma E.31 on $n - 1$ with \underline{cs} , \underline{v}''

and σ we get a derivation $\begin{array}{c} \mathcal{VBS} \\ \underline{cs}; \underline{v}''; \sigma \xrightarrow[\text{bu-visit}^*]{br}^{n-1} vres''; \sigma'' \end{array}$. Recall

that the output is well-typed from Lemma E.17.

By case analysis on $vres\star''$:

- Case $vres\star'' = vfres''$:
We proceed by case analysis on $vfres''$

- * Case $\underline{vfres}'' = \mathbf{success} \underline{v}'''$: We proceed by case analysis on br :
 - Case $br = \mathbf{break}$ then use EBU-BREAK-SUCS with \mathcal{VBS} and the derivation from applying Lemma E.19 on v and \underline{v}''' .
 - Case $br = \mathbf{no-break}$:
By applying Lemma E.19 on v and \underline{v}''' we get a derivation

$$\text{recons } v \text{ using } \underline{v}''' \text{ to } rcres \quad \mathcal{RC}$$
 By case analysis on $rcres$:
 - Case $rcres = \mathbf{success} v'$:
By induction hypothesis given by Lemma E.26 on $n - 1$ with \underline{cs} , v' and σ'' we get a derivation

$$\underline{cs}; v'; \sigma'' \xrightarrow[\text{cases}]{n-1} vres'; \sigma' \quad \mathcal{CS}$$
 By case analysis on $vres'$:
 - Case $vres' = \underline{vfres}'$ then use EBU-NO-BREAK-SUCS with \mathcal{VBS} , \mathcal{RC} and \mathcal{CS} .
 - Case $exres$ then use EBU-NO-BREAK-EXC with \mathcal{VBS} , \mathcal{RC} and \mathcal{CS} .
 - Case $rcres = \mathbf{error}$ then use EBU-NO-BREAK-ERR.
- * Case $\underline{vfres}'' = \mathbf{fail}$ then use EBU-FAIL-SUCS with the induction hypothesis given by Lemma E.26 on $n - 1$ with \underline{cs} , v and σ'' .
- Case $\underline{vres}\star'' = exres$ where $exres \neq \mathbf{fail}$ then use EBU-EXC with \mathcal{VBS} .

□

Lemma E.31. *It is possible to construct a derivation $\underline{cs}; \underline{v}; \sigma \xrightarrow[\text{bu-visit}\star]{br}^n vres \star ; \sigma'$ for any \underline{cs} , well-typed value sequence \underline{v} , well-typed store σ , breaking strategy br and fuel n .*

Proof. By induction on n :

- Case $n = 0$ then use the corresponding **timeout**-derivation.
- Case $n > 0$:
By case analysis on \underline{v} :

- Case $\underline{v} = \varepsilon$ then use EBUS-EMP.
- Case $\underline{v} = \underline{v'}, \underline{v''}$: By induction hypothesis given by Lemma E.30 on $n - 1$ with $\underline{cs}, \underline{v'}$ and σ we get a derivation

$$\frac{\mathcal{VB}}{\underline{cs}; \underline{v'}; \sigma \xrightarrow[\text{bu-visit}]{br}^{n-1} \underline{vres''}; \sigma''} .$$

By case analysis on $\underline{vres''}$:

* Case $\underline{vres''} = \underline{vfres''}$:

By case analysis on $\underline{vfres''}$ and br :

- Case $\underline{vfres''} = \mathbf{success} \underline{v''}$ and $br = \mathbf{break}$: then use EBUS-BREAK with \mathcal{VB} .
- Case $\underline{vfres''} \neq \mathbf{success} \underline{v''}$ or $br \neq \mathbf{break}$:
By Boolean logic we have $br = \mathbf{break} \Rightarrow \underline{vfres} = \mathbf{fail}$.
By induction hypothesis on $n - 1$ with $\underline{cs}, \underline{v''}$ and σ we

$$\text{get a derivation } \frac{\mathcal{VBS}}{\underline{cs}; \underline{v''}; \sigma'' \xrightarrow[\text{bu-visit}]{br}^{n-1} \underline{vres\star'}; \sigma' .}$$

By case analysis on $\underline{vres\star'}$:

- Case $\underline{vres\star'} = \underline{vfres\star'}$ then use EBUS-MORE with \mathcal{VB} and \mathcal{VBS}
 - Case $\underline{vres\star'} = \underline{exres}$ then use EBUS-Exc2 with \mathcal{VB} and \mathcal{VBS}
- * Case $\underline{vres''} = \underline{exres}$ then use EBUS-Exc1 with \mathcal{VB} .

□

Theorem 6.3 (Partial progress). *It is possible to construct a derivation $e; \sigma \xrightarrow[\text{expr}]{\implies^n} \underline{vres}; \sigma'$ for any input expression e , well-typed store σ and fuel n .*

Proof. By induction on n :

- Case $n = 0$ then use the corresponding **timeout**-derivation.
- Case $n > 0$:

By case analysis on syntax e :

- Case $e = vb$ then use T-BASIC.

– Case $e = x$

We proceed by checking $x \in \text{dom } \sigma$:

* Case $x \in \text{dom } \sigma$ then use E-VAR-SUCS.

* Case $x \notin \text{dom } \sigma$ then use E-VAR-ERR.

– Case $e = \ominus e'$:

By induction hypothesis on $n - 1$ with e' and σ we get a derivation $e'; \sigma \xrightarrow[\text{expr}]{}^{n-1} vres'; \sigma'$.

By case analysis on $vres'$:

* Case $vres' = \mathbf{success}v$ then use E-UN-SUCS with above derivation.

* Case $vres' = exres$ then use E-UN-EXC with above derivation.

– Case $e = e_1 \oplus e_2$: By induction hypothesis on $n - 1$ with e_1

and σ we get a derivation $e_1; \sigma \xrightarrow[\text{expr}]{}^{n-1} vres_1; \sigma'' \cdot$

By case analysis on $vres_1$:

* Case $vres_1 = \mathbf{success} v_1$: By induction hypothesis on $n - 1$ with e_2 and σ'' we get a derivation

$e_2; \sigma'' \xrightarrow[\text{expr}]{}^{n-1} vres_2; \sigma' \cdot$

By case analysis on $vres_2$:

· Case $vres_2 = \mathbf{success} v_2$ then use E-BIN-SUCS with \mathcal{E}_1 and \mathcal{E}_2 .

· Case $vres_2 = exres$ then use E-BIN-EXC2 with \mathcal{E}_1 and \mathcal{E}_2 .

* Case $vres_1 = exres$ then use E-BIN-EXC1 with \mathcal{E}_1 .

– Case $e = k(\underline{e}')$:

Recall that all derivations in our paper are assumed to be well-scoped so there must exist a corresponding data-type at that has $k(\underline{t})$ as a constructor.

By using induction hypothesis given by Lemma E.22 on $n - 1$

with \underline{e}' and σ we get a derivation $\underline{e}'; \sigma \xrightarrow[\text{expr}^*]{}^{n-1} vres^*; \sigma' \cdot$

By case analysis on $vres^*$:

* Case $vres\star' = \mathbf{success} \ \underline{v}'$:

By Lemma E.8 we know that \underline{v}' is well-typed, i.e. that we have $\underline{v}' : t'$ for some type sequence t' . We proceed to check whether all values in the sequence are non- \blacksquare and each have a type t'_i that is a subtype of the target type t_i .

- Case $\underline{v} \neq \blacksquare$ and $t' <: t$ then use E-CONS-SUCS with \mathcal{ES} and the required typing derivations.
- Case $v_i = \blacksquare$ or $t'_i <: t_i$ for some i then use E-CONS-ERR with \mathcal{ES} .

* Case $vres\star' = exres$ then use E-CONS-EXC with \mathcal{ES}

– Case $e = [\underline{e}']$:

By induction hypothesis given by Lemma E.22 on $n - 1$ with \underline{e}' and σ we get a derivation $\underline{e}'; \sigma \xrightarrow[\text{expr}\star]{n-1} vres\star'; \sigma'$.

By case analysis on $vres\star'$:

* Case $vres\star' = \mathbf{success} \ \underline{v}$:

We proceed by checking whether all values \underline{v} are non- \blacksquare :

- Case $\underline{v} \neq \blacksquare$ then use E-LIST-SUCS with above derivation.
- Case $v_i \neq \blacksquare$ for some i then use E-LIST-ERR with above derivation.

* Case $vres\star' = exres$ then use E-LIST-EXC with above derivation.

– Case $e = \{\underline{e}'\}$

By induction hypothesis given by Lemma E.22 on $n - 1$ with \underline{e}' and σ we get a derivation $\underline{e}'; \sigma \xrightarrow[\text{expr}\star]{n-1} vres\star'; \sigma'$.

By case analysis on $vres\star'$:

* Case $vres\star' = \mathbf{success} \ \underline{v}$:

We proceed by checking whether all values \underline{v} are non- \blacksquare :

- Case $\underline{v} \neq \blacksquare$ then use E-SET-SUCS with above derivation.
- Case $v_i \neq \blacksquare$ for some i then use E-SET-ERR with above derivation.

* Case $vres\star' = exres$ then use E-SET-EXC with above derivation.

– Case $e = (e' : e'')$

By induction hypothesis given by Lemma E.22 on $n - 1$ with e', e'' and σ we get a derivation $\underline{e', e''}; \sigma \xrightarrow[\text{expr}^*]{n-1} \text{vres}\star'; \sigma'$.

By case analysis on $\text{vres}\star'$:

* Case $\text{vres}\star' = \text{success } \underline{v, v'}$:

We proceed by checking whether all values \underline{v} and $\underline{v'}$ are non- \blacksquare :

- Case $\underline{v} \neq \blacksquare$ and $\underline{v'} \neq \blacksquare$ then use E-MAP-SUCS with above derivation.
- Case $v_i \neq \blacksquare$ or $v'_i \neq \blacksquare$ for some i then use E-MAP-ERR with above derivation.

* Case $\text{vres}\star' = \text{exres}$ then use E-MAP-EXC with above derivation.

– Case $e = e_1[e_2]$

By induction hypothesis on $n - 1$ with e_1 and σ we get a

derivation $e_1; \sigma \xrightarrow[\text{expr}]{\mathcal{E}, n-1} \text{vres}_1; \sigma''$.

