

Compositional Compiler Correctness Via Parametric Simulations

A dissertation submitted towards the degree
Doctor Engineering (Dr.-Ing)
of the Faculty of Mathematics and Computer Science
of Saarland University

by

Georg Neis, M.Sc.

Saarbrücken, 2018

Tag des Kolloquiums

7. Juni 2018

Dekan der Fakultät

Univ.-Prof. Dr. Sebastian Hack

Vorsitzender des Prüfungsausschusses

Univ.-Prof. Dr. Sebastian Hack

Erstgutachter / Doktorvater

Prof. Dr. Derek Dreyer

Zweitgutachter

Dr. Deepak Garg

Drittgutachter

Prof. Lars Birkedal, Ph.D.

Akademischer Beisitzer

Marco Patrignani, Ph.D.

Zusammenfassung

Die Verifikation von Compilern ist für die Konstruktion von vollständig verifizierter Software essenziell. Jedoch beschränkt sich die meiste bisherige Forschung (z.B. CompCert) auf das Szenario der Kompilierung von *ganzen* Programmen. Um separate Kompilierung und das Linken von Erzeugnissen *verschiedener* Compiler zu unterstützen, scheint es erforderlich einen kompositionalen Begriff von Compilerkorrektheit zu entwickeln, der modular (kompatibel mit Linken), transitiv (mehrphasige Kompilierung unterstützend), und flexibel ist (anwendbar auf Compiler mit verschiedenen Zwischensprachen oder unkonventionellen Transformationen).

In dieser Arbeit formalisieren wir einen solchen Korrektheitsbegriff basierend auf *parametrischen Simulationen*, und entwickeln auf diese Art einen neuartigen Ansatz zur kompositionalen Compilerverifikation.

1. Wir führen die grundlegende Idee von parametrischen (Bi-)simulationen ein und stellen eine konkrete Instanz vor. Diese bildet eine kompositionale Technik zum Beweisen von Programmäquivalenzen einer höheren ML-ähnlichen Sprache \mathcal{S} .
2. Dann betrachten wir eine Umgebung, in der \mathcal{S} -Programme in systemnahen Code einer Maschinensprache \mathcal{T} übersetzt werden. Wir verallgemeinern unsere Formulierung von PBs zu *parametrischen interlingualen Simulationen* (PILS), und demonstrieren dass PILS als Fundament von kompositionaler Compilerverifikation dienen können.

Unsere PB- und PILS-Entwicklungen wurden mittels des Beweissystems Coq durchgeführt und korrekt bewiesen.

Summary

Compiler verification is essential for the construction of fully verified software, but most prior work (such as CompCert) has focused on verifying whole-program compilers. To support separate compilation and to enable linking of results from *different* verified compilers, it is important to develop a compositional notion of compiler correctness that is modular (preserved under linking), transitive (supports multi-pass compilation), and flexible (applicable to compilers that use different intermediate languages or employ non-standard program transformations).

In this thesis, we formalize such a notion of correctness based on *parametric simulations*, thus developing a novel approach to compositional compiler verification.

The thesis is roughly divided into two parts.

1. We introduce the basic idea of parametric (bi-)simulations (PBs) and present a concrete instance of it. This instance constitutes a compositional technique for proving equivalences of programs written in the same higher-order ML-like language \mathcal{S} .
2. We then move from such a single-language setting to a setting where \mathcal{S} programs are being compiled into low-level code of some machine language \mathcal{T} . We generalize our PB formalization to *parametric inter-language simulations* (PILS) and demonstrate that PILS can serve as a foundation for compositional compiler verification.

Both our PB and PILS developments have been carried out and proven correct using the Coq proof assistant.

Note

This thesis is based primarily on the following publications.

- **The Marriage of Bisimulations and Kripke Logical Relations.**
Chung-Kil Hur, Derek Dreyer, Georg Neis, and Viktor Vafeiadis. In *2012 ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2012)*.
- **The Transitive Composability of Relation Transition Systems.**
Chung-Kil Hur, Georg Neis, Derek Dreyer, and Viktor Vafeiadis. *Technical Report MPI-SWS-2012-002, May 2012*.
- **Pilsner: A Compositionally Verified Compiler for a Higher-Order Imperative Language.**
Georg Neis, Chung-Kil Hur, Jan-Oliver Kaiser, Craig McLaughlin, Derek Dreyer, and Viktor Vafeiadis. In *2015 ACM SIGPLAN International Conference on Functional Programming (ICFP 2015)*.

Acknowledgements

It's been a long journey. I am immensely grateful to my advisor Derek Dreyer and my mentors Andreas Rossberg and Chung-Kil Hur. Thank you so much for your support, guidance, care, and wisdom!

I also want to thank Lars Birkedal and Viktor Vafeiadis for their support and insights; Neel Krishnaswami for excellent explanations; Claudio Russo and Aleks Nanevski for making me feel welcome at conferences; and everyone at MPI-SWS for the great environment.

Last but not least, I thank my family.

My PhD studies were partially funded by a Google European Doctoral Fellowship in Programming Technology. Thank you, Google!

Contents

Zusammenfassung	iii
Summary	v
Note	vii
Acknowledgements	ix
1 Introduction	1
1.1 Compiler Correctness	1
1.2 Compositionality	3
1.3 State of the Art	4
1.4 This Thesis	6
1.4.1 First Part	6
1.4.2 Second Part	7
1.4.3 Other Chapters	7
1.4.4 Proof Mechanization	8
2 Background: Logical Relations and Protocols	9
2.1 Kripke Logical Relations	10
2.2 State Transition Systems	11
2.2.1 The Limitation of Worlds as Memory Relations	11
2.2.2 The Power of Worlds as State Transition Systems	12
2.2.3 Private Transitions	12
2.3 Dealing With Recursive Language Features	14
2.3.1 Step-Indexing	14
3 Parametric Bisimulations	17
3.1 Overview	18
3.2 Notation	19
3.3 The Language $F^{\mu!}$	20
3.3.1 Syntax and Static Semantics	20
3.3.2 Dynamic Semantics	21
3.3.3 Contextual Equivalence	23

3.3.4	Deterministic Allocation	23
3.4	Global vs. Local Knowledge	27
3.5	Warmup: Parametric Bisimulations for λ^μ	31
3.5.1	Definitions	31
3.5.2	Example Proof	36
3.5.3	Basic Properties and Soundness	37
3.5.4	Transitivity	41
3.6	Parametric Bisimulations for $F^{\mu!}$	42
3.6.1	Worlds	42
3.6.2	Treatment of Universal and Existential Types	44
3.6.3	Treatment of Reference Types	44
3.6.4	Lifting and Separating Conjunction of Local Worlds	46
3.6.5	Program Equivalence	46
3.6.6	Expression and Continuation Equivalence	50
3.6.7	Living in a Different World	51
3.7	Metatheory	52
3.7.1	Basics	52
3.7.2	Symmetry	56
3.7.3	Compatibilities	58
3.7.4	Congruency	62
3.7.5	Soundness	62
3.8	Examples	63
3.8.1	Well-Bracketed State Change	63
3.8.2	Twin Abstraction	66
3.8.3	World Generator	68
3.8.4	Twin Abstraction, Alternate Proof	69
3.8.5	A Free Theorem	70
3.9	Transitivity	71
3.9.1	Structure of the Transitivity Proof	72
3.9.2	First Part: Constructing the Full World W	72
3.9.3	Second Part: Constructing the Corresponding Local World w	79
3.10	Stuttering Parametric Bisimulations	87
3.10.1	The Problem with Eta	88
3.10.2	Guardedness Revisited	89
3.10.3	Logical Reduction and the Stutter Budget	89
3.10.4	Eta Revisited	90
3.10.5	First-Class Continuations	91
3.10.6	Comparison to Step-Indexing	91
3.11	Greatest Local Knowledge	92
3.12	Comparison To Logical Relations	93

4	Parametric Inter-Language Simulations	97
4.1	Overview	98
4.1.1	Transitivity	100
4.2	Languages	100
4.2.1	Language-Generic Approach	100
4.2.2	Language Specification	101
4.2.3	Source Language \mathcal{S}	104
4.2.4	Intermediate Language \mathcal{I}	108
4.2.5	Target Language \mathcal{T}	113
4.3	Worlds	121
4.3.1	Queries	123
4.3.2	Value Closure	125
4.3.3	Lifting and Separating Conjunction of Local Worlds	126
4.4	Concrete Global Worlds	128
4.4.1	Global References	128
4.4.2	Unary Parts	129
4.5	Simulations	135
4.5.1	Stuttering according to algebraic well-founded orders.	141
4.5.2	The two modes of \mathbf{E} and cfg	143
4.5.3	A Note on the Untyped Model	145
4.5.4	Convenience Lemmas	145
4.6	Example	150
4.6.1	Modules	150
4.6.2	Proof	151
4.7	Metatheory	156
4.7.1	Basics	156
4.7.2	Adequacy	157
4.7.3	Modularity	159
4.8	Proof of Transitivity	165
4.8.1	Overview	165
4.8.2	Constructing the Local World	166
4.8.3	Discussion	175
5	The Pilsner Compiler and Its Verification	177
5.1	Overview	178
5.2	From \mathcal{S} to \mathcal{I} : CPS Transformation	180
5.2.1	Definition	180
5.2.2	Verification	185
5.3	Infrastructure for Optimizations	187
5.3.1	Relating Open Expressions	188
5.3.2	Annotating Expressions With Transformations	193
5.4	The Commute Pass	197
5.4.1	Transformation	197

5.4.2	Verification	198
5.4.3	Alternative Implementation.	199
5.5	The Dedup Pass	199
5.5.1	Transformation	199
5.5.2	Verification	199
5.6	The Hoist Pass	200
5.6.1	Transformation	200
5.6.2	Verification	201
5.7	The Dead Code Elimination Pass	202
5.7.1	Transformation	202
5.7.2	Verification	202
5.8	The Inline Pass	206
5.8.1	Transformation	206
5.8.2	Verification	206
5.8.3	Freshening	210
5.9	The Contify Pass	210
5.9.1	Transformation	210
5.9.2	Verification	213
5.10	The Codegen Pass	213
5.10.1	Transformation	213
5.10.2	Verification	221
5.11	The Full Pilsner Compiler	229
5.12	The Zwickel Compiler	230
5.12.1	Transformation	231
5.12.2	Verification	232
5.13	The Self-Modifying Awkward Example	236
5.14	Mechanization and Extraction	238
5.15	Putting It All Together	239
6	Related Work	241
6.1	Logical Relations	241
6.2	Bisimulations	243
6.3	Compositional Compiler Correctness	244
6.4	Miscellaneous	246
	Bibliography	249

Chapter 1

Introduction

Contents

1.1	Compiler Correctness	1
1.2	Compositionality	3
1.3	State of the Art	4
1.4	This Thesis	6
1.4.1	First Part	6
1.4.2	Second Part	7
1.4.3	Other Chapters	7
1.4.4	Proof Mechanization	8

1.1 Compiler Correctness

Compilers constitute some of the most critical infrastructure of software development. And yet they often have bugs. In 2011, for instance, Yang *et al.* [92] uncovered 325 bugs in widely used C compilers. In more recent work, Le *et al.* [46] identified 147 confirmed bugs in the industrial-strength GCC and LLVM compilers.

Let us distinguish between two classes of compiler bugs:

1. Bugs that cause the compiler to crash, to run forever, to reject valid source programs, or to produce inefficient code. Those are bad.
2. Bugs that cause the compiler to not preserve the semantics of its input, *i.e.*, to produce code that *behaves differently* than its source program. Those are terrible.

Compiler correctness, at least in the context of this thesis, is the absence of the second class of bugs. These bugs are so bad because they can go unnoticed until they blow up a system in production, with potentially fatal consequences. Moreover, they essentially nullify any efforts of formally verifying the program (usually a very

expensive endeavor), because typical verification tools operate on the source code and simply trust in the compiler being correct. In a sense, these bugs can make a correct program incorrect. Out of the 147 bugs mentioned above, 95 were violations of semantics preservation.

Establishing compiler correctness means providing a formal, machine-checked guarantee that a compiler preserves the semantics of its source programs, and has been an active research area ever since [50, 55, 20]. The standard way of expressing this preservation formally is roughly as follows. Suppose we are working with a distinguished *source* language \mathcal{S} and a distinguished *target* language \mathcal{T} . A compiler c from \mathcal{S} to \mathcal{T} , modelled as a partial function from \mathcal{S} programs to \mathcal{T} programs, is correct (semantics-preserving) iff for any given \mathcal{S} program $p_{\mathcal{S}}$, for which the compiler produces a \mathcal{T} program $p_{\mathcal{T}}$, this target program $p_{\mathcal{T}}$ *refines* its source program $p_{\mathcal{S}}$:

$$\forall p_{\mathcal{S}}, p_{\mathcal{T}}. c(p_{\mathcal{S}}) = p_{\mathcal{T}} \implies \text{Behav}(p_{\mathcal{T}}) \subseteq \text{Behav}(p_{\mathcal{S}})$$

Refinement means that any observable behavior of $p_{\mathcal{T}}$ is also a valid observable behavior of $p_{\mathcal{S}}$, for some common notion of “observable behavior” that both the \mathcal{S} and \mathcal{T} languages share (*e.g.*, termination, I/O events). In the case of \mathcal{S} being a deterministic language, this actually implies that $p_{\mathcal{S}}$ and $p_{\mathcal{T}}$ have exactly the same observable behavior. In the case of \mathcal{S} being non-deterministic, the compiler is permitted to translate away some of the non-determinism. We will see a concrete definition of *Behav* later on.

Research on compiler correctness has made great progress in the last years. Undoubtedly the most successful project so far has been CompCert [47], an optimizing compiler for most of the C language, targetting PowerPC, ARM and x86 processors. It was developed by Leroy and collaborators using the Coq proof assistant [19]. Indeed, Le *et al.* [46] report that, despite extensive testing, they were unable to uncover a single bug in CompCert. On the more functional side of languages, CakeML [42, 85] has become a serious optimizing ML compiler, implemented and verified in the HOL4 system. (See Chapter 6 for further references.)

However, the notion of semantics preservation stated earlier suffers from a fundamental limitation, and therefore so does CompCert as well as CakeML: it only applies to *whole program* compilation. The issue is that the formulation of refinement in terms of the observable behavior of a program implicitly assumes that the program is complete and thus can be meaningfully executed.