By case analysis on vres_1 :

* Case $\text{vres}_1 = \text{success } v_1$:

By case analysis on v_1 :

· Case $v_1 = (\underline{v'} : v'')$:

By induction hypothesis on $n - 1$ with e_2 and σ'' we get a derivation $e_2; \sigma'' \xrightarrow[\text{expr}]{\mathcal{E}', n-1} \text{vres}_2; \sigma'$.

By case analysis on vres_2 :

• Case $\text{vres}_2 = \text{success } v_2$:

We proceed to check whether $\exists i. v'_i = v_2$:

- Case $v'_i = v_2$ then use E-LOOKUP-SUCS with \mathcal{E} and \mathcal{E}' .
- Case $\nexists i. v'_i = v_2$ then use E-LOOKUP-NOKEY with \mathcal{E} and \mathcal{E}' .

• Case $\text{vres}_2 = \text{exres}$ then use E-LOOKUP-EXC2 with \mathcal{E} and \mathcal{E}' .

· Case $v_1 \neq (\underline{v'} : v'')$ then use E-LOOKUP-ERR with \mathcal{E} .

* Case $\text{vres}_1 = \text{exres}$ then use E-LOOKUP-EXC1 with \mathcal{E} .

– Case $e = e_1[e_2 = e_3]$

By induction hypothesis on $n - 1$ with e_1 and σ we get a

derivation
$$e_1; \sigma \xrightarrow[\text{expr}]{\mathcal{E}}^{n-1} vres_1; \sigma''' .$$

By case analysis on $vres_1$:

* Case $vres_1 = \mathbf{success} v_1$

By case analysis on v_1 :

· Case $v_1 = (\underline{v'} : v'')$:

By induction hypothesis on $n - 1$ with e_2 and σ''' we

get a derivation
$$e_2; \sigma''' \xrightarrow[\text{expr}]{\mathcal{E}'}}^{n-1} vres_2; \sigma'' .$$

By case analysis on $vres_2$:

• Case $vres_2 = \mathbf{success} v_2$:

By induction hypothesis on $n - 1$ with e_3 and σ'' we

get a derivation
$$e_3; \sigma'' \xrightarrow[\text{expr}]{\mathcal{E}''}}^{n-1} vres_3; \sigma' .$$

By case analysis on $vres_3$:

– Case $vres_3 = \mathbf{success} v_3$:

We proceed to check whether v_2 and v_3 are non-
■:

* Case $v_2 \neq \blacksquare$ and $v_3 \neq \blacksquare$ then use
E-UPDATE-SUCS with \mathcal{E} , \mathcal{E}' and \mathcal{E}'' .

* Case $v_2 = \blacksquare$ or $v_3 = \blacksquare$ then use
E-UPDATE-ERR2 with \mathcal{E} , \mathcal{E}' and \mathcal{E}'' .

– Case $vres_3 = \mathbf{exres}$ then use E-UPDATE-EXC3 with
 \mathcal{E} and \mathcal{E}' .

• Case $vres_2 = \mathbf{exres}$ then use E-UPDATE-EXC2 on \mathcal{E}
and \mathcal{E}' .

· Case $v_1 \neq (\underline{v'} : v'')$ then use E-UPDATE-ERR1 with \mathcal{E} .

* Case $vres_1 = \mathbf{exres}$ then use E-UPDATE-EXC1 with \mathcal{E} .

– Case $e = f(\underline{e'})$

By induction hypothesis given by Lemma E.22 on $n - 1$ with

$\underline{e'}$ and σ we get a derivation
$$\underline{e'}; \sigma \xrightarrow[\text{expr}^*]{\mathcal{ES}}^{n-1} vres^*; \sigma''$$

By case analysis on $vres^*$:

* Case $\overline{vres\star''} = \mathbf{success} \overline{v''}$

Recall that we assume that our function calls are well-scoped and so there must exist a corresponding function definition $\mathbf{fun} \ t' \ f(\underline{t \ x}) = e''$. By Lemma E.8 we know that we have $\overline{v''} : \underline{t''}$ for some type sequence $\underline{t''}$.

We proceed to check whether $\underline{t''} <: \underline{t}$:

· Case $\underline{t''} <: \underline{t}$

Let $\mathbf{global} \ t_y \ y$ represent all global variable definitions.

By induction hypothesis on $n - 1$ with e'' and $[\underline{y \mapsto \sigma''(y)}, \underline{x \mapsto v''}]$ we get a derivation

$$e'; [\underline{y \mapsto \sigma''(y)}, \underline{x \mapsto v''}] \xrightarrow[\text{expr}]{\mathcal{E}}^{n-1} \overline{vres'}; \sigma''' .$$

By case analysis on $\overline{vres'}$:

• Case $\overline{vres'} = \mathbf{success} \ v'$ or $\overline{vres'} = \mathbf{return} \ v'$.

By Theorem 6.2 we know that $v' : t'''$ for some type t''' . We proceed to check whether $t''' <: t'$:

– Case $t''' <: t'$ then use E-CALL-RES-SUCS with \mathcal{ES} and \mathcal{E} .

– Case $t''' \not<: t'$ then use E-CALL-RES-ERR1 with \mathcal{ES} and \mathcal{E} .

• Case $\overline{vres'} = \mathbf{throw} \ v'$ then use E-CALL-RES-EXC with \mathcal{ES} and \mathcal{E}

• Case $\overline{vres'} \in \{\mathbf{break}, \mathbf{continue}, \mathbf{fail}, \mathbf{error}\}$ then use E-CALL-RES-ERR2 with \mathcal{ES} and \mathcal{E} .

· Case $t_i'' \not<: t_i$ for some i then use E-CALL-ARG-ERR with \mathcal{ES} .

* Case $\overline{vres\star''} = \mathbf{exres}$ then use E-CALL-ARG-EXC with \mathcal{ES} .

– Case $e = \mathbf{return} \ e'$

By induction hypothesis on $n - 1$ with e' and σ we get a derivation $e'; \sigma \xrightarrow[\text{expr}]{}^{n-1} \overline{vres'}; \sigma'$.

* Case $\overline{vres'} = \mathbf{success} \ v$ then use E-RET-SUCS with above derivation.

* Case $\overline{vres'} = \mathbf{exres}$ then use E-RET-EXC with above derivation.

- Case $e = (x = e')$

Recall that our definitions are assumed to be well-scoped and so there must exist either a **local** $t\ x$ or **global** $t\ x$ declaration for the variable (with no overshadowing).

By induction hypothesis on $n - 1$ with e' and σ we get a derivation $e'; \sigma \xrightarrow[\text{expr}]{}^{n-1} vres'; \sigma''$.

By case analysis on $vres'$:

- * Case $vres' = \mathbf{success}\ v$: By Theorem 6.2 we know that there exists a t' such that $v : t'$.

We proceed by checking whether $t' <: t$:

- Case $t' <: t$ then use E-ASGN-SUCS with above evaluation and typing derivations.
- Case $t' \not<: t$ then use E-ASGN-ERR with above evaluation and typing derivations.

- * Case $vres' = \mathit{exres}$ then use E-ASGN-EXC with above derivation.

- Case $e = \mathbf{if}\ e_1\ \mathbf{then}\ e_2\ \mathbf{else}\ e_3$

By induction hypothesis on $n - 1$ with e_1 and σ we get a derivation $e_1; \sigma \xrightarrow[\text{expr}]{}^{n-1} vres''; \sigma''$ with \mathcal{E} .

By case analysis on $vres''$:

- * Case $vres'' = \mathbf{success}\ v''$

By case analysis on v'' :

- Case $v'' = \mathbf{false}()$ then use E-IF-FALSE with \mathcal{E} and the derivation from the induction hypothesis on $n - 1$ with e_2 and σ'' .
- Case $v'' = \mathbf{true}()$ then use E-IF-TRUE with \mathcal{E} and the derivation from the induction hypothesis on $n - 1$ with e_3 and σ'' .
- Case $v'' \neq \mathbf{true}()$ and $v'' \neq \mathbf{false}()$ then use E-IF-ERR with \mathcal{E} .

- * Case $vres'' = \mathit{exres}$ then use E-IF-EXC with \mathcal{E} .

- Case $e = \mathbf{switch}\ e'\ \mathbf{do}\ \underline{cs}$

By induction hypothesis on $n - 1$ with e' and σ we get a derivation $e'; \sigma \xrightarrow[\text{expr}]{}^{n-1} vres''; \sigma''$ with \mathcal{E} .

By case analysis on $vres''$:

* Case $vres'' = \mathbf{success} v''$

By induction hypothesis given by Lemma E.26 on $n - 1$ with \underline{cs} , v'' and σ'' we get a derivation

$$\frac{\underline{CS}}{\underline{cs}; v''; \sigma'' \xrightarrow[\text{cases}]{}^{n-1} vres'; \sigma'}$$

By case analysis on $vres'$:

- Case $vres' = \mathbf{success} v'$ then use E-SWITCH-SUCS with \mathcal{E} and \underline{CS} .
- Case $vres' = \mathbf{fail}$ then use E-SWITCH-FAIL with \mathcal{E} and \underline{CS} .
- Case $vres' = exres$ where $exres \neq \mathbf{fail}$ then use E-SWITCH-Exc2 with \mathcal{E} and \underline{CS} .

* Case $vres'' = exres$ then use E-SWITCH-Exc1 with \mathcal{E} .

– Case $e = st \mathbf{visit} e' \mathbf{do} \underline{cs}$:

By induction hypothesis on $n - 1$ with e' and σ we get a

$$\text{derivation } \frac{\mathcal{E}}{e'; \sigma \xrightarrow[\text{expr}]{}^{n-1} vres''; \sigma''}$$

By case analysis on $vres''$:

* Case $vres'' = \mathbf{success} v''$

By induction hypothesis given by Lemma E.27 on $n - 1$ with \underline{cs} , v'' and σ'' we get a derivation

$$\frac{\mathcal{V}}{\underline{cs}; v''; \sigma'' \xrightarrow[\text{visit}]{}^{n-1} vres'; \sigma'}$$

By case analysis on $vres'$:

- Case $vres' = \mathbf{success} v'$ then use E-VISIT-SUCS with \mathcal{E} and \mathcal{V} .
- Case $vres' = \mathbf{fail}$ then use E-VISIT-FAIL with \mathcal{E} and \mathcal{V} .
- Case $vres' = exres$ where $exres \neq \mathbf{fail}$ then use E-VISIT-Exc2 with \mathcal{E} and \mathcal{V} .