Consider the simple scenario where the source program $p_{\mathcal{S}}$ consists of two modules m_1 and m_2 :

$$p_{\mathcal{S}} = \text{link}(m_1, m_2)$$

The correctness statement then tells us that $c(p_{\mathcal{S}})$ refines $p_{\mathcal{S}}$. But what if we want to compile m_1 and m_2 separately and link them together? In that case the target program $p_{\mathcal{T}}$ is given as

$$p_{\mathcal{T}} = \text{link}(c(m_1), c(m_2))$$

and what we actually need to know is whether $p_{\mathcal{T}}$ refines $p_{\mathcal{S}}$, not whether $c(p_{\mathcal{S}})$ refines $p_{\mathcal{S}}$. While it is possible for $p_{\mathcal{T}}$ and $c(p_{\mathcal{S}})$ to be syntactically identical, this situation is

very rare (the compiler must not do cross-module optimizations and linking must be very simple). In general, they are different programs and thus the above correctness statement about $c(p_S)$ is useless.

In practice, of course, most programs are indeed linked together from multiple separately-produced modules. Some of these modules may as well be generated by different compilers or even hand-optimized in assembly. The limitation of compiler correctness to whole-program compilation is therefore a serious handicap.

1.2 Compositionality

The obvious question then is: can we define a more general notion of semantics preservation that says what it means for a single *module* in a program to be compiled correctly? And can this be done while assuming as little as possible about how the other modules in the program are compiled?

In particular, we articulate the following three desiderata concerning different aspects of compositionality:

- **Modularity:** To enable verified separate compilation, semantics preservation should be defined at the level of modules, not just whole programs, and it should be preserved under linking. Specifically, suppose that source (\mathcal{S}) module m_S is refined by target (\mathcal{T}) module m_T , and that \mathcal{S} module m'_S is refined by \mathcal{T} module m'_T . We should then be able to conclude that the \mathcal{S} -level linking of m_S and m'_S is refined by the \mathcal{T} -level linking of m_T and m'_T . (This is sometimes referred to as “horizontal compositionality”.)

Of course, in the special case of modules representing whole programs, this new notion of semantics preservation should imply the original one.

- **Transitivity:** Compilers are rarely expressed as direct translations but typically consist of several (if not dozens of) individual transformation passes that form a pipeline. Proofs of semantic preservation should therefore be transitively composable. That is, one should be able to prove a compiler correct by verifying refinement for its constituent passes independently and then chaining the results together by transitivity. (This is sometimes referred to as “vertical compositionality”.)
- **Flexibility:** It should be possible to prove semantics preservation for a range of different compilers and program transformations, so that the results of different verified compilers (which might employ different intermediate languages) can be safely linked together, and so that hand-optimized and hand-verified machine code can be safely linked with compiler-generated code.¹

¹**Note:** In our model of the semantics preservation problem, every module in a program is represented by both an \mathcal{S} and a \mathcal{T} version, and we aim to prove that the \mathcal{T} version refines its \mathcal{S} counterpart. For modules that are compiled by a verified compiler, the \mathcal{T} version is generated automatically by the compiler. But for any module that is hand-coded in \mathcal{T} , one must also manually

1.3 State of the Art

We are not the first to broach this question—it has been an active research topic in recent years. But we argue that all previously proposed solutions are lacking in some dimension of compositionality.

Contextual Refinement The natural starting point is the standard notion of *contextual refinement*: target module $m_{\mathcal{T}}$ contextually refines source module $m_{\mathcal{S}}$ if $C[m_{\mathcal{T}}]$ refines $C[m_{\mathcal{S}}]$ for all closing program contexts C . While contextual refinement is inherently modular and transitive, it only applies if the input and output languages are the same (we call this the *intra-language* setting): $m_{\mathcal{T}}$ and $m_{\mathcal{S}}$ are modules written in the same language, since they are placed into (think: linked with) the same program context C . Contextual refinement is therefore unsuitable for realistic compiler verification (unless combined with multi-language semantics, as discussed below).

Bisimulations *Bisimulations* originate in the study of process calculi [52, 72, 73] and have successfully been employed (and generalized) to characterize contextual equivalence in a variety of languages [29, 28, 41, 80, 79, 74, 43, 45, 7, 78]. Unfortunately, bisimulation methods rely on technical devices that prevent them (it seems) from scaling to the *inter-language* setting of transformations from one language to another. Specifically, in order to deal with higher-order functions, bisimulation methods employ various “syntactic” devices that restrict the applicability of the methods to intra-language reasoning (see Section 3.4 for details).

Logical Relations *Logical Relations* were originally used to prove fundamental properties of typed lambda calculi such as strong normalization [84, 76] and lambda-definability results [66, 67]. Later they were applied to relational reasoning [69, 82, 90, 65, 64, 2, 38] and used as an effective technique for proving correctness of a wide variety of program transformations in a wide variety of languages. For instance, Dreyer *et al.*’s recent Kripke logical relations [4, 22] support proofs with a very clean and intuitive high-level structure, based on the idea of establishing a *state transition system* that expresses directly how the “abstract state” of a module may evolve over time. We review this technique in the background Chapter 2, as it is important to the work in this thesis.

supply its \mathcal{S} counterpart, which serves as a “specification” that the hand-coded \mathcal{T} module is then proven to refine. (We will see an example of this in Section 5.13.) This means that we can only account for hand-coded \mathcal{T} modules that have *some* \mathcal{S} -level counterpart. This is somewhat of a restriction at present, since we focus here on the setting where \mathcal{S} is a high-level, ML-like language and \mathcal{T} a low-level, assembly-like language, and certainly not all assembly modules have an ML-level counterpart. However, we do not view this as a fundamental restriction: there is nothing in principle preventing us from generalizing our approach to a setting where \mathcal{S} itself supports interoperation between high- and low-level modules. We discuss this point further in Section 6.3.

Besides being highly flexible, logical relations are also inherently modular and, unlike contextual refinement and bisimulations, can be used to relate different source and target languages. For these reasons, Benton and Hur [6] proposed the idea of employing logical relations to define compositional semantics preservation.

Hur and Dreyer [32] developed this idea further by formalizing the compositional correctness of a simple, single-pass compiler from an ML-like source language to an idealized assembly language. They additionally demonstrated the flexibility of their inter-language logical relations by using them to verify a contrived but illustrative example, wherein a higher-order ML function was implemented in a rather baroque way by some tricky hand-written *self-modifying* assembly code. Thanks to the modularity of their logical relations method, this highly non-standard assembly code could nonetheless be safely linked with assembly modules produced by their verified compiler, with the resulting assembly program guaranteed to preserve the semantics of the corresponding linked source modules.

Unfortunately, it is not clear how to scale Hur *et al.*'s approach from single- to multi-pass compilers because, although logical relations are modular and flexible, they are not typically transitive. Chapter 2 provides a short review of logical relations, including the issue of transitivity.

Multi-Language Semantics Motivated by the goal of supporting compiler verification for programs that interoperate between different languages, Perconti and Ahmed [62] propose an approach based on *multi-language semantics* [49]. In particular, they define a “big-tent” language that comprises the source, target, and intermediate languages of a compiler, and provides “wrapping” operations for embedding terms of each language within the others. They then use logical relations to prove that every source module is contextually equivalent to a suitably wrapped version of the target module to which it is compiled. In this way, their method synthesizes the benefits of logical relations (modularity and different source and target languages) and contextual equivalence (transitivity).

One downside of their approach is that the intermediate languages (ILs) used in a compiler show up explicitly in the statement of compiler correctness. This leads to a loss of modularity: the semantics of source-level linking is not preserved when linking the results of compilers that have different ILs. Another downside concerns flexibility: the approach seems to be restricted to compilers that use *typed* intermediate and assembly languages. Lastly, Perconti and Ahmed have so far only applied their technique to a compiler for a purely functional source language.

Compositional Verification for CompCert Motivated by the goal of compositional compiler verification, Beringer *et al.* [9] propose an adaptation of the CompCert framework based on a novel operational semantics that differentiates between internal (intra-module) and external (inter-module) function calls. They introduce a notion of “logical simulation relation” that assumes little about the memory transformations performed by external function calls.

Beringer *et al.*'s approach is transitive, but lacking somewhat in modularity and flexibility. Concerning modularity, it does not yet (according to the authors) support fully separate compilation. Concerning flexibility, it depends on compiler passes only performing a restricted set of memory transformations—permitting additional transformations can potentially break the transitivity property. In addition, their method appears to be geared specifically toward compilers in the style of CompCert, which employ a uniform memory model across source, intermediate, and target languages. It is not clear how to generalize their technique to support richer (*e.g.*, ML-like) source languages, or compilers whose source and target languages have different memory models.

1.4 This Thesis

In this thesis, we develop a novel approach to compositional compiler verification in the form of *parametric inter-language simulations* (*PILS*). PILS formalize a compositional notion of semantics preservations that matches the criteria described in Section 1.2.

In order to ease the presentation of PILS, the thesis is divided into two main parts.

1.4.1 First Part

In the first part, consisting of Chapter 3, we introduce the basic idea of parametric inter-language simulations in the *intra-language* setting of an ML-like language \mathcal{S} . We call this development *parametric bisimulations* (PBs). Focusing on a single high-level language first lets us develop the core theory without getting lost in the details of lower-level languages and their implementation of higher-level concepts.

PBs marry together some of the most appealing aspects of KLRs and bisimulations. In particular, they synthesize bisimulations' support for reasoning about recursive features via coinduction with KLRs' support for reasoning about local state via state transition systems. Moreover, PBs are designed to avoid the limitations of KLRs and bisimulations that preclude their generalization to inter-language reasoning. To achieve this goal, we have to come up with a novel way of accounting for higher-order functions in the context of a coinductive bisimulation-like proof method, without relying on the “syntactic” devices that previous bisimulation methods use.

We explore PBs here in the setting of \mathcal{S} —a call-by-value λ -calculus with general recursive types, products, sums, universals, existentials, general references, and I/O—as this provides a relatively clear point of comparison with recent work on both KLRs [22] and bisimulations [79]. Technically, our result is a compositional technique for proving equivalences of \mathcal{S} programs, proven sound with respect to contextual equivalence. With one notable exception (see Section 3.10), we believe PBs are capable of reasoning effectively about all the challenging \mathcal{S} equivalences studied in the aforementioned papers, and we demonstrate PBs' effectiveness on several such

equivalences.

Crucially, the proof of transitivity does not simply exploit soundness and completeness with respect to contextual equivalence². Instead, it is developed from scratch and requires a number of technical innovations.

1.4.2 Second Part

The second main part of the thesis consists of Chapters 4 and 5. In it, we move from the previous single-language setting to a setting where \mathcal{S} programs are being compiled into low-level code of an idealized machine language \mathcal{T} . We generalize our PBs formalization to the full parametric inter-language simulations and demonstrate that PILS can serve as a foundation for compositional compiler verification.

Essentially, we are doing for PBs what Hur and Dreyer [32] did for logical relations: namely, taking a proof technique for single-language reasoning and showing how to scale it to inter-language reasoning between high- and low-level languages. The key difference is that, due to its reliance on a logical-relations model, Hur and Dreyer’s method only supports verification of single-pass compilation, whereas PILS support verification of multi-pass compilation.

Indeed, PILS satisfy all three compositionality criteria from Section 1.1.

- PILS are *modular*: they enable compiler verification in a way that supports separate compilation and is preserved under linking.
- PILS are *transitive*: we use them to verify **Pilsner**, a simple (but non-trivial) *multi-pass* optimizing compiler from \mathcal{S} to \mathcal{T} , going through a CPS-based intermediate language \mathcal{I} . After CPS conversion, Pilsner performs several optimizations at the \mathcal{I} level prior to code generation. These optimizations include function inlining, contification, dead code elimination, and hoisting (Figure 1.1). Although Pilsner is relatively simple—it is not nearly as realistic as the (whole-program) verified CakeML compiler, for instance [42]—it is **the first multi-pass compiler for a higher-order imperative language to be compositionally verified**.
- PILS are *flexible*: we use them to additionally verify (1) **Zwickel**, a direct (one-pass) non-optimizing compiler from \mathcal{S} to \mathcal{T} , and (2) Hur and Dreyer’s aforementioned self-modifying code example, programmed as a \mathcal{T} module and proven correct w.r.t. an \mathcal{S} -level specification. Thanks to PILS’ modularity, the output of Zwickel and the self-modifying \mathcal{T} module can then be safely linked together with the output of Pilsner.

1.4.3 Other Chapters

The two main parts are preceded by Chapter 2, which briefly reviews Kripke logical relations, and followed by Chapter 6, in which we discuss related work and conclude.

²In fact, our PB model is not complete, but this is not necessarily a “bad” thing.

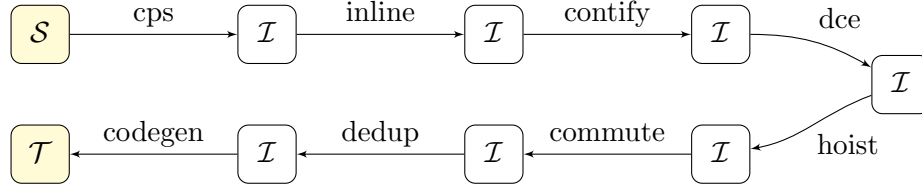


Figure 1.1: Structure of the Pilsner compiler

1.4.4 Proof Mechanization

Both PBs and PILS are formalized using the Coq proof assistant. The extensive development, available at <https://people.mpi-sws.org/~neis/thesis/>, includes machine-checked proofs of all theorems. We also provide glueing code to extract some definitions to OCaml, for instance in order to obtain an executable compiler.

Chapter 2

Background: Logical Relations and Protocols

Contents

2.1	Kripke Logical Relations	10
2.2	State Transition Systems	11
2.2.1	The Limitation of Worlds as Memory Relations	11
2.2.2	The Power of Worlds as State Transition Systems	12
2.2.3	Private Transitions	12
2.3	Dealing With Recursive Language Features	14
2.3.1	Step-Indexing	14

Logical relations are one of the best-known methods for local reasoning about observational equivalence¹ in higher-order, typed languages. The basic idea is to lift the notion of observable equivalence at base types to a compositional one at higher types by defining a notion of “logical equivalence” inductively on the type structure of the language. The Curry-Howard correspondence dictates which logical connective to use in assigning a relational operation to each type constructor. For instance, logical relatedness at *function type* is defined with the help of *implication* in terms of logical relatedness at argument type and logical relatedness at result type: two functions are logically related if relatedness of their arguments implies relatedness of their results; similarly, two *existential packages* are logically related if there *exists* a relational interpretation of their hidden type representations that is preserved by their operations; and so forth.

¹Or, more generally, for local reasoning about observational *refinement*. Here we talk in terms of equivalence, because this is what most of the work focusses on.

2.1 Kripke Logical Relations

Since Reynolds’ seminal paper on *relational parametricity* [69], which presented logical relations for reasoning about the pure, strongly normalizing System F, there has been a lot of work on generalizing and extending the method to handle increasingly realistic languages [64, 65, 51, 8, 3, 13]. Most relevant for this thesis is the line of work on modelling languages with state, where an equivalence may only hold due to particular invariants being imposed by the programs on parts of their memory. Naturally, these invariants must play a crucial role in any proof of that equivalence.

This is where *Kripke* logical relations come in. Kripke logical relations [65] are logical relations indexed by a *possible world* W , which codifies memory constraints. Roughly speaking, programs e_1 and e_2 are related under W only if they behave “the same” when run under any memories h_1 and h_2 , respectively, that satisfy the constraints of W .

When reasoning about programs that maintain some *local* state, *i.e.*, state that is encapsulated and cannot be freely accessed by the context, possible worlds allow us to impose whatever constraints on the local state we want, so long as we ensure that those constraints are preserved by the code that accesses the state. (Global state, *e.g.*, a reference that is passed to the context at a *ref* type, will have to obey the canonical invariants dictated by that type.)

In early Kripke logical relations, such as those of Pitts and Stark [65], worlds essentially take the form of simple memory relations, *i.e.*, memory invariants. To make things concrete, consider the following simple example (written in an ML-like language):

$$\begin{aligned}\tau &= (\text{unit} \rightarrow \text{unit}) \rightarrow \text{int} \\ e_1 &= \lambda f. (f \langle \rangle; 1) \\ e_2 &= \text{let } x = \text{ref } 1 \text{ in } \lambda f. (f \langle \rangle; !x)\end{aligned}$$

Each program evaluates to a higher-order function that, when applied, calls its argument f (a “callback” function) and then returns a number. In e_1 this number is simply 1. In the case of e_2 , however, this number is the value of reference x , which is initially set to 1. In this example, and in the variants to follow, keep in mind that the callback function f might itself call the higher-order function in question (here: the value of e_2).

We would like to show that e_1 and e_2 are observationally equivalent at type τ . The reason, intuitively, is obvious: x is kept private (*i.e.*, it is never leaked to the context) and thus references local state; since it is never modified by the function returned by e_2 , it will always point to 1.

To prove this using Kripke logical relations, we would set out to prove that e_1 and e_2 are related under an *arbitrary initial world* W . So suppose we evaluate the two terms under heaps h_1 and h_2 that satisfy W . Since the evaluation of e_1 results in the allocation of some *fresh* heap location l_x for x (*i.e.*, $l_x \notin \text{dom}(h_2)$), we know that the initial world W cannot already contain any constraints governing the contents of l_x . (If it contained such a constraint, h_1 would have had to satisfy it, and hence l_x

would have to be in $\text{dom}(h_1)$.) So we may extend W with a *new* invariant stating that $l_x \hookrightarrow 1$ (*i.e.*, l_x points to 1). More formally, this invariant can be represented as a binary heap relation of the following form:

$$\{(h_1, h_2) \mid h_2(l_x) = 1\}$$

(Note that while in this example no constraints are imposed on e_1 's heap, in general such a heap relation can talk about the contents of both heaps and even relate them to each other.)

It then remains to show that the two λ -abstractions are logically related under this extended world—*i.e.*, under the assumption that $l_x \hookrightarrow 1$. The key part of this concerns reasoning about the calls to f . We are guaranteed that if both calls return, the resulting memories again satisfy the extended world. This lets us conclude that both computations return 1 (assuming f terminates).

2.2 State Transition Systems

2.2.1 The Limitation of Worlds as Memory Relations

The applicability of worlds as growing collections of memory relations is fairly limited. To see why, let us look at Pitts and Stark's “awkward” example [65], a slight modification of the previous one.

$$\begin{aligned} \tau &= (\text{unit} \rightarrow \text{unit}) \rightarrow \text{int} \\ e_1 &= \lambda f. (f \langle \rangle; 1) \\ e_3 &= \text{let } x = \text{ref } 0 \text{ in } \lambda f. (x := 1; f \langle \rangle; !x) \end{aligned}$$

Only the second program has changed: e_3 is very similar to e_2 from the previous section, but (i) it initializes x with 0 instead of 1, and (ii) the returned function sets x to 1 before calling f .

Clearly, e_1 and e_3 are still equivalent. When the function returned by e_3 —say v_3 —is called for the very first time, it immediately writes 1 to x . Since there are no other writes, from that point on x will always hold 1 and so the value returned by the first and by any subsequent invocation of v_3 will always be 1 as well.

However, using the notion of worlds from the previous section, we cannot prove this equivalence. The reason is simple: the only “interesting” invariant concerning x is that it points to *either* 0 or 1, but this invariant is insufficient to rule out that after the call to f , the contents of x have changed back to 0. Hence a proof attempt would get stuck.

While this example is clearly contrived, it is also a minimal representative of a common class of programs in which changes to local state occur in some monotonic fashion. As Ahmed *et al.*[4] pointed out, this includes well-known *generative* (or *state-dependent*) abstract data types (ADTs), in which the interpretation of an abstract type grows over time in correspondence with changes to some local state.

2.2.2 The Power of Worlds as State Transition Systems

To address this limitation, Ahmed *et al.*[4] proposed generalizing possible worlds to include the ability for a memory relation to *evolve* over time. Dreyer *et al.*[22] later streamlined and extended Ahmed *et al.*'s approach in various ways, and cast Ahmed *et al.*'s possible worlds as collections of *state transition systems* (STS's), where each state determines a particular heap property, and where the transitions determine how the heap property may evolve. This approach lets one express complex protocols according to which programs manipulate their memory.

Let us illustrate how this works for the example above. Instead of extending the initial world W with a simple heap relation governing location l_x as we did in Section 2.1, we extend W with the following (very simple) STS with initial state s_0 :



This expresses that the value at l_x (in e_3 's heap) right after allocation is 0, but that it may eventually change to 1, from which point on it will stay 1 forever.

When showing that the returned λ -abstractions, say v_3 and v_4 , are logically related in this extended world, it is important that we not only consider their application when our STS resides in its initial state s_0 (in which these functions were returned), but also in any future state (here s_1). Intuitively, this is because we are effectively reasoning about an arbitrary application of v_3 and v_4 , not necessarily the first one (*i.e.*, l_x might already have been set to 1 by a previous call). If the calls happen in state s_0 , then we transition to state s_1 when v_4 assigns 1 to l_x . If they happen in s_1 , then the assignment doesn't change anything, so we just stay in that state. In either case, we conform to the protocol described by our STS and end up in s_1 after the assignment. Hence the call to f always happens when we are in the $l_x \mapsto 1$ state, and since there is no transition out of it, we know that f can only return in the same state. Consequently, v_4 's dereferencing l_x can only result in 1.

2.2.3 Private Transitions

A very useful extension to STSs, also due to Dreyer *et al.*[22], is the distinction between *public* and *private* transitions. Let us motivate this with the help of yet another modification of the “awkward” example. It is originally due to Jacob Thamsborg and has been dubbed the “very awkward” example or “well-bracketed state change” example in the literature.

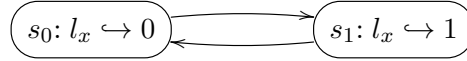
$$\begin{aligned}
 \tau &= (\text{unit} \rightarrow \text{unit}) \rightarrow \text{int} \\
 e_4 &= \lambda f. (f \langle \rangle; f \langle \rangle; 1) \\
 e_5 &= \text{let } x = \text{ref } 0 \text{ in } \lambda f. (x := 0; f \langle \rangle; x := 1; f \langle \rangle; !x)
 \end{aligned}$$

This time both programs have changed: they now make two calls to the callback f , instead of just one. Moreover, in e_5 , x is actually set to 0 before the first call, and only afterwards to 1, right before the second call.

Assuming that the language does not provide control effects such as `call/cc` or exceptions², e_4 and e_5 are in fact again equivalent. Intuitively, we can see that whenever x is set to 0, it will eventually be set to 1—no matter what the callback function does (as long as it terminates). Consequently, the final dereferencing will always yield 1.

But let's try to prove this equivalence using STSs. First, recall the STS that we used in order to prove the “awkward” example. It's easy to see why this STS is insufficient for our present purpose: suppose v_5 , the function value resulting from evaluating e_5 , is applied in the $l_x \hookrightarrow 1$ state. The first thing that happens is that l_x is set to 0. However, as there is no transition from state s_1 back to s_0 , there is no way we can continue the proof.

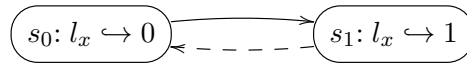
So how about adding that transition from $l_x \hookrightarrow 1$ to $l_x \hookrightarrow 0$, *i.e.*, using the following STS instead?



While doing so clears the first hurdle, it also erects a new one: according to the updated STS, it is now possible that, after the second call to f , we end up in state s_0 —even though this situation (l_x pointing to 0 after that call) cannot actually arise in reality. And indeed, if l_x could point to 0 at that point, our proof would be doomed (in fact, using this STS is effectively the same as using the insufficient simple invariant that l_x points to either 0 or 1). In summary, while we would like this transition to be *available to us*, we would at the same time like it to be *unavailable to the context*. This is the idea behind *private transitions*.

Private transitions are a new class of transitions in our STSs, separate from the ordinary transitions that we have seen so far (and which we henceforth call *public transitions*). The basic idea is very simple: when reasoning about the relatedness of terms, we must show that—when viewed *extensionally*—they appear to be making only public transitions, and correspondingly we may assume that the context only makes public transitions as well. Internally, however, *within* a computation, we may make use of both public and private transitions.

Concretely, we can use the following STS to prove our last example, where the dashed arrow denotes a private transition:



First, if v_5 is called in state s_1 , the presence of the private transition allows us to “lawfully” transition from s_1 to s_0 when l_x is set to 0. Second, we know that, because we are in the $l_x \hookrightarrow 1$ state before the second call to f and there is no *public* transition from there to any other state, we must still (or again) be in that same state when f returns. Hence we know that l_x holds 1 at that point, as desired. Lastly, although the body of v_5 makes a private transition internally when called in state s_1 , it appears

²The other examples hold even in the presence of such control features.

extensionally to make a public transition, since its final state (s_1) is obviously publicly accessible from itself.

This distinction of public and private transitions is also very useful in modelling and reasoning about concepts in low-level languages such as stack discipline, as we will see in Chapter 3.

2.3 Dealing With Recursive Language Features

The previous example programs store only integers in the heap, *i.e.*, they are all limited to first-order store. In a higher-order language, however, programs can also store references and functions in the heap. It turns out[2] that, in order for Kripke logical relations to support reasoning about higher-order store, one wants their heap relations to be world-dependent as well. This allows one, for instance, to express the constraint that whatever functions heaps h_1 and h_2 store (at some fixed pair of locations), these functions are logically related *in whatever the current world is*.

However, this immediately leads to a circularity in the construction of the whole model: the universe of worlds depends on the universe of heap relations and the universe of heap relations in turn depends on the universe of worlds, in an way that has no solution.

A different problem arises when trying to support general recursive types. One would like to define logical relatedness at such a type, say $\mu\alpha.\tau$, with the help of logical relatedness at the unfolded type $\tau[\mu\alpha.\tau/\alpha]$. Since the latter is generally larger, the logical relation can no longer be defined by recursion on types.

A third difficulty concerns reasoning about recursive functions. In the process of showing the relatedness of two recursive functions, one typically must establish the relatedness of the recursive calls, which in turn requires one to show the relatedness of the recursive functions—the very task one set out to do.

2.3.1 Step-Indexing

An elementary technique that solves all three problems at once is *step-indexing* [5, 2]. The idea is simply to stratify the construction of the logical relation by a natural number (or “step index”), representing roughly the number of steps of computation for which the programs in question behave in a related manner.

Regarding the first issue above, worlds of step-index n are then defined in terms of heap relations that can depend on worlds of step-index $n - 1$. Intuitively, this is sufficient because, in order for a program to make use of a value stored in the heap, it has to take at least one step (dereferencing a pointer). Regarding the second issue, one defines two values of recursive type to be related for n steps iff their unfoldings are related for $n - 1$ steps. Regarding the third issue, one can prove the relatedness of two recursive functions by induction on the step-index, effectively assuming that the functions are already related in any recursive calls.

Due to the tedious threading of step counting throughout proofs [21], it can be somewhat annoying to work with step-indexed logical relations (SILRs). There is a more serious catch, however: it seems fundamentally difficult to compose SILR proofs transitively.

Ahmed studied the transitivity problem in her first paper on binary SILRs [3]. There, she observes a serious problem in naively proving that Appel and McAllester’s original binary SILRs formed a PER. She proposes a way of regaining transitivity, but (i) she is concerned only with a very simple pure language, (ii) her proof relies on only working with syntactically well-typed terms, and (iii) it also depends on the second and third program (in the statement of transitivity) coming from the same language. Such an approach is unlikely to scale to reasoning about the intermediate and low-level languages of a compiler, which in general may be untyped. Moreover, we are not aware of any successful attempts to generalize her technique to SILRs for richer languages.

Birkedal & Bizjak [11] propose a different trick to obtain transitivity, but their technique is similarly unsuitable for an inter-language setting.

The lack of transitivity in many of the intra-language models is actually not a big issue, simply because they are only used to reason about contextual equivalence (or refinement), and so one can always transitively compose proofs at the level of contextual equivalence instead. However, when the purpose of the model is precisely to *define* a notion of equivalence—*e.g.*, because contextual equivalence does not make sense—then transitivity is critical.

Chapter 3

Parametric Bisimulations

Contents

3.1	Overview	18
3.2	Notation	19
3.3	The Language $F^{\mu!}$	20
3.3.1	Syntax and Static Semantics	20
3.3.2	Dynamic Semantics	21
3.3.3	Contextual Equivalence	23
3.3.4	Deterministic Allocation	23
3.4	Global vs. Local Knowledge	27
3.5	Warmup: Parametric Bisimulations for λ^{μ}	31
3.5.1	Definitions	31
3.5.2	Example Proof	36
3.5.3	Basic Properties and Soundness	37
3.5.4	Transitivity	41
3.6	Parametric Bisimulations for $F^{\mu!}$	42
3.6.1	Worlds	42
3.6.2	Treatment of Universal and Existential Types	44
3.6.3	Treatment of Reference Types	44
3.6.4	Lifting and Separating Conjunction of Local Worlds	46
3.6.5	Program Equivalence	46
3.6.6	Expression and Continuation Equivalence	50
3.6.7	Living in a Different World	51
3.7	Metatheory	52
3.7.1	Basics	52
3.7.2	Symmetry	56
3.7.3	Compatibilities	58
3.7.4	Congruency	62

3.7.5	Soundness	62
3.8	Examples	63
3.8.1	Well-Bracketed State Change	63
3.8.2	Twin Abstraction	66
3.8.3	World Generator	68
3.8.4	Twin Abstraction, Alternate Proof	69
3.8.5	A Free Theorem	70
3.9	Transitivity	71
3.9.1	Structure of the Transitivity Proof	72
3.9.2	First Part: Constructing the Full World W	72
3.9.3	Second Part: Constructing the Corresponding Local World w	79
3.10	Stuttering Parametric Bisimulations	87
3.10.1	The Problem with Eta	88
3.10.2	Guardedness Revisited	89
3.10.3	Logical Reduction and the Stutter Budget	89
3.10.4	Eta Revisited	90
3.10.5	First-Class Continuations	91
3.10.6	Comparison to Step-Indexing	91
3.11	Greatest Local Knowledge	92
3.12	Comparison To Logical Relations	93

3.1 Overview

In this chapter we present our technique for relational reasoning about programs, called **Parametric Bisimulations (PBs)**. PBs marry together some of the most appealing features of KLRs and bisimulations, while circumventing their limitations.