* Case $vres'' = exres$ then use E-VISIT-Exc1 with \mathcal{E} .

– Case $e = \mathbf{break}$ then use E-BREAK.

– Case $e = \mathbf{continue}$ then use E-CONTINUE.

– Case $e = \mathbf{fail}$ then use E-FAIL.

– Case $e = \mathbf{local} \ \underline{t} \ x \ \mathbf{in} \ e' \ \mathbf{end}$ then use either E-BLOCK-SUCS (when it produces a successful result) or E-BLOCK-EXC (otherwise) with the derivation of the induction hypothesis given by Lemma E.22 on $n - 1$ with e' and σ .

– Case $e = \mathbf{for} \ g \ e'$:

By induction hypothesis given by Lemma E.24 on $n - 1$ with g and σ we get a derivation $g; \sigma \xrightarrow[\text{gexpr}]{}^{n-1} \text{envres}; \sigma''$.

By case analysis on envres :

* Case $\text{envres} = \mathbf{success} \ \underline{\rho}$ then use E-FOR-SUCS with above derivation and the derivation from the induction hypothesis given by Lemma E.23 on $n - 1$ with e' , $\underline{\rho}$ and σ'' .

* Case $\text{envres} = \text{exres}$ then use E-FOR-EXC with above derivation.

– Case $e = \mathbf{while} \ e_1 \ e_2$

By induction hypothesis on $n - 1$ with e_1 and σ we get a

derivation: $e_1; \sigma \xrightarrow[\text{expr}]{}^{n-1} \text{vres}''; \sigma''$.

By case analysis on vres'' :

* Case $\text{vres}'' = \mathbf{success} \ v''$

By case analysis on v'' :

· Case $v'' = \mathbf{false}()$ then use E-WHILE-FALSE with \mathcal{E} .

· Case $v'' = \mathbf{true}()$:

By induction hypothesis on $n - 1$ with e_2 and σ'' we

get a derivation $e_2; \sigma'' \xrightarrow[\text{expr}]{}^{n-1} \text{vres}'''; \sigma'''$.

By case analysis on vres''' :

• Case $\text{vres}''' = \mathbf{success} \ v'''$ or $\text{vres}''' = \mathbf{continue}$ then use E-WHILE-TRUE-SUCS with \mathcal{E} , \mathcal{E}' and the derivation from the induction hypothesis on $n - 1$ with $\mathbf{while} \ e_1 \ e_2$ and σ''' .

• Case $\text{vres}''' = \mathbf{break}$ then use E-WHILE-TRUE-BREAK with \mathcal{E} and \mathcal{E}' .

• Case $\text{vres}''' = \text{exres} \in \{\mathbf{throw} \ v''', \mathbf{return} \ v''', \mathbf{fail}, \mathbf{error}\}$ then use E-WHILE-EXC2 with \mathcal{E} and \mathcal{E}' .

- Case $v'' \neq \text{true}()$ and $v'' \neq \text{false}()$ then use E-WHILE-ERR
 - * Case $vres'' = exres$ then use E-WHILE-EXC1 with \mathcal{E} .
- Case $e = \text{solve } \underline{x} \ e'$:
- By induction hypothesis on $n - 1$ with e' and σ we get a derivation
$$e'; \sigma \xrightarrow[\text{expr}]{\mathcal{E}, n-1} vres''; \sigma'' .$$
- By case analysis on $vres''$
- * Case $vres'' = \text{success } v''$:
We proceed by checking whether the variable sequence \underline{x} is both in old and new stores.
 - Case $\underline{x} \subseteq \text{dom } \sigma \cap \text{dom } \sigma''$:
We then check whether any value in \underline{x} has changed:
 - Case $\sigma(\underline{x}) = \sigma''(\underline{x})$ then use E-SOLVE-EQ with \mathcal{E} .
 - Case $\sigma(x_i) \neq \sigma''(x_i)$ for some i then use E-SOLVE-NEQ with \mathcal{E} and the derivation from the induction hypothesis on $n - 1$ with $\text{solve } \underline{x} \ e'$ and σ'' .
 - Case $x_i \notin \text{dom } \sigma \cap \text{dom } \sigma''$ for some i then use E-SOLVE-ERR with \mathcal{E} .
 - * Case $vres'' = exres$ then use E-SOLVE-EXC with \mathcal{E} .
- Case $e = \text{throw } e'$:
- By induction hypothesis on $n - 1$ with e' and σ we get a derivation
$$e'; \sigma \xrightarrow[\text{expr}]{\mathcal{E}, n-1} vres'; \sigma' .$$
- * Case $vres' = \text{success } v$ then use E-THR-SUCS with above derivation.
 - * Case $vres' = exres$ then use E-THR-EXC with above derivation.
- Case $e = \text{try } e_1 \ \text{catch } x \Rightarrow e_2$:
- By induction hypothesis on $n - 1$ with e_1 and σ we get a derivation
$$e_1; \sigma \xrightarrow[\text{expr}]{\mathcal{E}, n-1} vres_1; \sigma'' .$$
- By case analysis on $vres_1$:

- * Case $vres_1 = \mathbf{throw} v_1$ then use E-TRY-CATCH with \mathcal{E} and the derivation from the induction hypothesis on $n - 1$ with e_2 and $\sigma''[x \mapsto v_1]$.
 - * Case $vres_1 \neq \mathbf{throw} v_1$ then use E-TRY-ORD with \mathcal{E} .
- Case $e = \mathbf{try} e_1 \mathbf{finally} e_2$
- By induction hypothesis on $n - 1$ with e_1 and σ we get a derivation
$$e_1; \sigma \xrightarrow[\text{expr}]{\mathcal{E}}^{n-1} vres_1; \sigma'' .$$
- By induction hypothesis on $n - 1$ with e_2 and σ'' we get a derivation
$$e_2; \sigma'' \xrightarrow[\text{expr}]{\mathcal{E}'}^{n-1} vres_2; \sigma' .$$
- By case analysis on $vres_2$:
- * Case $vres_2 = \mathbf{success} v_2$ then use E-FIN-SUCS with \mathcal{E} and \mathcal{E}'
 - * Case $vres_2 = \mathbf{exres}$ then use E-FIN-EXC with \mathcal{E} and \mathcal{E}' .

□

Theorem 6.4 (Terminating expressions). *There exists n such that derivation
$$e_{\text{fin}}; \sigma \xrightarrow[\text{expr}]{\mathcal{E}}^n vres; \sigma'$$
 has a result $vres$ which is not **timeout** for expression e_{fin} in the terminating subset.*

Proof. The proof proceeds similarly to Theorem 6.3, except that instead of doing the induction on n —which we know need to provide—we do the induction on the relevant syntactic element starting with the e_{fin} for this theorem. The only major complication is that for the **bottom-up** visit rules, we need to do an inner well-founded induction on the \prec relation on values when traversing the children in order to terminate.

The result n is simply taking to be $n' = 1 + n$ where n is the maximal fuel used in a sub-term. □

Appendix F

Rascal Light Verification Proofs

Theorem 7.1. Assuming Galois connection for element domains $\wp(E) \begin{smallmatrix} \xleftarrow{\gamma_{\hat{E}}} \\ \xrightarrow{\alpha_{\hat{E}}} \end{smallmatrix} \hat{E}$ then we have a Galois connection for the set shape domain $\wp(\wp(E)) \begin{smallmatrix} \xleftarrow{\gamma_{\widehat{SS}}} \\ \xrightarrow{\alpha_{\widehat{SS}}} \end{smallmatrix} \widehat{\text{SetShape}}(\hat{E})$

Proof. Since $\widehat{\text{SetShape}}(\hat{E})$ is a reduced product domain, then this property follows directly from the element domain Galois connection, and the interval domain Galois connection. \square

Theorem 7.2. Assuming a Galois connection for the leaf elements $\wp(E) \begin{smallmatrix} \xleftarrow{\gamma_{\hat{E}}} \\ \xrightarrow{\alpha_{\hat{E}}} \end{smallmatrix} \hat{E}$ then we have a Galois connection for the data elements $\wp(\text{Data}) \begin{smallmatrix} \xleftarrow{\gamma_{\widehat{DS}}} \\ \xrightarrow{\alpha_{\widehat{DS}}} \end{smallmatrix} \widehat{\text{DataShape}}(\hat{E})$

Proof. Follows directly by lifting from Theorem 7.3. \square

Theorem 7.3. Given a Galois connection $\wp(E) \begin{smallmatrix} \xleftarrow{\gamma_{\hat{E}}} \\ \xrightarrow{\alpha_{\hat{E}}} \end{smallmatrix} \hat{E}$, then there is a Galois connection $\text{PData}(E) \begin{smallmatrix} \xleftarrow{\gamma_{\widehat{R}}} \\ \xrightarrow{\alpha_{\widehat{R}}} \end{smallmatrix} \widehat{\text{Refinement}}(\hat{E})$

Proof Sketch. In essence this Galois connection signifies a conversion between syntactic definitions and semantics of data type refinements, except that the contained values have to respectively be in domains \hat{E} and E respectively, whose Galois connection follows from the premise. Note, that in case of fixed-point domains used in \hat{E} there is a possible hidden recursion between the operations on \hat{E} and \widehat{R} , but these are in such case

also interpreted as least fixed-point equations whose validity depends on Theorem 7.5. \square

Theorem 7.4. *Given $\forall i \left(\wp(A_i) \xleftrightarrow[\alpha_{\hat{A}_i}]{\gamma_{\hat{A}_i}} \hat{A}_i \right)$ then $\wp(\bigsqcup \underline{A}) \xleftrightarrow[\alpha_{\oplus}]{\gamma_{\oplus}} \oplus(\hat{\underline{A}})$*

Proof. We have to show $eas \subseteq \gamma_{\oplus}(\hat{e}\hat{a})$ if and only if $\alpha_{\oplus}(eas) \sqsubseteq \hat{e}\hat{a}$. We do this by showing it holds in both directions, by case analysis on $\hat{e}\hat{a}$ and inversion of the definition \sqsubseteq_{\oplus} , then delegating to the relevant Galois connection. \square