In particular, PBs show how the use of state transition systems (from KLRs) can be synthesized with the coinductive, step-index-free style of reasoning (from bisimulations), thereby enabling clean and elegant proofs about local state *and* recursive features simultaneously. Thus, concerning the long-standing open question of whether there is a fundamental tradeoff between KLRs and bisimulations, we provide a definitive answer: no, there is not.

We explore PBs here in the setting of $F^{\mu!}$ —a call-by-value λ -calculus with general recursive types, products, sums, universals, existentials, and general references [4]—as this provides a clear point of comparison with work on both KLRs [22] and bisimulations [79]. With one notable exception (see Section 3.10), we believe PBs are capable of reasoning effectively about all the challenging $F^{\mu!}$ equivalences studied in these papers, and we demonstrate PBs’ effectiveness on several such equivalences.

In this chapter we focus on intra-language reasoning. In order to enable the generalization to inter-language reasoning, the subject of Chapter 4, PBs were designed

so as to avoid the use of any limiting technical devices. To achieve this goal, we had to come up with a novel way of accounting for higher-order functions in the context of a coinductive bisimulation-like proof method, without relying on the “syntactic” devices that previous bisimulation methods use. Our solution—a new technique we call **global vs. local knowledge**—is one of the major contributions of this work. Relying heavily on this new technique, we have proven that PB equivalence proofs *are* transitively composable, which suggests they may serve as a superior foundation to KLRs for inter-language reasoning.

The remainder of this chapter is structured as follows. In Section 3.3, we define $F^{\mu!}$, the language under consideration. In Section 3.4, we motivate our key novel technical idea of global vs. local knowledge. We then present the formal development of PBs. For pedagogical reasons, we begin in Section 3.5 with the presentation of a relational model for λ^{μ} (a pure subset of $F^{\mu!}$ with recursive types), and then proceed in Section 3.6 to extend that model to handle the full language $F^{\mu!}$. Section 3.7 details the meta theory of the full model. The proof of transitivity merits its own section, Section 3.9. In Section 3.8, we demonstrate the expressive power of our method by proving several challenging equivalences from the literature. Afterwards, in Section 3.10, we explore a variant of PBs that we call *stuttering* PBs, and discuss its applications. In Section 3.11, we revisit local knowledges and point out how applying *parametric coinduction* [34] to their consistency property enables a convenient proof style. Finally, in Section 3.12, we recap and conclude with a comparison between PBs and logical relations.

Because of this chapter’s focus on establishing contextual equivalences, it is convenient to develop PBs as symmetric relations. This is not essential, however. When turning to refinement and compiler correctness in the next chapter, we will develop PILS using an asymmetric formulation.

3.2 Notation

Let us briefly establish some basic notational conventions that we use in the remainder of this thesis (not necessarily in the next section).

- We sometimes write meta-level functions in lambda notation (*e.g.*, $\lambda x.x + 42$).
- For a binary relation R and a natural number n , we write R^n for the $n - 1$ -times iterated relational composition $R \circ \dots \circ R$.
- For a unary relation S and a natural number n , we write S^n for the $n - 1$ -times iterated cartesian product $S \times \dots \times S$.
- We write R^* for the reflexive transitive closure of relation R and R^+ for its transitive closure.
- Given sets X and Y , we write $X \uplus Y$ to denote $X \cup Y$ when we want to emphasize that X and Y are disjoint.

- We write X_\perp short for X extended with a distinct element \perp .
- We write $X \multimap Y$ short for $X \rightarrow Y_\perp$, *i.e.*, the set of partial functions from X to Y .
- Given $f \in X \multimap Y$, we write $\text{dom}(f)$ to denote $\{x \in X \mid f(x) \neq \perp\}$, *i.e.*, the largest subset of X for which f is total.
- We write $X \stackrel{\text{fin}}{\multimap} Y$ short for $\{f \in X \multimap Y \mid \text{dom}(f) \text{ is finite}\}$, *i.e.*, for the set of partial functions from X to Y whose domain is finite.
- Given $f, g \in X \multimap Y$, we write $f \sqcup g$ to denote either the merge of f and g (if their domains are disjoint) or \perp (if they aren't):

$$f \sqcup g \in (X \multimap Y)_\perp$$

$$f \sqcup g = \begin{cases} \perp & \text{if } \text{dom}(f) \cap \text{dom}(g) \neq \emptyset \\ h & \text{otherwise} \end{cases}$$

$$\text{where } h(x) = \begin{cases} f(x) & \text{if } x \in \text{dom}(f) \\ g(x) & \text{if } x \in \text{dom}(g) \\ \perp & \text{otherwise} \end{cases}$$

- For better readability, we sometimes write properties (*e.g.*, lemmas) in the form of proof rules. Moreover, in these rules, we often write multiple conclusions in the same style as one writes multiple premises. For instance, the rule

$$\frac{P_1 \quad P_2 \quad P_3}{C_1 \quad C_2}$$

represents the property $P_1 \wedge P_2 \wedge P_3 \implies C_1 \wedge C_2$.

3.3 The Language $F^{\mu!}$

3.3.1 Syntax and Static Semantics

Figure 3.2 presents the surface syntax of $F^{\mu!}$, a call-by-value PCF-like language extended with features such as polymorphism and state. We assume a countably infinite set of type variables α and term variables x .

Types σ consist of variables (α), base types (including natural numbers), binary product and sum types, higher-order function types, universal and existential types, (iso-)recursive types, and general (higher-order) references.

Programs p consists of variables (x) and introduction and elimination constructs for each type form (some of these carrying explicit typing annotations):

- For the unit type: the unit value $(\langle \rangle)$.

$$\begin{aligned}
\sigma \in \text{Ty} &::= \alpha \mid \text{unit} \mid \text{nat} \mid \sigma_1 \times \sigma_2 \mid \sigma_1 + \sigma_2 \mid \sigma_1 \rightarrow \sigma_2 \mid \forall \alpha. \sigma \mid \exists \alpha. \sigma \mid \mu \alpha. \sigma \mid \text{ref } \sigma \\
p \in \text{Prg} &::= x \mid \langle \rangle \mid n \mid p_1 \odot p_2 \mid \text{ifnz } p_0 \text{ then } p_1 \text{ else } p_2 \mid \langle p_1, p_2 \rangle \mid p.1 \mid p.2 \mid \text{inl}_\sigma p \mid \\
&\quad \text{inr}_\sigma p \mid \text{case } p (x. p_1) (x. p_2) \mid \text{fix } f(x:\sigma_1):\sigma_2. p \mid p_1 p_2 \mid \\
&\quad \Lambda \alpha. p \mid p [\sigma] \mid \text{pack } \langle \sigma, p \rangle \text{ as } \exists \alpha. \sigma' \mid \text{unpack } p_1 \text{ as } \langle \alpha, x \rangle \text{ in } p_2 \mid \\
&\quad \text{roll}_\sigma p \mid \text{unroll } p \mid \text{ref } p \mid !p \mid p_1 := p_2 \mid p_1 == p_2
\end{aligned}$$

Figure 3.1: Surface syntax of $F^{\mu!}$.

- For the natural number type: constants (n), standard binary arithmetic operations ($p_1 \odot p_2$) such as addition and multiplication, and an “if not zero” conditional ($\text{ifnz } p_0 \text{ then } p_1 \text{ else } p_2$).
- For product types: product formation ($\langle p_1, p_2 \rangle$) and projections ($p.1, p.2$).
- For sum types: injections ($\text{inl}_\sigma p, \text{inr}_\sigma p$) and case analysis ($\text{case } p (x. p_1) (x. p_2)$).
- For function types: (recursive) abstraction ($\text{fix } f(x:\sigma_1):\sigma_2. p$) and application ($p_1 p_2$).
- For universal types: generalization ($\Lambda \alpha. p$) and instantiation ($p [\sigma]$).
- For existential types: packing ($\text{pack } \langle \sigma, p \rangle \text{ as } \exists \alpha. \sigma'$) and unpacking ($\text{unpack } p_1 \text{ as } \langle \alpha, x \rangle \text{ in } p_2$).
- For recursive types: folding ($\text{roll}_\sigma p$) and unfolding ($\text{unroll } p$).
- For reference types: allocation ($\text{ref } p$), reading ($!p$), writing ($p_1 := p_2$), and comparing ($p_1 == p_2$).

The static semantics of $F^{\mu!}$, given in Figure 3.2, is completely standard. It defines the program typing judgment $\Delta; \Gamma \vdash p : \sigma$, stating that program p has type σ in environments Δ and Γ . Δ lists the type variables that may appear free in p , while Γ lists the term variables and their types. We write $\sigma[\sigma'/\alpha]$ for the standard capture-avoiding substitution of σ' for variable α in σ .

3.3.2 Dynamic Semantics

The operational semantics of $F^{\mu!}$ is shown in Figure 3.3 and defined in terms of a standard substitution-based small-step reduction, \hookrightarrow , of run-time entities. These run-time entities, representing programs in execution, are pairs consisting of a *heap* h and an *expression* e . Initially, an expression is obtained simply by erasing the type annotations from a program p in the obvious way (written $|p|$). Expressions, however, can also contain heap *locations* l , as created by an allocation step (we

Type environments $\Delta ::= \cdot \mid \Delta, \alpha$

Term environments $\Gamma ::= \cdot \mid \Gamma, x:\sigma$

$\Delta; \Gamma \vdash p : \sigma$

$$\frac{\Delta \vdash \Gamma \quad x:\sigma \in \Gamma}{\Delta; \Gamma \vdash x : \sigma} \quad \frac{\Delta \vdash \Gamma}{\Delta; \Gamma \vdash n : \text{nat}} \quad \frac{\Delta; \Gamma \vdash p_1 : \text{nat} \quad \Delta; \Gamma \vdash p_2 : \text{nat}}{\Delta; \Gamma \vdash p_1 \odot p_2 : \text{nat}}$$

$$\frac{\Delta; \Gamma \vdash p_0 : \text{nat} \quad \Delta; \Gamma \vdash p_1 : \sigma \quad \Delta; \Gamma \vdash p_2 : \sigma}{\Delta; \Gamma \vdash \text{ifnz } p_0 \text{ then } p_1 \text{ else } p_2 : \sigma}$$

$$\frac{\Delta; \Gamma \vdash p_1 : \sigma_1 \quad \Delta; \Gamma \vdash p_2 : \sigma_2}{\Delta; \Gamma \vdash \langle p_1, p_2 \rangle : \sigma_1 \times \sigma_2} \quad \frac{\Delta; \Gamma \vdash p : \sigma_1 \times \sigma_2}{\Delta; \Gamma \vdash p.1 : \sigma_1} \quad \frac{\Delta; \Gamma \vdash p : \sigma_1 \times \sigma_2}{\Delta; \Gamma \vdash p.2 : \sigma_2}$$

$$\frac{\Delta; \Gamma \vdash p : \sigma_1 \quad \Delta; \Gamma \vdash \sigma_2}{\Delta; \Gamma \vdash \text{inl}_{\sigma_2} p : \sigma_1 + \sigma_2} \quad \frac{\Delta; \Gamma \vdash \sigma_1 \quad \Delta; \Gamma \vdash p : \sigma_2}{\Delta; \Gamma \vdash \text{inr}_{\sigma_1} p : \sigma_1 + \sigma_2}$$

$$\frac{\Delta; \Gamma \vdash p_0 : \sigma_1 + \sigma_2 \quad \Delta; \Gamma, x:\sigma_1 \vdash p_1 : \sigma \quad \Delta; \Gamma, x:\sigma_2 \vdash p_2 : \sigma}{\Delta; \Gamma \vdash \text{case } p_0 (x. p_1) (x. p_2) : \sigma}$$

$$\frac{\Delta; \Gamma, f:(\sigma_1 \rightarrow \sigma_2), x:\sigma_1 \vdash p : \sigma_2}{\Delta; \Gamma \vdash \text{fix } f(x:\sigma_1):\sigma_2. p : \sigma_1 \rightarrow \sigma_2} \quad \frac{\Delta; \Gamma \vdash p_1 : \sigma_1 \rightarrow \sigma_2 \quad \Delta; \Gamma \vdash p_2 : \sigma_1}{\Delta; \Gamma \vdash p_1 p_2 : \sigma_2}$$

$$\frac{\Delta, \alpha; \Gamma \vdash p : \sigma}{\Delta; \Gamma \vdash \Lambda \alpha. p : \forall \alpha. \sigma} \quad \frac{\Delta; \Gamma \vdash p : \forall \alpha. \sigma_1 \quad \Delta; \Gamma \vdash \sigma_2}{\Delta; \Gamma \vdash p [\sigma_2] : \sigma_1[\sigma_2/\alpha]}$$

$$\frac{\Delta; \Gamma \vdash \sigma_1 \quad \Delta; \Gamma \vdash p : \sigma_2[\sigma_1/\alpha]}{\Delta; \Gamma \vdash \text{pack } \langle \sigma_1, p \rangle \text{ as } \exists \alpha. \sigma_2 : \exists \alpha. \sigma_2}$$

$$\frac{\Delta; \Gamma \vdash p_1 : \exists \alpha. \sigma_1 \quad \Delta, \alpha; \Gamma, x:\sigma_1 \vdash p_2 : \sigma_2 \quad \Delta; \Gamma \vdash \sigma_2}{\Delta; \Gamma \vdash \text{unpack } p_1 \text{ as } \langle \alpha, x \rangle \text{ in } p_2 : \sigma_2}$$

$$\frac{\Delta; \Gamma \vdash p : \sigma[\mu \alpha. \sigma/\alpha]}{\Delta; \Gamma \vdash \text{roll}_{\mu \alpha. \sigma} p : \mu \alpha. \sigma} \quad \frac{\Delta; \Gamma \vdash p : \mu \alpha. \sigma}{\Delta; \Gamma \vdash \text{unroll } p : \sigma[\mu \alpha. \sigma/\alpha]}$$

$$\frac{\Delta; \Gamma \vdash p : \sigma}{\Delta; \Gamma \vdash \text{ref } p : \text{ref } \sigma} \quad \frac{\Delta; \Gamma \vdash p_1 : \text{ref } \sigma \quad \Delta; \Gamma \vdash p_2 : \sigma}{\Delta; \Gamma \vdash p_1 := p_2 : \text{unit}}$$

$$\frac{\Delta; \Gamma \vdash p : \text{ref } \sigma}{\Delta; \Gamma \vdash !p : \sigma} \quad \frac{\Delta; \Gamma \vdash p_1 : \text{ref } \sigma \quad \Delta; \Gamma \vdash p_2 : \text{ref } \sigma}{\Delta; \Gamma \vdash p_1 == p_2 : \text{nat}}$$

Figure 3.2: Static semantics of $F^{\mu l}$

assume a countably infinite set of locations Loc). A heap is a mapping from a finite subset of Loc to *values* v , which are a syntactic subset of expressions.