Theorem 7.5. *If that for continuous parametrized domains F, \hat{F} it holds that for all E, \hat{E} such that $\wp(E) \xleftrightarrow[\alpha_{\hat{E}}]{\gamma_{\hat{E}}} \hat{E}$ then $\wp(F(E)) \xleftrightarrow[\alpha_{\hat{F}}]{\gamma_{\hat{F}}} \hat{F}(\hat{E})$, then $\wp(\text{lfp } X.F(X)) \xleftrightarrow[\alpha_{\hat{F}\hat{X}}]{\gamma_{\hat{F}\hat{X}}} \text{lfp } \hat{X}.\hat{F}(\hat{X})$*

Proof. By definition of fixed-points[Scott, 1976; Smyth and Plotkin, 1982] we have that

$$\text{lfp } X.F(X) = \bigsqcup_{n \in \mathbb{N}} F^{(n)}(\emptyset)$$

where \emptyset is the empty domain. It suffices to show that for all n , there exists α' and γ' such that $\wp(F^{(n)}(\emptyset)) \xleftrightarrow[\alpha']{\gamma'} \hat{F}^{(n)}(\emptyset)$, which is done by induction over n :

- Case $n = 0$: We have to show $\wp(\emptyset) \xleftrightarrow[\alpha']{\gamma'} \emptyset$ which holds vacuously.
- Case $n = k + 1$: By induction hypothesis we have

$$\wp(F^{(k)}(\emptyset)) \xleftrightarrow[\alpha'']{\gamma''} \hat{F}^{(k)}(\emptyset)$$

and we have to show

$$(F^{(k+1)}(\emptyset)) \xleftrightarrow[\alpha']{\gamma'} \hat{F}^{(k+1)}(\emptyset)$$

which holds by application of the premise on the induction hypothesis. \square

Corollary 7.1. *If given a Galois connection for the parameter element domain $\wp(\mathbf{E}) \xleftrightarrow[\alpha_{\widehat{\mathbf{E}}}]^{\gamma_{\widehat{\mathbf{E}}}} \widehat{\mathbf{E}}$, then we have a Galois connection for the parameterized value domain $\wp(\text{PValue}(\mathbf{E})) \xleftrightarrow[\alpha_{\widehat{\text{PVS}}}]^{\gamma_{\widehat{\text{PVS}}}} \widehat{\text{PValueShape}}(\widehat{\mathbf{E}})$*

Proof. Follows directly from Proposition 7.1, Theorem 7.1, Theorem 7.2 and Theorem 7.4. \square

Corollary 7.2. *We have a Galois connection $\wp(\text{Value}) \xleftrightarrow[\alpha_{\widehat{\text{VS}}}]^{\gamma_{\widehat{\text{VS}}}} \widehat{\text{ValueShape}}$*

Proof. Follows directly from Corollary 7.1 and Theorem 7.5. \square

Theorem 7.6. *We have a Galois connection $\wp(\text{Store}) \xleftrightarrow[\alpha_{\widehat{\text{Store}}}]^{\gamma_{\widehat{\text{Store}}}} \widehat{\text{Store}}$*

Proof. Follows from lifting of the reduced product domain of the Boolean domain and $\widehat{\text{Value}}$ abstract domain, the latter which holds from Corollary 7.2. \square

Theorem 7.7. *We have a Galois connection*

$$\wp(\text{ValueResult} \times \text{Store}) \xleftrightarrow[\alpha_{\widehat{\text{RS}}}]^{\gamma_{\widehat{\text{RS}}}} \widehat{\text{ResultSet}}$$

Proof Sketch. We are essentially representing an infinite set of value result-store pairs in a finite set containing abstract values partitioned by their kind of control flow. The proof can therefore be done in two steps, one mapping directly the concrete value result-store pairs to abstract value result-store pairs in the domain $\wp(\widehat{\text{ValueResult}} \times \widehat{\text{Store}})$ —which is holds by lifting of Corollary 7.2 and Theorem 7.6)—and then show that this domain has a Galois connection to $\widehat{\text{ResultSet}}$ which is straightforward since it simply takes the least upper bound of all subsets with a particular kind of control flow. \square

Theorem 7.8 (Soundness). *For all valid expressions e , concrete stores σ and over-approximating abstract stores $\widehat{\sigma}$, so $\sigma \in \gamma_{\widehat{\text{Store}}}(\widehat{\sigma})$, where we have a concrete evaluation derivation $e; \widehat{\sigma} \xrightarrow[\text{expr}]{} vres; \sigma'$ and corresponding abstract evaluation derivation $e; \widehat{\sigma} \xrightarrow[\text{a-expr}]{} \widehat{\text{Res}}$ then it holds that the abstract result set properly over-approximates the concrete result value result and store, i.e., $(vres, \sigma') \in \gamma_{\text{RS}}(\widehat{\text{Res}})$.*

Proof Sketch. The soundness of our abstract interpreter concretely depends on three aspects:

1. That the abstract evaluation rules over-approximate all possible corresponding concrete evaluations. Because of the close correspondence between abstract and concrete evaluation rules in Schmidt-style abstract interpretation, this is done by showing that there is a homomorphism [Schmidt, 1998] from the concrete rules to the abstract rules, i.e. that for each concrete rule there exists a set of abstract rules that covers its evaluation.
2. That the abstract conditions and semantic operations properly over-approximate the concrete conditions and semantic operations. This is done using classical abstract domain operation analysis, as done for the theorems above.
3. That the chosen memoization strategies produce monotone results. The argument largely follows the one presented in Rosendahl [2013], but we must take into account the extension to infinite inputs. This extension is sound since we on recursion maintain monotonicity by using the least upper bound with the previous occurrence of inputs from the same partition and terminating because the partitioning is finite and the additional widening on input.

□

Appendix G

Rascal Subject programs

Negation Normal Form:

```
1 module NNF
2
3 data Formula = atom(str nm)
4             | and(Formula l, Formula r)
5             | or(Formula l, Formula r)
6             | imp(Formula l, Formula r)
7             | neg(Formula f);
8
9 Formula nnf(Formula phi) = top-down visit(phi) {
10 case neg(or(l,r)) => and(neg(l), neg(r))
11 case neg(and(l,r)) => or(neg(l), neg(r))
12 case neg(imp(l,r)) => and(l, neg(r))
13 case neg(neg(f)) => nnf(f)
14 case imp(l,r) => or(neg(l), r)
15 };
```

Rename Field of Struct:

```
1 module RenameStructField
2
3 data Nominal = nfn() | ofn() | other();
4 data Package = package(map[str, Struct] structures,
5                       map[str, Function] functions);
6 data Struct = struct(str name, map[Nominal, Field] fields
7                    ↪ );
8 data Field = field(Nominal name, str typ);
9 data Function = function(str name, str return_typ,
10                        list[Parameter] parameters, Stmt
11                       ↪ body);
12 data Parameter = parameter(str typ, str name);
13 data Stmt = ifstmt(Expr cond, Stmt thenb, Stmt elseb)
14           | returnstmt(Expr val)
15           | assignstmt(Expr lhs, Expr rhs)
16           | block(list[Stmt] stmts)
17           ;
18 data Expr = fieldaccessexpr(Expr target, Nominal
19                            ↪ fieldname)
```

```

17         | varexpr(str name)
18         | functioncallexpr(Expr target, str methodname,
19             ↪ list[Expr] args)
20     ;
21 Package renameField(Package pkg, str st,
22     Nominal oldFieldName, Nominal
23     ↪ newFieldName) {
24     assert (st in pkg.structures) &&
25     (oldFieldName in pkg.structures[st].fields) &&
26     (newFieldName notin pkg.structures[st].fields)
27     ↪ ;
28     Field fieldDef = pkg.structures[st].fields[
29     ↪ oldFieldName];
30     fieldDef.name = newFieldName;
31     pkg.structures[st].fields =
32     delete(pkg.structures[st].fields, oldFieldName)
33     + (newFieldName: fieldDef);
34     return top-down visit(pkg) {
35     case fieldaccessexpr(target, oldFieldName) =>
36         fieldaccessexpr(target, newFieldName)
37     };
38 }

```

Desugar Oberon

```

1 module DesugarOberon
2
3 data Module =
4     \mod(Ident name, Declarations decls, list[Statement]
5     ↪ body, Ident endName)
6     ;
7 data Declarations
8     = decls(list[ConstDecl] consts, list[TypeDecl] types,
9     ↪ list[VarDecl] vars)
10    ;
11 data ConstDecl
12     = constDecl(Ident name, Expression \value)
13     ;
14
15 data TypeDecl
16     = typeDecl(Ident name, Type \type)
17     ;
18
19 data VarDecl
20     = varDecl(list[Ident] names, Type \type)
21     ;
22
23 data Type
24     = user(Ident name)
25     ;
26
27 data MaybeExpression
28     = nothing()
29     | just(Expression e)

```

```

30   ;
31
32 data Statement
33   = assign(Ident var, Expression exp)
34   | ifThen(Expression condition, list[Statement] body,
35             list[ElseIf] elseIfs, list[Statement] elsePart
36             ↪ )
37   | whileDo(Expression condition, list[Statement] body)
38   | skip()
39   | forDo(Ident name, Expression from, Expression to,
40           MaybeExpression by, list[Statement] body)
41   | caseOf(Expression exp, list[Case] cases, list[
42             ↪ Statement] elsePart)
43   | begin(list[Statement] body)
44 ;
45
46 data Expression
47   = nat(int val)
48   | \trueE()
49   | \falseE()
50   | lookup(Ident var)
51   | neg(Expression exp)
52   | pos(Expression exp)
53   | not(Expression exp)
54   | mul(Expression lhs, Expression rhs)
55   | div(Expression lhs, Expression rhs)
56   | amp(Expression lhs, Expression rhs)
57   | add(Expression lhs, Expression rhs)
58   | sub(Expression lhs, Expression rhs)
59   | or(Expression lhs, Expression rhs)
60   | eq(Expression lhs, Expression rhs)
61   | neq(Expression lhs, Expression rhs)
62   | lt(Expression lhs, Expression rhs)
63   | gt(Expression lhs, Expression rhs)
64   | leq(Expression lhs, Expression rhs)
65   | geq(Expression lhs, Expression rhs)
66 ;
67
68 data Ident
69   = id(str name)
70 ;
71
72 data Case
73   = guard(Expression guard, list[Statement] body)
74 ;
75
76 data ElseIf = elseif(Expression condition, list[Statement
77 ↪ ] body);
78
79 public Module desugar(Module \mod) {
80   \mod.body = flattenBegin(case2ifs(for2while((\mod.
81 ↪ body))));
82   return \mod;
83 }

```

```

82 Statement cases2if(Expression e, list[Case] cs, list[
    ↪ Statement] es) {
83     switch (cs) {
84         case []: return begin(es);
85         case [c, *cs2]: {
86             list[Elseif] eis = [];
87             for (c2 <- cs2)
88                 eis = eis + [elseif(eq(e, c2.guard), c.
    ↪ body)];
89             return ifThen(eq(e, c.guard), c.body, eis, es
    ↪ );
90         }
91     };
92 }
93
94 public list[Statement] case2ifs(list[Statement] stats) {
95     return visit (stats) {
96         case caseOf(e, cs, es) => cases2if(e, cs, es)
97     }
98 }
99
100 public list[Statement] for2while(list[Statement] stats) {
101     return visit (stats) {
102         case forDo(n, f, t, just(by), b) =>
103             begin([assign(n, f), whileDo(geq(lookup(n), t
    ↪ ),
104                 b + [assign(n, add(lookup(n), by))])
    ↪ ])
105         case forDo(n, f, t, nothing(), b) =>
106             begin([assign(n, f), whileDo(geq(lookup(n), t
    ↪ ),
107                 b + [assign(n, add(lookup(n), nat(1))
    ↪ )])])
108     }
109 }
110
111 public list[Statement] flattenBegin(list[Statement] stats
    ↪ ) {
112     return innermost visit (stats) {
113         case list[Statement] ss =>
114             ({
115                 list[Statement] res = [];
116                 for (s <- ss) {
117                     switch (s) {
118                         case begin(b): res = res + b;
119                         case _: res = res + [s];
120                     };
121                 };
122                 res;
123             })
124     }
125 }

```

Glagol-to-PHP Expression Translation:

```

1 module Glagol2PHP
2

```

```

3 data Unsupported = unsupported();
4
5 data Expression
6   = integer(int intValue)
7   | string(str strValue)
8   | boolean(bool boolValue)
9   | \list(list[Expression] values)
10  | arrayAccess(Expression variable, Expression
11    ↪ arrayIndexKey)
12  | \map(map[Expression key, Expression \value])
13  | variable(str name)
14  | \bracket(Expression expr)
15  | product(Expression lhs, Expression rhs)
16  | remainder(Expression lhs, Expression rhs)
17  | division(Expression lhs, Expression rhs)
18  | addition(Expression lhs, Expression rhs)
19  | subtraction(Expression lhs, Expression rhs)
20  | greaterThanOrEq(Expression lhs, Expression rhs)
21  | lessThanOrEq(Expression lhs, Expression rhs)
22  | lessThan(Expression lhs, Expression rhs)
23  | greaterThan(Expression lhs, Expression rhs)
24  | equals(Expression lhs, Expression rhs)
25  | nonEquals(Expression lhs, Expression rhs)
26  | and(Expression lhs, Expression rhs)
27  | or(Expression lhs, Expression rhs)
28  | negative(Expression expr)
29  | positive(Expression expr)
30  | ternary(Expression condition, Expression ifThen,
31    ↪ Expression \else)
32  | new(str artifact, list[Expression] args)
33  | get(Type t)
34  | invoke(str methodName, list[Expression] args)
35  | invoke2(Expression prev, str methodName, list[
36    ↪ Expression] args)
37  | fieldAccess(str field)
38  | fieldAccess2(Expression prev, str field)
39  | emptyExpr()
40  | this()
41  ;
42
43 data Type
44   = integer2()
45   | string2()
46   | voidValue()
47   | boolean2()
48   | list2(Type \type)
49   | map2(Type key, Type v)
50   | artifact(str name)
51   | repository(str name)
52   | selfie()
53   ;
54
55 public data PhpOptionExpr = phpSomeExpr(PhpExpr expr) |
56   ↪ phpNoExpr();
57
58 public data PhpOptionName = phpSomeName(PhpName name) |
59   ↪ phpNoName();

```

```

55
56 public data PhpOptionElse = phpSomeElse(PhpElse e) |
    ↪ phpNoElse();
57
58 public data PhpActualParameter
59     = phpActualParameter(PhpExpr expr, bool byRef)
60     | phpActualParameter2(PhpExpr expr, bool byRef, bool
    ↪ isVariadic);
61
62 public data PhpConst = phpConst(str name, PhpExpr
    ↪ constValue);
63
64 public data PhpArrayElement = phpArrayElement(
    ↪ PhpOptionExpr key, PhpExpr val, bool byRef);
65
66 public data PhpName = phpName(str name);
67
68 public data PhpNameOrExpr = phpName2(PhpName name) |
    ↪ phpExpr(PhpExpr expr);
69
70 public data PhpCastType = phpIntCast() | phpBoolCast() |
    ↪ phpStringCast() | phpArrayCast() | phpObjectCast
    ↪ () | phpUnsetCast();
71
72 public data PhpClosureUse = phpClosureUse(PhpExpr varName
    ↪ , bool byRef);
73
74 public data PhpIncludeType = phpInclude() |
    ↪ phpIncludeOnce() | phpRequire() | phpRequireOnce
    ↪ ();
75
76 public data PhpExpr
77     = phpArray(list[PhpArrayElement] items)
78     | phpFetchArrayDim(PhpExpr var, PhpOptionExpr dim)
79     | phpFetchClassConst(PhpNameOrExpr className, str
    ↪ constantName)
80     | phpAssign(PhpExpr assignTo, PhpExpr assignExpr)
81     | phpAssignWOp(PhpExpr assignTo, PhpExpr assignExpr,
    ↪ PhpOp operation)
82     | phpListAssign(list[PhpOptionExpr] assignsTo,
    ↪ PhpExpr assignExpr)
83     | phpRefAssign(PhpExpr assignTo, PhpExpr assignExpr)
84     | phpBinaryOperation(PhpExpr left, PhpExpr right,
    ↪ PhpOp operation)
85     | phpUnaryOperation(PhpExpr operand, PhpOp operation)
86     | phpNew(PhpNameOrExpr className, list[
    ↪ PhpActualParameter] parameters)
87     | phpCast(PhpCastType castType, PhpExpr expr)
88     | phpClone(PhpExpr expr)
89     | phpClosure(list[PhpStmt] statements, list[PhpParam]
    ↪ params, list[PhpClosureUse] closureUses,
    ↪ bool byRef, bool static)
90     | phpFetchConst(PhpName name)
91     | phpEmpty(PhpExpr expr)
92     | phpSuppress(PhpExpr expr)
93     | phpEval(PhpExpr expr)

```



```

94 | phpExit(PhpOptionExpr exitExpr)
95 | phpCall(PhpNameOrExpr funName, list[
    ↪ PhpActualParameter] parameters)
96 | phpMethodCall(PhpExpr target, PhpNameOrExpr
    ↪ methodName, list[PhpActualParameter]
    ↪ parameters)
97 | phpStaticCall(PhpNameOrExpr staticTarget,
    ↪ PhpNameOrExpr methodName, list[
    ↪ PhpActualParameter] parameters)
98 | phpIncludeExpr(PhpExpr expr, PhpIncludeType
    ↪ includeType)
99 | phpInstanceOf(PhpExpr expr, PhpNameOrExpr toCompare
    ↪ )
100 | phpIsSet(list[PhpExpr] exprs)
101 | phpPrint(PhpExpr expr)
102 | phpPropertyFetch(PhpExpr target, PhpNameOrExpr
    ↪ propertyName)
103 | phpShellExec(list[PhpExpr] parts)
104 | phpTernary(PhpExpr cond, PhpOptionExpr ifBranch,
    ↪ PhpExpr elseBranch)
105 | phpStaticPropertyFetch(PhpNameOrExpr className,
    ↪ PhpNameOrExpr propertyName)
106 | phpScalar(PhpScalar scalarVal)
107 | phpVar(PhpNameOrExpr varName)
108 | phpYield(PhpOptionExpr keyExpr, PhpOptionExpr
    ↪ valueExpr)
109 | phpListExpr(list[PhpOptionExpr] listExprs)
110 | phpBracket(PhpOptionExpr bracketExpr)
111 ;
112
113 public data PhpOp = phpBitwiseAnd() | phpBitwiseOr() |
    ↪ phpBitwiseXor() | phpConcat() | phpDiv()
114 | phpMinus() | phpMod() | phpMul() |
    ↪ phpPlus() | phpRightShift() |
    ↪ phpLeftShift()
115 | phpBooleanAnd() | phpBooleanOr() |
    ↪ phpBooleanNot() | phpBitwiseNot()
116 | phpGt() | phpGeq() | phpLogicalAnd() |
    ↪ phpLogicalOr() | phpLogicalXor()
117 | phpNotEqual() | phpNotIdentical() |
    ↪ phpPostDec() | phpPreDec() |
    ↪ phpPostInc()
118 | phpPreInc() | phpLt() | phpLeq() |
    ↪ phpUnaryPlus() | phpUnaryMinus()
119 | phpEqual() | phpIdentical() ;
120
121 public data PhpParam
122 = phpParam(str paramName, PhpOptionExpr paramDefault,
    ↪ PhpOptionName paramType, bool byRef, bool
    ↪ isVariadic)
123 ;
124
125 public data PhpScalar
126 = phpClassConstant()
127 | phpDirConstant()
128 | phpFileConstant()