The reduction relation follows an eager left-to-right evaluation strategy and is defined the usual way using *evaluation contexts* K . We write $K[e]$ for filling the hole \bullet in K with expression e (defined in the obvious way). We also write $e[v/x]$ for the standard capture-avoiding substitution of value v for variable x in e , and $[l \mapsto v]$ for the singleton heap storing v at location l .

3.3.3 Contextual Equivalence

Our definition of *contextual equivalence* of $F^{\mu l}$ programs is also fairly standard. Roughly, two programs of the same type are contextually equivalent iff putting them in the same arbitrary (but well-typed) program context yields the same termination behavior.

Formally, we first define program contexts C (programs with a hole) and their typing, analogous to programs. This is shown in Figure 3.4. Context typing has the property that if $\vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma')$ and $\Delta; \Gamma \vdash p : \sigma$, then $\Delta'; \Gamma' \vdash C[p] : \sigma'$.

To express the termination behavior it suffices to define *divergence*. Intuitively, $\langle h; e \rangle \uparrow$ holds iff the repeated reduction of $\langle h; e \rangle$ never results in a value and never gets stuck, *i.e.*, goes on forever.

Definition 1 (Divergence). Coinductively, using a single recursive rule:

$$\frac{\langle h; e \rangle \hookrightarrow \langle h'; e' \rangle \quad \langle h'; e' \rangle \uparrow}{\langle h; e \rangle \uparrow}$$

Contextual equivalence is then defined as follows.

Definition 2 (Contextual equivalence).

$$\boxed{\Delta; \Gamma \vdash p_1 \sim_{\text{ctx}} p_2 : \sigma} \stackrel{\text{def}}{\iff} \Delta; \Gamma \vdash p_1 : \sigma \wedge \Delta; \Gamma \vdash p_2 : \sigma \wedge \forall C, \sigma', h. \\ \vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\epsilon; \epsilon; \sigma') \implies \\ (\langle h; |C[p_1]| \rangle \uparrow \iff \langle h; |C[p_2]| \rangle \uparrow)$$

3.3.4 Deterministic Allocation

Notice that reduction (3.3) is deterministic except for the rule for reference allocation (the $\text{ref } e$ expression form), which is completely non-deterministic. For technical reasons, however, it is more convenient to work with a deterministic allocation rule.

Given an allocation function $\text{alloc} \in \text{Heap} \rightarrow \text{Loc}$, let $\sim_{\text{ctx}}^{\text{alloc}}$ be the contextual equivalence of a deterministic version of $F^{\mu l}$, obtained by replacing the allocation rule with the following one (note the modified side condition):

$$\langle h; K[\text{ref } v] \rangle \hookrightarrow \langle h \sqcup [l \mapsto v]; K[l] \rangle \quad (l = \text{alloc}(h))$$

 $l \in \text{Loc}$

$$e \in \text{Exp} ::= x \mid \langle \rangle \mid n \mid e_1 \odot e_2 \mid \text{ifnz } e_0 \text{ then } e_1 \text{ else } e_2 \mid \langle e_1, e_2 \rangle \mid e.1 \mid e.2 \mid \text{inl } e \mid \text{inr } e \mid \text{case } e (x. e_1) (x. e_2) \mid \text{fix } f(x). e \mid e_1 e_2 \mid \Lambda. e \mid e [] \mid \text{pack } e \mid \text{unpack } e_1 \text{ as } x \text{ in } e_2 \mid \text{roll } e \mid \text{unroll } e \mid \text{ref } e \mid !e \mid e_1 := e_2 \mid e_1 == e_2 \mid l$$

$$v \in \text{Val} ::= \langle \rangle \mid n \mid \langle v_1, v_2 \rangle \mid \text{inl } v \mid \text{inr } v \mid \text{fix } f(x). e \mid \Lambda. e \mid \text{pack } v \mid \text{roll } v \mid l$$

$$K \in \text{Cont} ::= \bullet \mid K \odot e \mid v \odot K \mid \text{ifnz } K \text{ then } e_1 \text{ else } e_2 \mid \langle K, e \rangle \mid \langle v, K \rangle \mid K.1 \mid K.2 \mid \text{inl } K \mid \text{inr } K \mid \text{case } K (x. e_1) (x. e_2) \mid K e \mid v K \mid K [] \mid \text{pack } K \mid \text{unpack } K \text{ as } x \text{ in } e \mid \text{roll } K \mid \text{unroll } K \mid \text{ref } K \mid !K \mid K := e \mid v := K \mid K == e \mid v == K$$

$$h \in \text{Heap} ::= \text{Loc} \xrightarrow{\text{fin}} \text{Val}$$

Reduction: $\langle h; e \rangle \hookrightarrow \langle h'; e' \rangle$

$\langle h; K[n_1 \odot n_2] \rangle$	$\hookrightarrow \langle h; K[n] \rangle$	$(n = \llbracket \odot \rrbracket (n_1, n_2))$
$\langle h; K[\text{ifnz } 0 \text{ then } e_1 \text{ else } e_2] \rangle$	$\hookrightarrow \langle h; K[e_2] \rangle$	
$\langle h; K[\text{ifnz } n \text{ then } e_1 \text{ else } e_2] \rangle$	$\hookrightarrow \langle h; K[e_1] \rangle$	$(n \neq 0)$
$\langle h; K[\langle v_1, v_2 \rangle.1] \rangle$	$\hookrightarrow \langle h; K[v_1] \rangle$	
$\langle h; K[\langle v_1, v_2 \rangle.2] \rangle$	$\hookrightarrow \langle h; K[v_2] \rangle$	
$\langle h; K[\text{case } (\text{inl } v) (x. e_1) (x. e_2)] \rangle$	$\hookrightarrow \langle h; K[e_1[v/x]] \rangle$	
$\langle h; K[\text{case } (\text{inr } v) (x. e_1) (x. e_2)] \rangle$	$\hookrightarrow \langle h; K[e_2[v/x]] \rangle$	
$\langle h; K[(\text{fix } f(x). e) v] \rangle$	$\hookrightarrow \langle h; K[e[\text{fix } f(x). e/f][v/x]] \rangle$	
$\langle h; K[(\Lambda. e) []] \rangle$	$\hookrightarrow \langle h; K[e] \rangle$	
$\langle h; K[\text{unpack } (\text{pack } v) \text{ as } x \text{ in } e] \rangle$	$\hookrightarrow \langle h; K[e[v/x]] \rangle$	
$\langle h; K[\text{unroll } (\text{roll } v)] \rangle$	$\hookrightarrow \langle h; K[v] \rangle$	
$\langle h; K[\text{ref } v] \rangle$	$\hookrightarrow \langle h \sqcup [l \mapsto v]; K[l] \rangle$	$(l \notin \text{dom}(h))$
$\langle h \sqcup [l \mapsto v]; K[l := v'] \rangle$	$\hookrightarrow \langle h \sqcup [l \mapsto v']; K[\langle \rangle] \rangle$	$(l \notin \text{dom}(h))$
$\langle h \sqcup [l \mapsto v]; K[!l] \rangle$	$\hookrightarrow \langle h \sqcup [l \mapsto v]; K[v] \rangle$	$(l \notin \text{dom}(h))$
$\langle h; K[l_1 == l_2] \rangle$	$\hookrightarrow \langle h; K[1] \rangle$	$(l_1 = l_2)$
$\langle h; K[l_1 == l_2] \rangle$	$\hookrightarrow \langle h; K[0] \rangle$	$(l_1 \neq l_2)$

Figure 3.3: Dynamic semantics of $F^{\mu!}$.

$\vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma')$	
$\frac{\Delta; \Gamma \vdash \sigma}{\vdash \bullet : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta; \Gamma; \sigma)}$	
$\frac{\vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \text{nat}) \quad \Delta'; \Gamma' \vdash p_2 : \text{nat}}{\vdash C \odot p_2 : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \text{nat})}$	
$\frac{\Delta'; \Gamma' \vdash p_1 : \text{nat} \quad \vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \text{nat})}{\vdash p_1 \odot C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \text{nat})}$	
$\frac{\vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \text{nat}) \quad \Delta'; \Gamma' \vdash p_1 : \sigma' \quad \Delta'; \Gamma' \vdash p_2 : \sigma'}{\vdash \text{ifnz } C \text{ then } p_1 \text{ else } p_2 : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma')}$	
$\frac{\Delta'; \Gamma' \vdash p_0 : \text{nat} \quad \vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma') \quad \Delta'; \Gamma' \vdash p_2 : \sigma'}{\vdash \text{ifnz } p_0 \text{ then } C \text{ else } p_2 : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma')}$	
$\frac{\Delta'; \Gamma' \vdash p_0 : \text{nat} \quad \Delta'; \Gamma' \vdash p_1 : \sigma' \quad \vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma')}{\vdash \text{ifnz } p_0 \text{ then } p_1 \text{ else } C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma')}$	
$\frac{\vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma_1) \quad \Delta'; \Gamma' \vdash p_2 : \sigma_2}{\vdash \langle C, p_2 \rangle : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma_1 \times \sigma_2)}$	
$\frac{\vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma_1 \times \sigma_2)}{\vdash C.1 : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma_1)}$	$\frac{\vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma_1 \times \sigma_2)}{\vdash C.2 : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma_2)}$
$\frac{\Delta'; \Gamma' \vdash \sigma_2 \quad \vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma_1)}{\vdash \text{inl}_{\sigma_2} C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma_1 + \sigma_2)}$	$\frac{\Delta'; \Gamma' \vdash \sigma_1 \quad \vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma_2)}{\vdash \text{inr}_{\sigma_1} C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma_1 + \sigma_2)}$
$\frac{\vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma_1 + \sigma_2) \quad \Delta'; \Gamma', x:\sigma_1 \vdash p_1 : \sigma' \quad \Delta'; \Gamma', x:\sigma_2 \vdash p_2 : \sigma'}{\vdash \text{case } C (x. p_1) (x. p_2) : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma')}$	
$\frac{\Delta'; \Gamma' \vdash p_0 : \sigma_1 + \sigma_2 \quad \vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma', x:\sigma_1; \sigma') \quad \Delta'; \Gamma', x:\sigma_2 \vdash p_2 : \sigma'}{\vdash \text{case } p_0 (x. C) (x. p_2) : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma')}$	
$\frac{\Delta'; \Gamma' \vdash p_0 : \sigma_1 + \sigma_2 \quad \Delta'; \Gamma', x:\sigma_1 \vdash p_1 : \sigma' \quad \vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma', x:\sigma_2; \sigma')}{\vdash \text{case } p_0 (x. p_1) (x. C) : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma')}$	
$\frac{\vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma', f:(\sigma_1 \rightarrow \sigma_2), x:\sigma_1; \sigma_2)}{\vdash \text{fix } f(x:\sigma_1):\sigma_2. C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma_1 \rightarrow \sigma_2)}$	

Figure 3.4: Program contexts in $F^{\mu l}$.

$$\begin{array}{c}
\frac{\vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma_1 \rightarrow \sigma_2) \quad \Delta'; \Gamma' \vdash p_2 : \sigma_1}{\vdash C \ p_2 : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma_2)} \\
\\
\frac{\Delta'; \Gamma' \vdash p_1 : \sigma_1 \rightarrow \sigma_2 \quad \vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma_1)}{\vdash p_1 \ C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma_2)} \\
\\
\frac{\vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \alpha; \Gamma'; \sigma_1)}{\vdash \Lambda \alpha. C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \forall \alpha. \sigma_1)} \quad \frac{\vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \forall \alpha. \sigma_1) \quad \Delta'; \Gamma' \vdash \sigma_2}{\vdash C \ [\sigma_2] : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma_1[\sigma_2/\alpha])} \\
\\
\frac{\Delta'; \Gamma' \vdash \sigma_1 \quad \vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma_2[\sigma_1/\alpha])}{\vdash \text{pack } \langle \sigma_1, C \rangle \text{ as } \exists \alpha. \sigma_2 : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \exists \alpha. \sigma_2)} \\
\\
\frac{\vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \exists \alpha. \sigma_1) \quad \Delta', \alpha; \Gamma', x:\sigma_1 \vdash p_2 : \sigma_2 \quad \Delta'; \Gamma' \vdash \sigma_2}{\vdash \text{unpack } C \text{ as } \langle \alpha, x \rangle \text{ in } p_2 : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma_2)} \\
\\
\frac{\Delta'; \Gamma' \vdash p_1 : \exists \alpha. \sigma_1 \quad \vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta', \alpha; \Gamma', x:\sigma_1; \sigma_2) \quad \Delta'; \Gamma' \vdash \sigma_2}{\vdash \text{unpack } p_1 \text{ as } \langle \alpha, x \rangle \text{ in } C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma_2)} \\
\\
\frac{\vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma'[\mu \alpha. \sigma'/\alpha])}{\vdash \text{roll}_{\mu \alpha. \sigma'} C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \mu \alpha. \sigma')} \\
\\
\frac{\vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \mu \alpha. \sigma')}{\vdash \text{unroll } C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma'[\mu \alpha. \sigma'/\alpha])} \\
\\
\frac{\vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma_1)}{\vdash \text{ref } C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \text{ref } \sigma_1)} \quad \frac{\vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \text{ref } \sigma')}{\vdash !C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma')} \\
\\
\frac{\vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \text{ref } \sigma_1) \quad \Delta'; \Gamma' \vdash p_2 : \sigma_1}{\vdash C := p_2 : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \text{unit})} \\
\\
\frac{\Delta'; \Gamma' \vdash p_1 : \text{ref } \sigma_1 \quad \vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma_1)}{\vdash p_1 := C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \text{unit})} \\
\\
\frac{\vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \text{ref } \sigma') \quad \Delta'; \Gamma' \vdash p_2 : \text{ref } \sigma'}{\vdash C == p_2 : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \text{nat})} \\
\\
\frac{\Delta'; \Gamma' \vdash p_1 : \text{ref } \sigma' \quad \vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \text{ref } \sigma')}{\vdash p_1 == C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \text{nat})}
\end{array}$$

Figure 3.4 (cont.): Program contexts in $F^{\mu!}$.