```

```

129 | phpFuncConstant()
130 | phpLineConstant()
131 | phpMethodConstant()
132 | phpNamespaceConstant()
133 | phpTraitConstant()
134 | phpNull()
135 | phpInteger(int intVal)
136 | phpString(str strVal)
137 | phpBoolean(bool boolVal)
138 | phpEncapsed(list[PhpExpr] parts)
139 | phpEncapsedStringPart(str strVal)
140 | ;
141
142 public data PhpStmt
143 = phpBreak(PhpOptionExpr breakExpr)
144 | phpClassDef(PhpClassDef classDef)
145 | phpConsts(list[PhpConst] consts)
146 | phpContinue(PhpOptionExpr continueExpr)
147 | phpDeclare(list[PhpDeclaration] decls, list[PhpStmt
    ↪ ] body)
148 | phpDo(PhpExpr cond, list[PhpStmt] body)
149 | phpEcho(list[PhpExpr] exprs)
150 | phpExprstmt(PhpExpr expr)
151 | phpFor(list[PhpExpr] inits, list[PhpExpr] conds,
    ↪ list[PhpExpr] exprs, list[PhpStmt] body)
152 | phpForeach(PhpExpr arrayExpr, PhpOptionExpr keyvar,
    ↪ bool byRef, PhpExpr asVar, list[PhpStmt]
    ↪ body)
153 | phpFunction(str name, bool byRef, list[PhpParam]
    ↪ params, list[PhpStmt] body, PhpOptionName
    ↪ returnType)
154 | phpGlobal(list[PhpExpr] exprs)
155 | phpGoto(str label)
156 | phpHaltCompiler(str remainingText)
157 | phpIf(PhpExpr cond, list[PhpStmt] body, list[
    ↪ PhpElseIf] elseIfs, PhpOptionElse elseClause)
158 | phpInlineHTML(str htmlText)
159 | phpInterfaceDef(PhpInterfaceDef interfaceDef)
160 | phpTraitDef(PhpTraitDef traitDef)
161 | phpLabel(str labelName)
162 | phpNamespace(PhpOptionName nsName, list[PhpStmt]
    ↪ body)
163 | phpNamespaceHeader(PhpName namespaceName)
164 | phpReturn(PhpOptionExpr returnExpr)
165 | phpStaticVars(list[PhpStaticVar] vars)
166 | phpSwitch(PhpExpr cond, list[PhpCase] cases)
167 | phpThrow(PhpExpr expr)
168 | phpTryCatch(list[PhpStmt] body, list[PhpCatch]
    ↪ catches)
169 | phpTryCatchFinally(list[PhpStmt] body, list[
    ↪ PhpCatch] catches, list[PhpStmt] finallyBody)
170 | phpUnset(list[PhpExpr] unsetVars)
171 | phpUseExpr(set[PhpUse] uses)
172 | phpWhile(PhpExpr cond, list[PhpStmt] body)
173 | phpEmptyStmt()
174 | phpBlock(list[PhpStmt] body)

```

```

175     | phpNewLine()
176     ;
177
178 public data PhpDeclaration = phpDeclaration(str key,
    ↪ PhpExpr val);
179
180 public data PhpCatch = phpCatch(PhpName xtype, str
    ↪ varName, list[PhpStmt] body);
181
182 public data PhpCase = phpCase(PhpOptionExpr cond, list[
    ↪ PhpStmt] body);
183
184 public data PhpElseIf = phpElseIf(PhpExpr cond, list[
    ↪ PhpStmt] body);
185
186 public data PhpElse = phpElse(list[PhpStmt] body);
187
188 public data PhpUse = phpUse(PhpName importName,
    ↪ PhpOptionName asName);
189
190 public data PhpClassItem
191     = phpPropertyCI(set[PhpModifier] modifiers, list[
    ↪ PhpProperty] prop)
192     | phpConstCI(list[PhpConst] consts)
193     | phpMethod(str name, set[PhpModifier] modifiers,
    ↪ bool byRef, list[PhpParam] params, list[
    ↪ PhpStmt] body, PhpOptionName returnType)
194     | phpTraitUse(list[PhpName] traits, list[
    ↪ PhpAdaptation] adaptations)
195     ;
196
197 public data PhpAdaptation
198     = phpTraitAlias(PhpOptionName traitName, str methName
    ↪ , set[PhpModifier] newModifiers,
    ↪ PhpOptionName newName)
199     | phpTraitPrecedence(PhpOptionName traitName, str
    ↪ methName, set[PhpName] insteadOf)
200     ;
201
202 public data PhpProperty = phpProperty(str propertyName,
    ↪ PhpOptionExpr defaultValue);
203
204 public data PhpModifier = phpPublic() | phpPrivate() |
    ↪ phpProtected() | phpStatic() | phpAbstract() |
    ↪ phpFinal();
205
206 public data PhpClassDef = phpClass(str className,
207     set[PhpModifier] modifiers,
208     PhpOptionName extends,
209     list[PhpName] implements,
210     list[PhpClassItem] members);
211
212 public data PhpInterfaceDef = phpInterface(str
    ↪ interfaceName,
213     list[PhpName] extends
    ↪ ,

```

```

214                                     list[PhpClassItem]
                                         ↪ members);
215
216 public data PhpTraitDef = phpTrait(str traitName, list[
                                         ↪ PhpClassItem] members);
217
218 public data PhpStaticVar = phpStaticVar(str name,
                                         ↪ PhpOptionExpr defaultValue);
219
220 public data PhpScript = phpScript(list[PhpStmt] body) |
                                         ↪ phpErrscript(str err);
221
222 public data PhpAnnotation
223     = phpAnnotation(str key)
224     | phpAnnotation(str key, PhpAnnotation v)
225     | phpAnnotationVal(map[str k, PhpAnnotation v])
226     | phpAnnotationVal(str string)
227     | phpAnnotationVal(int integer)
228     | phpAnnotationVal(bool boolean)
229     | phpAnnotationVal(list[PhpAnnotation] items)
230     | phpAnnotationVal(PhpAnnotation v)
231 ;
232
233 // Changed because of lack of string abstractions of
                                         ↪ built-in library functions
234 public str toLowerCaseFirstChar(str text) = text;
235
236 PhpExpr toPhpExpr(Expression expr) {
237     switch(expr) {
238
239         // literals
240         case integer(int i): return phpScalar(phpInteger(
                                         ↪ i));
241         case string(str s): return phpScalar(phpString(s)
                                         ↪ );
242         case boolean(bool b): return phpScalar(phpBoolean
                                         ↪ (b));
243
244         // arrays
245         case \list(list[Expression] items): {
246             list[PhpArrayElement] phpItems = [];
247             for (i <- items)
248                 phpItems = phpItems + [phpArrayElement(
                                         ↪ phpNoExpr(), toPhpExpr(i), false)];
249             return phpNew(phpName2(phpName("Vector")),
250                 [phpActualParameter(phpArray(phpItems
                                         ↪ ), false)]);
251         }
252
253         case arrayAccess(Expression variable, Expression
                                         ↪ arrayIndexKey):
254             return phpFetchArrayDim(toPhpExpr(variable),
                                         ↪ phpSomeExpr(toPhpExpr(arrayIndexKey))
                                         ↪ );
255         case \map(map[Expression key, Expression \value]
                                         ↪ m): {

```

```

256     list[PhpActualParameter] elements = [];
257     for (k <- m)
258         elements = elements +
259             [phpActualParameter(phpNew(
260                 ↪ phpName2(phpName("
261                     ↪ Pair")),
262                 [phpActualParameter(
263                     ↪ toPhpExpr(k),
264                     ↪ false),
265                 phpActualParameter(
266                     ↪ toPhpExpr(m[k
267                     ↪ ]), false)],
268                     ↪ false)];
269     return phpStaticCall(phpName2(phpName("
270         ↪ MapFactory")), phpName2(phpName("
271         ↪ createFromPairs")), elements);
272 }
273
274 case get(artifact(str name)):
275     return phpPropertyFetch(phpVar(phpName2(
276         ↪ phpName("this")), phpName2(phpName("
277         ↪ _" + toLowerCaseFirstChar(name))));
278
279 case variable(str name):
280     return phpVar(phpName2(phpName(name)));
281
282 case ternary(Expression condition, Expression
283     ↪ ifThen, Expression \else):
284     return phpTernary(toPhpExpr(condition),
285         ↪ phpSomeExpr(toPhpExpr(ifThen)),
286         ↪ toPhpExpr(\else));
287
288 case new(str artifact, list[Expression] args): {
289     list[PhpActualParameter] phpParams = [];
290     for (arg <- args)
291         phpParams = phpParams + [
292             ↪ phpActualParameter(toPhpExpr(arg)
293             ↪ , false)];
294     phpNew(phpName2(phpName(artifact)), phpParams
295         ↪ );
296 }
297
298 // Binary operations
299 case equals(Expression l, Expression r):
300     return phpBinaryOperation(toPhpExpr(l),
301         ↪ toPhpExpr(r), phpIdentical());
302 case greaterThan(Expression l, Expression r):
303     return phpBinaryOperation(toPhpExpr(l),
304         ↪ toPhpExpr(r), phpGt());
305 case product(Expression lhs, Expression rhs):
306     return phpBinaryOperation(toPhpExpr(lhs),
307         ↪ toPhpExpr(rhs), phpMul());
308 case remainder(Expression lhs, Expression rhs):
309     return phpBinaryOperation(toPhpExpr(lhs),
310         ↪ toPhpExpr(rhs), phpMod());
311 case division(Expression lhs, Expression rhs):

```

```

291     return phpBinaryOperation(toPhpExpr(lhs),
292                               ↪ toPhpExpr(rhs), phpDiv());
293 case addition(Expression lhs, Expression rhs):
294     return phpBinaryOperation(toPhpExpr(lhs),
295                               ↪ toPhpExpr(rhs), phpPlus());
296 case subtraction(Expression lhs, Expression rhs):
297     return phpBinaryOperation(toPhpExpr(lhs),
298                               ↪ toPhpExpr(rhs), phpMinus());
299 case \bracket(Expression e):
300     return phpBracket(phpSomeExpr(toPhpExpr(e)));
301 case greaterThanOrEq(Expression lhs, Expression
302     ↪ rhs):
303     return phpBinaryOperation(toPhpExpr(lhs),
304                               ↪ toPhpExpr(rhs), phpGeq());
305 case lessThanOrEq(Expression lhs, Expression rhs)
306     ↪ :
307     return phpBinaryOperation(toPhpExpr(lhs),
308                               ↪ toPhpExpr(rhs), phpLeq());
309 case lessThan(Expression lhs, Expression rhs):
310     return phpBinaryOperation(toPhpExpr(lhs),
311                               ↪ toPhpExpr(rhs), phpLt());
312 case greaterThan(Expression lhs, Expression rhs):
313     return phpBinaryOperation(toPhpExpr(lhs),
314                               ↪ toPhpExpr(rhs), phpGt());
315 case equals(Expression lhs, Expression rhs):
316     return phpBinaryOperation(toPhpExpr(lhs),
317                               ↪ toPhpExpr(rhs), phpIdentical());
318 case nonEquals(Expression lhs, Expression rhs):
319     return phpBinaryOperation(toPhpExpr(lhs),
320                               ↪ toPhpExpr(rhs), phpNotIdentical());
321 case and(Expression lhs, Expression rhs):
322     return phpBinaryOperation(toPhpExpr(lhs),
323                               ↪ toPhpExpr(rhs), phpLogicalAnd());
324 case or(Expression lhs, Expression rhs):
325     return phpBinaryOperation(toPhpExpr(lhs),
326                               ↪ toPhpExpr(rhs), phpLogicalOr());
327
328 // Unary operations
329 case negative(Expression e):
330     return phpUnaryOperation(toPhpExpr(e),
331                               ↪ phpUnaryMinus());
332 case positive(Expression e):
333     return phpUnaryOperation(toPhpExpr(e),
334                               ↪ phpUnaryPlus());
335
336 case invoke(str methodName, list[Expression] args
337     ↪ ): {
338     list[PhpActualParameter] phpParams = [];
339     for (arg <- args)
340         phpParams = phpParams + [
341             ↪ phpActualParameter(toPhpExpr(arg)
342             ↪ , false)];
343     return phpMethodCall(phpVar(phpName2(phpName(
344     ↪ "this"))), phpName2(phpName(
345     ↪ methodName)), phpParams);