Here, *alloc* determines the choice of the fresh location.

Fortunately, it is easy to show that if two programs are contextually equivalent under all deterministic allocators (*i.e.*, according to $\sim_{\text{ctx}}^{\text{alloc}}$), then they are contextually equivalent under a non-deterministic allocator (*i.e.*, according to \sim_{ctx}).

Theorem 1. If

$$\Delta; \Gamma \vdash p_1 \sim_{\text{ctx}}^{\text{alloc}} p_2 : \sigma$$

for any $\text{alloc} \in \text{Heap} \rightarrow \text{Loc}$ satisfying $\forall h. \text{alloc}(h) \notin \text{dom}(h)$, then:

$$\Delta; \Gamma \vdash p_1 \sim_{\text{ctx}} p_2 : \sigma$$

Proof. We are given a heap h and a suitably typed context C and must show that $\langle h; C[p_1] \rangle \uparrow$ iff $\langle h; C[p_2] \rangle \uparrow$. We sketch the “only if” direction here, the proof of the “if” direction is symmetric.

Suppose $\langle h; C[p_1] \rangle \uparrow$. We can look at the execution trace and observe the heap locations that are allocated (and the order in which they are allocated) and choose any deterministic allocator *alloc* that implements this allocation “strategy”.

Instantiating the premise then tells us that $\Delta; \Gamma \vdash p_1 \sim_{\text{ctx}}^{\text{alloc}} p_2 : \sigma$. Since $\langle h; C[p_1] \rangle$ diverges according to the *alloc*-determinized semantics (by choice of *alloc*), we learn that $\langle h; C[p_2] \rangle$ diverges according to the determinized semantics as well. It therefore also diverges according to the non-deterministic semantics, *i.e.*, $\langle h; C[p_2] \rangle \uparrow$. \square

In the remainder of this chapter, we therefore assume a valid allocator *alloc* as an axiom and work with the correspondingly determinized version of $\mathbf{F}^{\mu!}$. From now on, we simply write \sim_{ctx} for $\sim_{\text{ctx}}^{\text{alloc}}$ and never talk about the original \sim_{ctx} again.

3.4 Global vs. Local Knowledge

Our method of parametric bisimulations (PBs) is essentially *coinductive*, following the style of previous bisimulation techniques in many respects. Coinductive reasoning makes it easy to deal with recursive features—such as recursive types and higher-order store—without requiring the use of step-indexed constructions.

The two main ways in which PBs differ from previous bisimulation techniques are in their treatment of:

Local state. From recent work on KLRs [4, 22], we borrow the idea of using state transition systems (STSs) to establish invariants on how a module’s local state may evolve over time. STSs enable one to encode more flexible state invariants than are expressible using “environmental” bisimulations [74, 79].

Higher-order functions. In order to reason about higher-order functions in a coinductive style, but without confining ourselves to single-language reasoning, we employ a novel technical idea: *global vs. local knowledge*.

The treatment of local state using state transition systems follows prior work very closely, which is already discussed in Chapter 2. Here we focus attention on motivating our idea of global vs. local knowledge.

Coinductive Reasoning Recall the way we defined contextual equivalence in the previous section. An alternative but equivalent way of formulating contextual equivalence is as the *largest adequate congruence relation* [64]. Being adequate means that if two terms of base type are related, then either they both diverge (run forever) or they both evaluate to the same value (*e.g.*, if one term evaluates to 3, then the other must evaluate to 3 as well). Being a congruence means the relation is closed under all the constructs of the language (*e.g.*, if f_1 and f_2 are related at $\tau' \rightarrow \tau$, and e_1 and e_2 are related at τ' , then $f_1 e_1$ and $f_2 e_2$ are related at τ).

To prove using coinduction that two terms e_1 and e_2 are contextually equivalent at type τ , one must exhibit a (type-indexed) term relation L that contains (τ, e_1, e_2) and then prove that L is an adequate congruence. The relation L serves as a “generalized coinduction hypothesis”, by which one proves equivalence for all pairs of terms related by L simultaneously. However, while it is possible for one to employ this kind of “brute-force” coinductive proof, it is typically not very pleasant, because proving a relation to be a congruence directly can be incredibly tedious.

Bisimulation techniques help make coinductive proofs manageable by lightening the congruence proof burden. Typically, this is achieved by only requiring one to show that L is closed under type-directed *uses* (*i.e.*, evaluation or deconstruction) of the terms it relates. This results in proof obligations that look like the following (we are not being totally formal here):

- (1) If $(\tau, e_1, e_2) \in L$, then either $e_1 \uparrow$ and $e_2 \uparrow$,
or $\exists v_1, v_2. e_1 \hookrightarrow^* v_1$ and $e_2 \hookrightarrow^* v_2$ and $(\tau, v_1, v_2) \in L$.
- (2) If $(\text{int}, v_1, v_2) \in L$, then $\exists n. v_1 = v_2 = n$.
- (3) If $(\tau' \times \tau'', v_1, v_2) \in L$, then $\exists v'_1, v''_1, v'_2, v''_2. v_i = \langle v'_i, v''_i \rangle$
and $(\tau', v'_1, v'_2) \in L$ and $(\tau'', v''_1, v''_2) \in L$.
- (4) If $(\mu\alpha. \tau, v_1, v_2) \in L$, then $\exists v'_1, v'_2. v_i = \text{roll } v'_i$
and $(\tau[\mu\alpha. \tau/\alpha], v'_1, v'_2) \in L$.

The most problematic proof obligation is the one for function values. It usually looks something like this (simplifying fix to λ):

- (5) If $(\tau' \rightarrow \tau, v_1, v_2) \in L$, then $\exists x, e_1, e_2. v_i = \lambda x. e_i$ and
 $\forall v'_1, v'_2. (\tau', v'_1, v'_2) \in \boxed{G} \Rightarrow (\tau, e_1[v'_1/x], e_2[v'_2/x]) \in L$.

In other words, if L relates function values v_1 and v_2 , then applying them to any “equivalent arguments” v'_1 and v'_2 should produce results that are also related by L . The big question is: what is this relation G from which the arguments v'_1 and v'_2 are drawn?

Global vs. Local Knowledge First, some (non-standard) terminology: There exist many equivalent terms, but when we do a bisimulation proof, we only make a claim about some of them. So let us make a distinction between “local” and “global”

knowledge about term equivalence. The relation L describes our *local knowledge*: these are the terms whose equivalence we aim to validate in our proof. The relation G , on the other hand, embodies the *global knowledge* about all terms out there that are equivalent. In proof obligation (5), we draw equivalent function arguments from G (rather than L) since they might indeed originate from “somewhere else” in the program (some unknown client code), and thus our local knowledge L may not be sufficient to justify their equivalence. This leaves us with the question of how to define G .

Whence Global Knowledge? Coming up with a sound (and practically usable) choice for G is far from obvious, and existing bisimulation methods make a variety of different choices. For example:

- *Applicative* bisimulations [1] define G to be the syntactic identity relation on closed values.

This is a nice, simple choice, which works well for pure λ -calculus. Unfortunately, for higher-order stateful languages like $F^{\mu!}$, it is unsound [40], so more advanced approaches are needed:

- *Environmental* bisimulations [81, 41, 74, 79] take G to be the “context closure” of L , *i.e.*, the relation that extends the syntactic identity relation on closed values by including closures of open values v with pairs of values (w_1, w_2) that are related by L (formally: $\{(\sigma, v[w_1/y], v[w_2/y]) \mid \{(\sigma', w_1, w_2)\} \subseteq L\}$).¹
- *Normal form* (or *open*) bisimulations [43, 78, 44, 45] sidestep the whole question by choosing a fresh variable name x and representing equivalent arguments by the same x . As a result, these bisimulations are built over open terms, and proof obligation (1) above must be updated to account for the possibility that the evaluations of e_1 and e_2 get stuck trying to deconstruct the same free variable x (more about that below).

All of these methods define global knowledge in a very “syntactic” way that is well suited to proving contextual equivalences. However, we wish to develop a method that will be capable of generalizing to the setting of inter-language reasoning, where G may relate different languages. We therefore seek an account of global knowledge that is more “semantic”.

Parameterizing Over Global Knowledge The essential difficulty in choosing G has to do with higher-order functions: if the argument type τ' is (or contains) a function type, then equivalence at τ' is very hard to characterize directly.² Our

¹We are glossing over a lot of details here. To be precise, environmental bisimulations are actually *sets* of L ’s. For more details, see Section 6.2.

²Conversely, if τ' were arrow-free (*e.g.*, in a first-order language), it would be easy to characterize equivalence at τ' directly.

solution is simple: we don't try to define the global knowledge at all; instead, we take G to be a *parameter* of our model!

Our key observation is that it is not necessary to pin down exactly what G is, so long as we make our coinductive proof for L as parametric as possible with respect to it. (We will clarify what “as parametric as possible” means in Section 3.6.) This parametricity makes our proofs quite robust by allowing G to be instantiated in a variety of different ways. In particular, we make no assumptions whatsoever about the values that G relates at function type. For all we know, G might even include “garbage” like $(\text{int} \rightarrow \text{int}, 4, \langle \rangle)$.³

Our approach can be viewed as a more semantic account of the idea behind normal form bisimulations (see above), which is to model “equivalent arguments” as black boxes about which nothing is known. Consequently, just as for normal form bisimulations, we need to adapt proof obligation (1) above to account for the possibility that e_1 and e_2 get stuck. For normal form bisimulations, e_1 and e_2 may get stuck if they try to deconstruct a free variable x , and so normal form bisimulations loosen proof obligation (1) to allow e_1 and e_2 to reduce to terms of the form $K_1[x\ v_1]$ and $K_2[x\ v_2]$, where K_1 and K_2 are equivalent continuations and v_1 and v_2 are equivalent values. In our case, e_1 and e_2 may get stuck if they try to apply some bogus functions that are equivalent according to the global knowledge but that turn out (like 4 and $\langle \rangle$) to not even be functions. Hence, we will allow e_1 and e_2 to reduce to terms of the form $K_1[f_1\ v_1]$ and $K_2[f_2\ v_2]$, where K_1 and K_2 are equivalent continuations, and where $\{(\tau' \rightarrow \tau, f_1, f_2), (\tau', v_1, v_2)\} \subseteq G$. In this way, the parameter G serves as a semantic analogue of free variables in normal form bisimulations.

Intuitively, although the idea of parameterizing over the global knowledge may seem surprising at first, we find it to be comfortingly reminiscent of Girard's method for modeling System F [27]. In Girard's method, a potential cycle in the definition of the logical relation for impredicative universal types $\forall \alpha. \tau$ is avoided by parameterizing over an arbitrary relational interpretation of the abstract type α . In our scenario, the problem of how to define the global knowledge is avoided by parameterizing over an arbitrary relational interpretation of *function types*. In essence, we are treating a function type $\tau_1 \rightarrow \tau_2$ as an unusual kind of abstract type: the coinductive proofs about different “modules” in a program all treat the global interpretation of $\tau_1 \rightarrow \tau_2$ abstractly, while simultaneously they each contribute to defining it.

Parameterizing over the global knowledge turns out to be *very* useful. First and foremost, it makes it easy to soundly compose our coinductive proofs for different “modules” together (and hence prove soundness of our method w.r.t. contextual equivalence). Second, it enables us to reason about open terms (Section 3.5) and higher-order state invariants (Section 3.6), replacing the use of context closure or free variables for those purposes in environmental and normal form bisimulations, respectively. Finally, it is the key to establishing transitivity for our proof method (Sections 3.5.4 and 3.9).

³The ability to instantiate G with a “trashy” relation is surprisingly useful. We will make critical use of it in our transitivity proof.

3.5 Warmup: Parametric Bisimulations for λ^μ

To ease the presentation of parametric bisimulations (PBs), we begin in this section by using the idea of global vs. local knowledge, motivated in the previous section, to define a relational model for λ^μ , a pure fragment of $F^{\mu!}$. This sub-language is obtained by restricting the definitions in Section 3.3 to base, function, product, sum, and recursive types (*i.e.*, by excluding universal, existential, and reference types, as well as the associated constructs and rules). Regarding reduction, the heap component h becomes unused and is erased. Regarding typing, the type variable environment Δ becomes unused and is erased.

Due to the simplicity of the language, the model developed here does not feature STSs. These will come into play when dealing with state in Section 3.6. However, by ignoring the transition-systems aspect of PBs for the time being, we can focus attention on other aspects of the model.

3.5.1 Definitions

Figure 3.5 lists the various semantic domains we will be using. Here, CTy denotes the set of closed types of λ^μ (*i.e.*, types with no free type variables α), CVal denotes the set of closed values (*i.e.*, values with no free term variables x), etc. The first four definitions are standard: relations on values, expressions, continuations, and heaps, indexed by the relevant types (in case of KRel , input and output types).

Next, we define what we call the *flexible* types, CTyF , along with the *flexible relations*, VRelF , which are just relations on closed values indexed by such flexible types. Whereas bisimulation methods typically allow terms of arbitrary type to be included in the bisimulation, we find it useful to restrict local and global knowledges to relate only values of “flexible” types. Intuitively, these are the types at which value equivalence may depend on module-specific knowledge. In $F^{\mu!}$, there will be several kinds of flexible types, but in λ^μ , the only flexible types are function types.

In contrast, value equivalence at the remaining types—base, product, sum, and recursive types, which we call *rigid*—is fixed and agreed upon by all modules once the meaning of the flexible types is defined. This is achieved by a closure operation that takes a relation $R \in \text{VRelF}$ and returns its value closure $\bar{R} \in \text{VRel}$. It is defined as the least fixed-point of the set of equations in Figure 3.6. For instance, two pairs are related by \bar{R} if and only if they are related componentwise. Note that R only occurs covariantly⁴, so \bar{R} is inductively well defined, even though the type gets bigger on the r.h.s. in the case of $\mu\alpha. \tau$.

Local and Global Knowledge As shown in Figure 3.7, a local knowledge $L \in \text{LK}$ is essentially a flexible relation, except that this relation is actually itself parameterized by the global knowledge, G . In effect, $L(G)$ describes the values that we wish

⁴We don’t need to worry about functions introducing contravariance because the relation on function types is already given by R itself.

\mathbf{VRel}	$:= \mathbf{CTy} \rightarrow \mathcal{P}(\mathbf{CVal} \times \mathbf{CVal})$
\mathbf{ERel}	$:= \mathbf{CTy} \rightarrow \mathcal{P}(\mathbf{CExp} \times \mathbf{CExp})$
\mathbf{KRel}	$:= \mathbf{CTy} \times \mathbf{CTy} \rightarrow \mathcal{P}(\mathbf{CCont} \times \mathbf{CCont})$
\mathbf{HRel}	$:= \mathcal{P}(\mathbf{Heap} \times \mathbf{Heap})$
\mathbf{CTyF}	$:= \{(\tau_1 \rightarrow \tau_2) \in \mathbf{CTy}\}$
\mathbf{VRelF}	$:= \mathbf{CTyF} \rightarrow \mathcal{P}(\mathbf{CVal} \times \mathbf{CVal})$

Figure 3.5: Semantic domains for λ^μ .

$\overline{R}(\tau)$	$:= R(\tau) \quad \text{if } \tau \in \mathbf{CTyF}$
$\overline{R}(\mathbf{unit})$	$:= \{(\langle \rangle, \langle \rangle)\}$
$\overline{R}(\mathbf{nat})$	$:= \{(n, n)\}$
$\overline{R}(\tau_1 \times \tau_2)$	$:= \{(\langle v_1, v'_1 \rangle, \langle v_2, v'_2 \rangle) \mid (v_1, v_2) \in \overline{R}(\tau_1) \wedge (v'_1, v'_2) \in \overline{R}(\tau_2)\}$
$\overline{R}(\tau_1 + \tau_2)$	$:= \{(\mathbf{inl } v_1, \mathbf{inl } v_2) \mid (v_1, v_2) \in \overline{R}(\tau_1)\} \cup \{(\mathbf{inr } v_1, \mathbf{inr } v_2) \mid (v_1, v_2) \in \overline{R}(\tau_2)\}$
$\overline{R}(\mu\alpha. \tau)$	$:= \{(\mathbf{roll } v_1, \mathbf{roll } v_2) \mid (v_1, v_2) \in \overline{R}(\tau[\mu\alpha. \tau/\alpha])\}$

Figure 3.6: Value closure for λ^μ (if $R \in \mathbf{VRelF}$, then $\overline{R} \in \mathbf{VRel}$).

$\mathbf{step}(e)$	$:= \begin{cases} e' & \text{if } e \hookrightarrow e' \\ \perp & \text{otherwise} \end{cases}$
\mathbf{FixVal}	$:= \{f \in \mathbf{CVal} \mid \forall v. \mathbf{step}(f v) \neq \perp\}$
$R' \supseteq R$	$:= \forall \tau. R'(\tau) \supseteq R(\tau)$
\mathbf{LK}	$:= \{L \in \mathbf{VRelF} \rightarrow \mathbf{VRelF} \mid L \text{ is monotone w.r.t. } \subseteq \wedge \\ \forall G, \tau. L(G)(\tau) \subseteq \mathbf{FixVal} \times \mathbf{FixVal}\}$
$\mathbf{GK}(L)$	$:= \{G \in \mathbf{VRelF} \mid G \supseteq L(G)\}$

Figure 3.7: Specification of local and global knowledges for λ^μ .

to prove are equivalent, assuming that G correctly represents the global knowledge. This parameterization is necessary in order to reason about open terms, and we will see its utility below in the proof of compatibility for fix. We require that L is monotone w.r.t. G : intuitively, passing in a larger global knowledge should never result in fewer terms being related by L .

We also require that the values related by the local knowledge at function type are indeed functions, in the sense that their application to an arbitrary value should not be stuck, but should be reducible at least for one step. This is a technical requirement that is used in our transitivity proof in Section 3.5.4. The idea is that there should be classes of values that may be related by the global knowledge but not by the local knowledge.

We want to restrict attention to global knowledges that are closed w.r.t. the local knowledge in question: For a particular L , we define $\mathbf{GK}(L)$ to be the set of flexible relations G s.t. $G \supseteq L(G)$. This requirement makes sense since the global knowledge must by definition be a superset of any local knowledge. Observe, however, that we do not restrict what other values G relates. Indeed, G may relate values at function type that are not actually functions, or that are obviously inequivalent (e.g., 4 and $\langle \rangle$, cf. Section 3.4).

Relating Expressions and Continuations So far, we have seen how local knowledges describe equivalence between values of flexible type, and how this equivalence can be lifted to arbitrary types.

Figure 3.8 shows how equivalence is defined for (closed) expressions, e , and continuations, K . Specifically, we introduce two new relations, $\mathbf{E} \in \mathbf{VRelF} \rightarrow \mathbf{ERel}$ and $\mathbf{K} \in \mathbf{VRelF} \rightarrow \mathbf{KRel}$, which are defined coinductively.

Given a type τ , a local knowledge $L \in \mathbf{LK}$, and a global knowledge $G \in \mathbf{GK}(L)$, we say that two expressions are “locally” equivalent, written $(e_1, e_2) \in \mathbf{E}(G)(\tau)$, if they either both diverge or both terminate producing related values. Along the way, however, they may make calls to “external” functions, that is, functions that are related by G , but not necessarily by the local knowledge $L(G)$. More precisely, we say two closed expressions are locally equivalent if and only if one of the following three cases holds:

Case DIV. Both expressions diverge (run forever).

Case EVAL. Both expressions run successfully to completion, producing related values.

Case CALL. Both expressions reduce after some number of steps to some expressions of the form $K_i[f_i v_i]$, where both the f_i and v_i are related by the global knowledge G at the appropriate types, and the continuations, K_1 and K_2 , are equivalent. We say that two continuations are equivalent if instantiating them with equivalent values (according to the global knowledge G) yields equivalent expressions.

$$\begin{aligned}
\mathbf{E}(G)(\tau) &:= \{(e_1, e_2) \mid \\
&\quad (\text{DIV}) \quad e_1 \uparrow \wedge e_2 \uparrow \\
&\quad \vee \quad (\text{EVAL}) \quad \exists v_1, v_2. e_1 \hookrightarrow^* v_1 \wedge e_2 \hookrightarrow^* v_2 \wedge (v_1, v_2) \in \overline{G}(\tau) \\
&\quad \vee \quad (\text{CALL}) \quad \exists \tau', K_1, K_2, e'_1, e'_2. \\
&\quad \quad e_1 \hookrightarrow^* K_1[e'_1] \wedge e_2 \hookrightarrow^* K_2[e'_2] \wedge \\
&\quad \quad (e'_1, e'_2) \in \mathbf{U}(G, G)(\tau') \wedge \\
&\quad \quad (K_1, K_2) \in \mathbf{K}(G)(\tau', \tau)\} \\
\\
\mathbf{K}(G)(\tau_1, \tau_2) &:= \{(K_1, K_2) \mid \forall (v_1, v_2) \in \overline{G}(\tau_1). (K_1[v_1], K_2[v_2]) \in \mathbf{E}(G)(\tau_2)\} \\
\\
\mathbf{U}(R, R')(\tau) &:= \{(v_1 \ v'_1, v_2 \ v'_2) \mid \exists \tau'. (v_1, v_2) \in R(\tau' \rightarrow \tau) \wedge (v'_1, v'_2) \in \overline{R'}(\tau')\} \\
\\
\text{consistent}(L) &:= \forall G \in \mathbf{GK}(L). \forall \tau. \forall (e_1, e_2) \in \mathbf{U}(L(G), G)(\tau). \\
&\quad (\text{step}(e_1), \text{step}(e_2)) \in \mathbf{E}(G)(\tau) \\
\\
\Gamma \vdash e_1 \sim_L e_2 : \tau &:= \text{consistent}(L) \wedge \forall G \in \mathbf{GK}(L). \forall \gamma_1, \gamma_2 \in \text{dom}(\Gamma) \rightarrow \mathbf{CVal}. \\
&\quad (\forall x: \tau' \in \Gamma. (\gamma_1 x, \gamma_2 x) \in \overline{G}(\tau')) \implies (\gamma_1 e_1, \gamma_2 e_2) \in \mathbf{E}(G)(\tau) \\
\\
\Gamma \vdash e_1 \sim e_2 : \tau &:= \exists L. \Gamma \vdash e_1 \sim_L e_2 : \tau \\
\\
\Gamma \vdash p_1 \sim p_2 : \tau &:= \Gamma \vdash |p_1| \sim |p_2| : \tau
\end{aligned}$$

Figure 3.8: Mutually coinductive definitions of expression equivalence, $\mathbf{E} \in \mathbf{VRelF} \rightarrow \mathbf{ERel}$, and continuation equivalence, $\mathbf{K} \in \mathbf{VRelF} \rightarrow \mathbf{KRel}$, and definitions of consistency and program equivalence for λ^μ .

Intuitively, expressions that are locally equivalent w.r.t global knowledge G do not necessarily have the same observable behavior, but they can be understood to have the same local behavior, *i.e.*, behave equivalently modulo what happens during calls to functions related by G . If G happens to relate functions that behave differently, this is G 's fault, not ours.

Consistency and Program Equivalence We say that a local knowledge L is **consistent** (in Figure 3.8) if and only if any functions that it declares equivalent yield in fact equivalent expressions when applied to equivalent arguments. In the formal definition, we parameterize over an arbitrary global knowledge $G \in \mathbf{GK}(L)$; the auxiliary construction \mathbf{U} (characterizing the “uses” of functions) draws the functions being tested for equivalence from the local knowledge, $L(G)$, while the arguments to which they are applied are drawn from the global knowledge, G .

In order for PB's coinductive reasoning to be sound, we must be quite careful in the definition of consistency: for each pair of functions (v_1, v_2) related by $L(G)$, and arguments (v'_1, v'_2) related by \overline{G} , we cannot simply require $(v_1 v'_1, v_2 v'_2)$ to be related by $\mathbf{E}(G)$ because, via **CALL**, this is a tautology! Instead, we demand that $v_i v'_i \hookrightarrow e_i$, and that e_1 and e_2 are related by $\mathbf{E}(G)$ (see the use of *step* in the formal definition). Requiring the terms to take a step of reduction at this point ensures that “progress” is made in the coinductive argument, *i.e.*, that the coinduction is *guarded*. We will say more about this guardedness condition in Section 3.10.

We call two expressions equivalent at type τ in context Γ , written $\Gamma \vdash e_1 \sim e_2 : \tau$, if and only if there exists a consistent local knowledge L that shows that $\gamma_1 e_1$ and $\gamma_2 e_2$ are equivalent at type τ for value substitutions γ_1 and γ_2 and an arbitrary global knowledge G for L . (We write γe for substituting in e according to γ .) These substitutions are also arbitrary except that they must be pointwise related by G at Γ .

Finally, two programs are equivalent simply if and only if their type-erased versions are equivalent expressions.

A note on coinduction. As \mathbf{E} and \mathbf{K} are defined mutually dependent over a complete lattice (powerset lattice lifted pointwise to function spaces) and all operations involved are monotone, we can take the meaning of these definitions to be either the least or the greatest fixed point. We choose the greatest fixed point, corresponding to coinduction, because this can in principle relate more terms and is somewhat easier to work with. This means that PBs actually come with *two* coinductive reasoning principles:

1. The first is enabled by \mathbf{E} 's **CALL** case and the use of the global knowledge, which contains the local knowledge. Here, the coinductive assumption is about L (and thus G). This coinduction principle is used all the time in examples, for instance in Section 3.5.2 below. For a quick intuition, think of reasoning about recursive functions: we construct a local knowledge relating them and then need to show its consistency. In the proof of consistency, after taking a

step, we may get to recursive calls. At that point, we are allowed to make use of the **CALL** since we already know that the functions are related by the global knowledge (after all, it contains the local knowledge that we constructed).

2. The second is enabled by the definition of **E** as a greatest fixed point, as we just discussed above. Here, the coinductive assumption is about **E**. This coinduction principle is only used in the metatheory (*e.g.*, Lemma 3 and its generalization Lemma 11), not in any example proofs. The reason is that in λ^μ and F^μ there is no way for a program to loop other than via function calls, so we can always use the first principle. Were we to extend the language with other forms of recursion (such as while loops or primitive recursion), the coinductive interpretation of **E** would become essential.

When presenting PILS in Chapter 4, we will see that it is possible to get away with a single form of coinductive reasoning, without losing expressivity.

3.5.2 Example Proof

Consider the following example concerning streams as functions (taken from Sumii and Pierce [81]):

$$\begin{aligned}
 \tau &:= \mu\alpha. \text{unit} \rightarrow \text{nat} \times \alpha \\
 \text{ones} : \text{unit} \rightarrow \text{nat} \times \tau &:= \text{fix } f(x). \langle 1, \text{roll } f \rangle \\
 \text{twos} : \text{unit} \rightarrow \text{nat} \times \tau &:= \text{fix } f(x). \langle 2, \text{roll } f \rangle \\
 \text{succ} : \tau \rightarrow \tau &:= \text{fix } f(s). \text{let } \langle n, s' \rangle = \text{unroll } s \text{ in} \\
 &\quad \text{roll } \lambda x. \langle n+1, f \ s' \rangle
 \end{aligned}$$

(Here and elsewhere we write $\lambda x. e$ short for $\text{fix } g(x). e$, where g does not occur free in e .) As the names suggest, *ones* encodes the stream consisting of only 1's (*i.e.*, 111...), while *twos* encodes the stream consisting of only 2's. The goal is to show that incrementing each component of *ones* by one results in a stream equivalent to *twos*:

$$\vdash \text{roll } \text{twos} \sim \text{succ } (\text{roll } \text{ones}) : \tau$$

Constructing a Suitable Local Knowledge. Note that we have

$$\text{succ } (\text{roll } \text{ones}) \hookrightarrow^* \text{roll } \text{twos}'$$

for $\text{twos}' := \lambda x. \langle 1+1, \text{succ } (\text{roll } \text{ones}) \rangle$. We define a local knowledge L that relates exactly *twos* and *twos'*:

$$L(G)(\tau') := \{(\text{twos}, \text{twos}') \mid \tau' = \text{unit} \rightarrow \text{nat} \times \tau\}$$

Proving Its Consistency. Suppose $(e_1, e_2) \in \mathbf{U}(L(G), G)(\tau')$, where $G \in \mathbf{GK}(L)$. By construction of L we know that $e_1 = \text{twos } \langle \rangle$, $e_2 = \text{twos}' \langle \rangle$, and $\tau' = \text{nat} \times \tau$. Hence we must show:

$$(\langle 2, \text{roll twos} \rangle, \langle 1+1, \text{succ } (\text{roll ones}) \rangle) \in \mathbf{E}(G)(\text{nat} \times \tau)$$

Using the **EVAL** case in \mathbf{E} and the definition of \overline{G} , this reduces to showing $(2, 2) \in \overline{G}(\text{nat})$ and $(\text{roll twos}, \text{roll twos}') \in \overline{G}(\tau)$. The former is trivial. The latter is equivalent to $(\text{twos}, \text{twos}') \in \overline{G}(\text{unit} \rightarrow \text{nat} \times \tau)$, which holds by construction because G extends $L(G)$.