```

```

327     }
328
329     case invoke2(Expression prev, str methodName,
330     ↪ list[Expression] args): {
331         list[PhpActualParameter] phpParams = [];
332         for (arg <- args)
333             phpParams = phpParams + [
334                 ↪ phpActualParameter(toPhpExpr(arg)
335                 ↪ , false)];
336
337         return phpMethodCall(toPhpExpr(prev),
338                 ↪ phpName2(phpName(methodName)),
339                 ↪ phpParams);
340     }
341
342     // Property fetch
343     case fieldAccess(str name):
344         return phpPropertyFetch(phpVar(phpName2(
345             ↪ phpName("this")), phpName2(phpName(
346             ↪ name)));
347
348     case fieldAccess2(Expression prev, str name):
349         return phpPropertyFetch(toPhpExpr(prev),
350             ↪ phpName2(phpName(name)));
351
352     case this():
353         return phpVar(phpName2(phpName("this")));
354
355     case _: throw unsupported();
356 };
357 }

```


Appendix H

Rascal Verification Output

Negation Normal Form:

```
1 refine Formula#19 = atom(string)
2 refine Formula#2045 = and(Formula#2045, Formula#2045) |
  ↪ atom(string) | neg(Formula#19) | or(Formula#2045,
  ↪ Formula#2045)
3
4 result:
5 success Formula#2045
6 store:
7 [phi -> Formula]
```

Rename Struct Field

```
1 refine Expr#564 = fieldaccessexpr(Expr#564, Nominal#442)
  ↪ | functioncallexpr(Expr#564, string, list[Expr
  ↪ #564]) | varexpr(string)
2 refine Function#6215 = function(string, string, list[
  ↪ Parameter], Stmt#3305)
3 refine Nominal#442 = nfn() | other()
4 refine Nominal#nfn = nfn()
5 refine Nominal#ofn = ofn()
6 refine Package#23593 = package(map[string, Struct], map[
  ↪ string, Function#6215])
7 refine Stmt#3305 = assignstmt(Expr#564, Expr#564) | block
  ↪ (list[Stmt#3305]) | ifstmt(Expr#564, Stmt#3305,
  ↪ Stmt#3305) | returnstmt(Expr#564)
8
9 result:
10 throw NoKey
11 store:
12 [st -> string,pkg -> Package,newFieldName -> Nominal#nfn
  ↪ ,oldFieldName -> Nominal#ofn]
13
14 result:
15 success Package#23593
16 store:
17 [st -> string,pkg -> Package,newFieldName -> Nominal#nfn
  ↪ ,oldFieldName -> Nominal#ofn]
```

Desugar Oberon-0

```

1 refine ElseIf#104075 = elseif(Expression, list[Statement
  ↪ #104076])
2 refine Module#207238 = mod(Ident, Declarations, list[
  ↪ Statement#104076], Ident)
3 refine Statement#104076 = assign(Ident, Expression) |
  ↪ begin(list[Statement#104076]) | ifThen(Expression
  ↪ , list[Statement#104076], list[ElseIf#104075],
  ↪ list[Statement#104076]) | skip() | whileDo(
  ↪ Expression, list[Statement#104076])
4
5 result:
6 error // Due to over-approximation
7 store:
8 [mod -> Module]
9
10 result:
11 success Module#207238
12 store:
13 [mod -> Module#207238]

```

Glagol-to-PHP Simple Expressions:

```

1 refine Expression#0 = addition(Expression#0, Expression
  ↪ #0) | and(Expression#0, Expression#0) | boolean(
  ↪ Bool) | bracket(Expression#0) | division(
  ↪ Expression#0, Expression#0) | emptyExpr() |
  ↪ equals(Expression#0, Expression#0) | greaterThan(
  ↪ Expression#0, Expression#0) | greaterThanOrEq(
  ↪ Expression#0, Expression#0) | integer(int) |
  ↪ lessThan(Expression#0, Expression#0) |
  ↪ lessThanOrEq(Expression#0, Expression#0) |
  ↪ negative(Expression#0) | nonEquals(Expression#0,
  ↪ Expression#0) | or(Expression#0, Expression#0) |
  ↪ positive(Expression#0) | product(Expression#0,
  ↪ Expression#0) | remainder(Expression#0,
  ↪ Expression#0) | string(string) | subtraction(
  ↪ Expression#0, Expression#0) | ternary(Expression
  ↪ #0, Expression#0, Expression#0) | variable(string
  ↪ )
2 refine PhpExpr#187 = phpBracket(PhpOptionExpr#185) |
  ↪ phpScalar(PhpScalar#90) | phpTernary(PhpExpr#187,
  ↪ PhpOptionExpr#185, PhpExpr#187) |
  ↪ phpUnaryOperation(PhpExpr#187, PhpOp#160) |
  ↪ phpVar(PhpNameOrExpr#18)
3 refine PhpExpr#9207 = phpBracket(PhpOptionExpr#9206) |
  ↪ phpScalar(PhpScalar#90) | phpUnaryOperation(
  ↪ PhpExpr#9207, PhpOp#160) | phpVar(PhpNameOrExpr
  ↪ #18)
4 refine PhpExpr#9794 = phpBinaryOperation(PhpExpr#9794,
  ↪ PhpExpr#9794, PhpOp#9795) | phpBracket(
  ↪ PhpOptionExpr#9796) | phpScalar(PhpScalar#90) |
  ↪ phpUnaryOperation(PhpExpr#9794, PhpOp#160) |
  ↪ phpVar(PhpNameOrExpr#18)
5 refine PhpExpr#9810 = phpBinaryOperation(PhpExpr#9794,
  ↪ PhpExpr#9794, PhpOp#9795) | phpBracket(

```

```

    ↪ PhpOptionExpr#9206) | phpScalar(PhpScalar#90) |
    ↪ phpTernary(PhpExpr#187, PhpOptionExpr#185,
    ↪ PhpExpr#187) | phpUnaryOperation(PhpExpr#9207,
    ↪ PhpOp#160) | phpVar(PhpNameOrExpr#18)
6 refine PhpNameOrExpr#18 = phpName2(PhpName)
7 refine PhpOp#160 = phpUnaryMinus() | phpUnaryPlus()
8 refine PhpOp#9795 = phpDiv() | phpGeq() | phpGt() |
    ↪ phpIdentical() | phpLeq() | phpLogicalAnd() |
    ↪ phpLogicalOr() | phpLt() | phpMinus() | phpMod()
    ↪ | phpMul() | phpNotIdentical() | phpPlus()
9 refine PhpOptionExpr#185 = phpSomeExpr(PhpExpr#187)
10 refine PhpOptionExpr#9206 = phpSomeExpr(PhpExpr#9207)
11 refine PhpOptionExpr#9796 = phpSomeExpr(PhpExpr#9794)
12 refine PhpScalar#90 = phpBoolean(Bool) | phpInteger(int)
    ↪ | phpString(string)
13 result:
14 throw Unsupported
15 store:
16 [expr -> Expression#0]
17
18 result:
19 error // Due to over-approximation
20 store:
21 [expr -> Expression#0]
22
23 result:
24 success PhpExpr#9810
25 store:
26 [expr -> Expression#0]

```

Glagol-to-PHP No Unary Expression

```

1 refine Bool#130 = false()
2 refine Expression#0 = addition(Expression#0, Expression
    ↪ #0) | and(Expression#0, Expression#0) |
    ↪ arrayAccess(Expression#0, Expression#0) | boolean
    ↪ (Bool) | bracket(Expression#0) | division(
    ↪ Expression#0, Expression#0) | emptyExpr() |
    ↪ equals(Expression#0, Expression#0) | fieldAccess(
    ↪ string) | fieldAccess2(Expression#0, string) |
    ↪ get(Type) | greaterThan(Expression#0, Expression
    ↪ #0) | greaterThanOrEq(Expression#0, Expression#0)
    ↪ | integer(int) | invoke(string, list[Expression
    ↪ #0]) | invoke2(Expression#0, string, list[
    ↪ Expression#0]) | lessThan(Expression#0,
    ↪ Expression#0) | lessThanOrEq(Expression#0,
    ↪ Expression#0) | list(list[Expression#0]) | map(
    ↪ map[Expression#0, Expression#0]) | new(string,
    ↪ list[Expression#0]) | nonEquals(Expression#0,
    ↪ Expression#0) | or(Expression#0, Expression#0) |
    ↪ product(Expression#0, Expression#0) | remainder(
    ↪ Expression#0, Expression#0) | string(string) |
    ↪ subtraction(Expression#0, Expression#0) | ternary
    ↪ (Expression#0, Expression#0, Expression#0) | this
    ↪ () | variable(string)
3 refine PhpActualParameter#2847 = phpActualParameter(
    ↪ PhpExpr#2848, Bool#130)

```

```

4 refine PhpActualParameter#2850 = phpActualParameter(
  ↪ PhpExpr#2852, Bool#130)
5 refine PhpActualParameter#37391 = phpActualParameter(
  ↪ PhpExpr#37393, Bool#130)
6 refine PhpActualParameter#37395 = phpActualParameter(
  ↪ PhpExpr#37388, Bool#130)
7 refine PhpActualParameter#38501 = phpActualParameter(
  ↪ PhpExpr#38510, Bool#130)
8 refine PhpActualParameter#38505 = phpActualParameter(
  ↪ PhpExpr#38509, Bool#130)
9 refine PhpActualParameter#54145 = phpActualParameter(
  ↪ PhpExpr#54146, Bool#130)
10 refine PhpActualParameter#54148 = phpActualParameter(
  ↪ PhpExpr#54142, Bool#130)
11 refine PhpActualParameter#54658 = phpActualParameter(
  ↪ PhpExpr#54659, Bool#130)
12 refine PhpActualParameter#54661 = phpActualParameter(
  ↪ PhpExpr#54655, Bool#130)
13 refine PhpActualParameter#598 = phpActualParameter(
  ↪ PhpExpr#607, Bool#130)
14 refine PhpActualParameter#602 = phpActualParameter(
  ↪ PhpExpr#603, Bool#130)
15 refine PhpActualParameter#62410 = phpActualParameter(
  ↪ PhpExpr#62411, Bool#130)
16 refine PhpActualParameter#62412 = phpActualParameter(
  ↪ PhpExpr#62417, Bool#130)
17 refine PhpActualParameter#62413 = phpActualParameter(
  ↪ PhpExpr#62414, Bool#130)
18 refine PhpArrayElement#2845 = phpArrayElement(
  ↪ PhpOptionExpr#483, PhpExpr#2848, Bool#130)
19 refine PhpArrayElement#37389 = phpArrayElement(
  ↪ PhpOptionExpr#483, PhpExpr#37393, Bool#130)
20 refine PhpArrayElement#38503 = phpArrayElement(
  ↪ PhpOptionExpr#483, PhpExpr#38509, Bool#130)
21 refine PhpArrayElement#54143 = phpArrayElement(
  ↪ PhpOptionExpr#483, PhpExpr#54146, Bool#130)
22 refine PhpArrayElement#54660 = phpArrayElement(
  ↪ PhpOptionExpr#483, PhpExpr#54655, Bool#130)
23 refine PhpArrayElement#600 = phpArrayElement(
  ↪ PhpOptionExpr#483, PhpExpr#603, Bool#130)
24 refine PhpArrayElement#62415 = phpArrayElement(
  ↪ PhpOptionExpr#483, PhpExpr#62417, Bool#130)
25 refine PhpExpr#2848 = phpBracket(PhpOptionExpr#2842) |
  ↪ phpMethodCall(PhpExpr#39, PhpNameOrExpr#38, list[
  ↪ PhpActualParameter#2847]) | phpNew(PhpNameOrExpr
  ↪ #38, list[PhpActualParameter#2850]) |
  ↪ phpPropertyFetch(PhpExpr#39, PhpNameOrExpr#38) |
  ↪ phpScalar(PhpScalar#140) | phpStaticCall(
  ↪ PhpNameOrExpr#38, PhpNameOrExpr#38, list[void]) |
  ↪ phpTernary(PhpExpr#2848, PhpOptionExpr#2842,
  ↪ PhpExpr#2848) | phpVar(PhpNameOrExpr#38)
26 refine PhpExpr#2852 = phpArray(list[PhpArrayElement
  ↪ #2845])
27 refine PhpExpr#37388 = phpArray(list[PhpArrayElement
  ↪ #37389])
28 refine PhpExpr#37393 = phpBracket(PhpOptionExpr#37386) |
  ↪ phpMethodCall(PhpExpr#39, PhpNameOrExpr#38, list[

```

```

    ↪ PhpActualParameter#37391]) | phpNew(PhpNameOrExpr
    ↪ #38, list[PhpActualParameter#37395]) |
    ↪ phpPropertyFetch(PhpExpr#37393, PhpNameOrExpr#38)
    ↪ | phpScalar(PhpScalar#140) | phpStaticCall(
    ↪ PhpNameOrExpr#38, PhpNameOrExpr#38, list[void]) |
    ↪ phpVar(PhpNameOrExpr#38)
29 refine PhpExpr#38509 = phpBracket(PhpOptionExpr#38500) |
    ↪ phpMethodCall(PhpExpr#38509, PhpNameOrExpr#38,
    ↪ list[PhpActualParameter#38505]) | phpNew(
    ↪ PhpNameOrExpr#38, list[PhpActualParameter#38501])
    ↪ | phpPropertyFetch(PhpExpr#39, PhpNameOrExpr#38)
    ↪ | phpScalar(PhpScalar#140) | phpStaticCall(
    ↪ PhpNameOrExpr#38, PhpNameOrExpr#38, list[void]) |
    ↪ phpVar(PhpNameOrExpr#38)
30 refine PhpExpr#38510 = phpArray(list[PhpArrayElement
    ↪ #38503])
31 refine PhpExpr#39 = phpVar(PhpNameOrExpr#38)
32 refine PhpExpr#54142 = phpArray(list[PhpArrayElement
    ↪ #54143])
33 refine PhpExpr#54146 = phpBracket(PhpOptionExpr#54140) |
    ↪ phpMethodCall(PhpExpr#39, PhpNameOrExpr#38, list[
    ↪ PhpActualParameter#54145]) | phpNew(PhpNameOrExpr
    ↪ #38, list[PhpActualParameter#54148]) |
    ↪ phpPropertyFetch(PhpExpr#39, PhpNameOrExpr#38) |
    ↪ phpScalar(PhpScalar#140) | phpStaticCall(
    ↪ PhpNameOrExpr#38, PhpNameOrExpr#38, list[void]) |
    ↪ phpVar(PhpNameOrExpr#38)
34 refine PhpExpr#54655 = phpBinaryOperation(PhpExpr#54655,
    ↪ PhpExpr#54655, PhpOp#54656) | phpBracket(
    ↪ PhpOptionExpr#54657) | phpMethodCall(PhpExpr#39,
    ↪ PhpNameOrExpr#38, list[PhpActualParameter#54661])
    ↪ | phpNew(PhpNameOrExpr#38, list[
    ↪ PhpActualParameter#54658]) | phpPropertyFetch(
    ↪ PhpExpr#39, PhpNameOrExpr#38) | phpScalar(
    ↪ PhpScalar#140) | phpStaticCall(PhpNameOrExpr#38,
    ↪ PhpNameOrExpr#38, list[void]) | phpVar(
    ↪ PhpNameOrExpr#38)
35 refine PhpExpr#54659 = phpArray(list[PhpArrayElement
    ↪ #54660])
36 refine PhpExpr#603 = phpBracket(PhpOptionExpr#597) |
    ↪ phpFetchArrayDim(PhpExpr#603, PhpOptionExpr#597)
    ↪ | phpMethodCall(PhpExpr#39, PhpNameOrExpr#38,
    ↪ list[PhpActualParameter#602]) | phpNew(
    ↪ PhpNameOrExpr#38, list[PhpActualParameter#598]) |
    ↪ phpPropertyFetch(PhpExpr#39, PhpNameOrExpr#38) |
    ↪ phpScalar(PhpScalar#140) | phpStaticCall(
    ↪ PhpNameOrExpr#38, PhpNameOrExpr#38, list[void]) |
    ↪ phpVar(PhpNameOrExpr#38)
37 refine PhpExpr#607 = phpArray(list[PhpArrayElement#600])
38 refine PhpExpr#62411 = phpNew(PhpNameOrExpr#38, list[
    ↪ PhpActualParameter#62412])
39 refine PhpExpr#62414 = phpArray(list[PhpArrayElement
    ↪ #62415])
40 refine PhpExpr#62417 = phpBracket(PhpOptionExpr#62409) |
    ↪ phpMethodCall(PhpExpr#39, PhpNameOrExpr#38, list[
    ↪ PhpActualParameter#62412]) | phpNew(PhpNameOrExpr

```

```

↪ #38, list[PhpActualParameter#62413]) |
↪ phpPropertyFetch(PhpExpr#39, PhpNameOrExpr#38) |
↪ phpScalar(PhpScalar#140) | phpStaticCall(
↪ PhpNameOrExpr#38, PhpNameOrExpr#38, list[
↪ PhpActualParameter#62410]) | phpVar(PhpNameOrExpr
↪ #38)
41 refine PhpExpr#90001 = phpBinaryOperation(PhpExpr#54655,
↪ PhpExpr#54655, PhpOp#54656) | phpBracket(
↪ PhpOptionExpr#54140) | phpFetchArrayDim(PhpExpr
↪ #603, PhpOptionExpr#597) | phpMethodCall(PhpExpr
↪ #38509, PhpNameOrExpr#38, list[PhpActualParameter
↪ #38505]) | phpNew(PhpNameOrExpr#38, list[
↪ PhpActualParameter#54148]) | phpPropertyFetch(
↪ PhpExpr#37393, PhpNameOrExpr#38) | phpScalar(
↪ PhpScalar#140) | phpStaticCall(PhpNameOrExpr#38,
↪ PhpNameOrExpr#38, list[PhpActualParameter#62410])
↪ | phpTernary(PhpExpr#2848, PhpOptionExpr#2842,
↪ PhpExpr#2848) | phpVar(PhpNameOrExpr#38)
42 refine PhpNameOrExpr#38 = phpName2(PhpName)
43 refine PhpOp#54656 = phpDiv() | phpGeq() | phpGt() |
↪ phpIdentical() | phpLeq() | phpLogicalAnd() |
↪ phpLogicalOr() | phpLt() | phpMinus() | phpMod()
↪ | phpMul() | phpNotIdentical() | phpPlus()
44 refine PhpOptionExpr#2842 = phpSomeExpr(PhpExpr#2848)
45 refine PhpOptionExpr#37386 = phpSomeExpr(PhpExpr#37393)
46 refine PhpOptionExpr#38500 = phpSomeExpr(PhpExpr#38509)
47 refine PhpOptionExpr#483 = phpNoExpr()
48 refine PhpOptionExpr#54140 = phpSomeExpr(PhpExpr#54146)
49 refine PhpOptionExpr#54657 = phpSomeExpr(PhpExpr#54655)
50 refine PhpOptionExpr#597 = phpSomeExpr(PhpExpr#603)
51 refine PhpOptionExpr#62409 = phpSomeExpr(PhpExpr#62417)
52 refine PhpScalar#140 = phpBoolean(Bool) | phpInteger(int)
↪ | phpString(string)

53
54 result:
55 throw Unsupported
56 store:
57 [expr -> Expression#0]
58
59 result:
60 error // Due to over approximation
61 store:
62 [expr -> Expression#0]
63
64 result:
65 success PhpExpr#90001?
66 store:
67 [expr -> Expression#0]

```