Showing the Programs Related By It. It remains to show:

$$\forall G \in \mathbf{GK}(L). (\text{roll twos}, \text{succ } (\text{roll ones})) \in \mathbf{E}(G)(\tau)$$

Again using **EVAL**, we end up having to show $(\text{roll twos}, \text{roll twos}') \in \overline{G}(\tau)$, which we have already done above.

3.5.3 Basic Properties and Soundness

We move on to some properties of our constructions. These lemmas will be generalized in a later section (3.7) to the full $\mathbf{F}^{\mu!}$ setting.

The first lemma says that a global knowledge that is valid for the union of two local knowledges is also valid for each local knowledge in isolation. (The reason why we don't express this more generally using \subseteq will become clear in Section 3.7.)

Lemma 1. If $L_1, L_2 \in \mathbf{LK}$ and $G \in \mathbf{GK}(L_1 \cup L_2)$, then $G \in \mathbf{GK}(L_1) \cap \mathbf{GK}(L_2)$.

The following lemma states that consistency of local knowledges is preserved under (pointwise) union. This is important for ensuring that equivalence proofs for different subterms, which rely on different local knowledges, can be soundly composed.

Lemma 2. If $\text{consistent}(L)$ and $\text{consistent}(L')$, then $\text{consistent}(L \cup L')$.

Figure 3.9 shows some of the basic properties of our program equivalence relation. First, we have a set of rules stating that equivalence is compatible with all the language constructs. These rules state that if two terms start with the same term constructor and their immediate subterms are component-wise equivalent, then so are the composite terms. For brevity, we just present the rules for recursive function definition (**FIX**) and function application (**APP**), whose proofs are the most interesting.

The proof of **FIX** is particularly interesting because it requires coinductive reasoning: when showing the consistency of our local knowledge, *i.e.*, when proving applications of the recursive functions related, we deal with the recursion by resorting to the fact that the local knowledge relates the functions.

$$\begin{array}{c}
\frac{\Gamma, f:(\tau' \rightarrow \tau), x:\tau' \vdash e_1 \sim e_2 : \tau}{\Gamma \vdash \text{fix } f(x). e_1 \sim \text{fix } f(x). e_2 : \tau' \rightarrow \tau} \text{FIX} \\
\\
\frac{\Gamma \vdash e_1 \sim e_2 : \tau' \rightarrow \tau \quad \Gamma \vdash e'_1 \sim e'_2 : \tau'}{\Gamma \vdash e_1 e'_1 \sim e_2 e'_2 : \tau} \text{APP} \\
\\
\frac{\Gamma \vdash p : \tau}{\Gamma \vdash |p| \sim |p| : \tau} \text{REFL} \quad \frac{\Gamma \vdash e_2 \sim e_1 : \tau}{\Gamma \vdash e_1 \sim e_2 : \tau} \text{SYMM} \\
\\
\frac{\Gamma \vdash e_1 \sim e_2 : \tau \quad \Gamma \vdash e_2 \sim e_3 : \tau}{\Gamma \vdash e_1 \sim e_3 : \tau} \text{TRANS} \\
\\
\frac{\Gamma \vdash e_1 \sim e_2 : \tau \quad \vdash C : (\Gamma; \tau) \rightsquigarrow (\Gamma'; \tau')}{\Gamma' \vdash |C|[e_1] \sim |C|[e_2] : \tau'} \text{CONG} \\
\\
\frac{\Gamma, x:\tau' \vdash e_1 \sim e_2 : \tau \quad \Gamma \vdash v_1 \sim v_2 : \tau'}{\Gamma \vdash e_1[v_1/x] \sim e_2[v_2/x] : \tau} \text{SUBST} \\
\\
\frac{\Gamma \vdash e'_1 \sim e'_2 : \tau \quad \forall \gamma \in \text{dom}(\Gamma) \rightarrow \text{CVal}. \gamma e_1 \hookrightarrow^* \gamma e'_1 \quad \forall \gamma \in \text{dom}(\Gamma) \rightarrow \text{CVal}. \gamma e_2 \hookrightarrow^* \gamma e'_2}{\Gamma \vdash e_1 \sim e_2 : \tau} \text{EXPAND} \\
\\
\frac{\Gamma, x:\tau' \vdash e_1 \sim e_2 : \tau \quad \Gamma \vdash v_1 \sim v_2 : \tau'}{\Gamma \vdash (\lambda x. e_1) v_1 \sim e_2[v_2/x] : \tau} \text{BETA}
\end{array}$$

Figure 3.9: Some basic properties of our equational model.

More formally, from the premise there exists L such that $\Gamma, f:(\tau' \rightarrow \tau), x:\tau' \vdash e_1 \sim_L e_2 : \tau$ and $\text{consistent}(L)$. We define:

$$L'(G)(\tau_1) := \{(\gamma_1 \text{fix } f(x). e_1, \gamma_2 \text{fix } f(x). e_2) \mid \\ \tau_1 = \tau' \rightarrow \tau \wedge \gamma_i \in \text{dom}(\Gamma) \rightarrow \text{CVal} \wedge \\ \forall y:\tau_2 \in \Gamma. (\gamma_1(y), \gamma_2(y)) \in \overline{G}(\tau_2)\}$$

Note here how the parameterization of L' over G provides it with a source from which to construct the closing substitutions γ_1 and γ_2 .

The goal now is to prove

$$\Gamma \vdash \text{fix } f(x). e_1 \sim \text{fix } f(x). e_2 : \tau' \rightarrow \tau.$$

The local knowledge that we choose is $L \cup L'$. Showing that L' (and thus $L \cup L'$) relates any appropriately closed instances of the two functions is simply a matter of unfolding definitions (by construction of L'). It therefore remains to establish $\text{consistent}(L \cup L')$. By Lemma 2, this boils down to showing

$$(\gamma'_1 e_1, \gamma'_2 e_2) \in \mathbf{E}(G)(\tau)$$

for any $G \in \mathbf{GK}(L \cup L')$, where

- $\gamma_i \in \text{dom}(\Gamma) \rightarrow \text{CVal}$,
- $\forall y:\tau' \in \Gamma. (\gamma_1(y), \gamma_2(y)) \in \overline{G}(\tau')$,
- $(v_1, v_2) \in \overline{G}(\tau')$,
- $\gamma'_i = \gamma_i, f \mapsto (\gamma_i \text{fix } f(x). e_i), x \mapsto v_i$.

Since $(\gamma_1 \text{fix } f(x). e_1, \gamma_2 \text{fix } f(x). e_2) \in L'(G)(\tau' \rightarrow \tau) \subseteq \overline{G}(\tau' \rightarrow \tau)$ and by Lemma 1 $G \in \mathbf{GK}(L)$, we can instantiate $\Gamma, f:\tau' \rightarrow \tau, x:\tau' \vdash e_1 \sim_L e_2 : \tau$ with γ'_i and are done.

The proof of rule APP relies on Lemma 2 as well, in order to show that the consistent local knowledges for its two premises combine to form a consistent local knowledge for the conclusion. In addition, the proof relies on the following lemma about plugging equivalent expressions or continuations into equivalent continuations, proved by mutual coinduction and case analysis.

Lemma 3. If $(K_1, K_2) \in \mathbf{K}(G)(\tau', \tau)$, then:

1. $(e_1, e_2) \in \mathbf{E}(G)(\tau') \implies (K_1[e_1], K_2[e_2]) \in \mathbf{E}(G)(\tau)$
2. $(K'_1, K'_2) \in \mathbf{K}(G)(\tau'', \tau') \implies (K_1[K'_1], K_2[K'_2]) \in \mathbf{K}(G)(\tau'', \tau)$

In the proof of APP, we apply the first case of this lemma with $e_i := \gamma_i e_i$ and $K_i := \bullet \gamma_i e'_i$, which leaves us to prove K_1 and K_2 to be equivalent according to \mathbf{K} . Unfolding the definition of \mathbf{K} , we have to show that for arbitrary equivalent values v_1 and v_2 , $(v_1 \gamma_1 e'_1, v_2 \gamma_2 e'_2)$ is in \mathbf{E} , for which we apply Lemma 3 again, this time

with $e_i := \gamma_i e'_i$ and $K_i := v_i \bullet$. Then we are left to prove $v_1 \bullet$ and $v_2 \bullet$ equivalent, *i.e.*, that $(v_1 v'_1, v_2 v'_2)$ is in \mathbf{E} for arbitrary equivalent values v'_1 and v'_2 , which follows from the third disjunct of the \mathbf{E} definition.

As a consequence of these “compatibility” rules, by a straightforward induction on the typing derivation, we can show that equivalence is reflexive on well-typed programs (rule REFL). This corresponds to the “fundamental property” of logical relations. Equivalence is also symmetric: this follows trivially from the symmetric nature of our definition. Transitivity holds as well, but its much more subtle and warrants its own section, namely Section 3.5.4.

By induction on the typing derivation of contexts, we can show that our equivalence is a congruence: if two equivalent terms are placed in the same contexts, the resulting compositions are equivalent. Next, we have a substitutivity property for values, an expansion law for pure execution steps, and finally a direct corollary of these two, namely β -equivalence (on value arguments).

We move to a key lemma about \mathbf{E} . Given a consistent local knowledge L , if the global knowledge extends L with some additional external knowledge \mathcal{R} , then the **CALL** case in the definition of \mathbf{E} can be restricted so that it applies only to *external* function calls (*i.e.*, calls to functions related by \mathcal{R}).

Lemma 4 (External call). For any $\text{consistent}(L)$, any $G \in \text{GK}(L)$ and $\mathcal{R} \in \text{VRelF}$, we have:

$$G = L(G) \cup \mathcal{R} \implies \mathbf{E}(G) = \mathbf{E}^{\mathcal{R}}(G)$$

where the definition of $\mathbf{E}^{\mathcal{R}}$ is the same as \mathbf{E} except that, in the third disjunct, $\mathbf{U}(G, G)$ is replaced by $\mathbf{U}(\mathcal{R}, G)$.

The \supseteq follows directly from the observation that $\mathcal{R} \subseteq G$. To prove the other direction, we essentially have to eliminate all uses of the third disjunct of \mathbf{E} where the functions being invoked are related by $G \setminus \mathcal{R}$. Since all such functions are by definition in $L(G)$, and since we know $\text{consistent}(L)$, we can in fact always “inline” the equivalence proofs for all such function calls.

For $G = \mu R.L(R)$ (denoting the least fixed point of L) and $\mathcal{R} = \emptyset$, Lemma 4 yields *adequacy*, which says that equivalent closed terms either both diverge or both terminate returning proper values. Note that this implies safety: the terms never get stuck during evaluation.

Lemma 5 (Adequacy). If $\vdash e_1 \sim e_2 : \tau$, then:

$$(e_1 \uparrow \wedge e_2 \uparrow) \vee \exists v_1, v_2. e_1 \hookrightarrow^* v_1 \wedge e_2 \hookrightarrow^* v_2$$

Finally, combining adequacy and congruence, we show our main soundness theorem: for well-typed programs, our equivalence relation is included in contextual equivalence.

Theorem 2 (Soundness). If $\Gamma \vdash p_1 : \tau$ and $\Gamma \vdash p_2 : \tau$, then:

$$\Gamma \vdash p_1 \sim p_2 : \tau \implies \Gamma \vdash p_1 \sim_{\text{ctx}} p_2 : \tau$$

3.5.4 Transitivity

We now briefly sketch our proof of transitivity. We will go into more details in Section 3.9 presenting the proof of transitivity of our PB model for full F^μ .

It suffices to show the following lemma:

Lemma 6. If $\Gamma \vdash e_1 \sim_{L_1} e_2 : \tau$ and $\Gamma \vdash e_2 \sim_{L_2} e_3 : \tau$, then there exists L such that $\Gamma \vdash e_1 \sim_L e_3 : \tau$.

Naturally, we can expect L to be some sort of composition of the given local knowledges L_1 and L_2 . Defining this composition is, however, quite subtle. The problem is that the local knowledge takes the global knowledge G as a parameter, but then what global knowledges G_{L_1} and G_{L_2} should be passed on to L_1 and L_2 , respectively, in constructing L ? Assuming for now we somehow pick G_{L_1} and G_{L_2} appropriately, L can naturally be defined as follows:

$$L(G)(\tau) := L_1(G_{L_1})(\tau) \circ L_2(G_{L_2})(\tau)$$

where \circ stands for ordinary relational composition.

The key part of the proof is showing transitivity of \mathbf{E} :

$$\forall G \in \mathbf{GK}(L). (e'_1, e'_2) \in \mathbf{E}(G_{L_1})(\tau) \wedge (e'_2, e'_3) \in \mathbf{E}(G_{L_2})(\tau) \implies (e'_1, e'_3) \in \mathbf{E}(G)(\tau)$$

In order to prove this, we want the disjunct of \mathbf{E} by which e'_1 and e'_2 are related to match the disjunct of \mathbf{E} by which e'_2 and e'_3 are related (recall the three disjuncts in the definition of \mathbf{E}). To illustrate, say e'_1 and e'_2 are related because they reduce to related values (second disjunct). Now consider the three cases regarding e'_2 and e'_3 :

Case DIV. They both diverge. Fortunately, this contradicts our assumption about e'_2 , so this case cannot arise.

Case EVAL. They are related for the same reason as e'_1 and e'_2 are—*i.e.*, e'_2 and e'_3 reduce to related values. This is the “good” case. Relying on determinacy of reduction, we are done if we can show transitivity of the value relation. Formally, we need to show that $\overline{G_{L_1}}(\tau) \circ \overline{G_{L_2}}(\tau) \subseteq \overline{G}(\tau)$.

Case CALL. They reduce to related function calls with related continuations. It is unclear how to make progress in this situation, so we would very much like to rule it out!

In order to make case **EVAL** straightforward to show, while simultaneously ruling out case **CALL** from consideration, we will need to define G_{L_1} and G_{L_2} carefully.

The key idea is as follows: for each pair of function values $(f_1, f_3) \in G$, we come up with a value, v_2 , that (i) uniquely identifies f_1 and f_3 , and that (ii) is not a normal function, but rather a “bad” value that gets stuck when applied to an argument (and hence does not belong to FixVal). The first requirement allows us to ensure $G_{L_1}(\tau) \circ G_{L_2}(\tau) = G(\tau)$, as needed in proving transitivity of the value relation, by

$$\begin{aligned}
\tau \in \text{Ty} &::= \dots \mid \mathbf{n} \\
\text{CTyF} &::= \{(\tau_1 \rightarrow \tau_2) \in \text{CTy}\} \cup \{(\forall \alpha. \tau) \in \text{CTy}\} \cup \{(\text{ref } \tau) \in \text{CTy}\} \cup \{\mathbf{n} \in \text{CTy}\}
\end{aligned}$$

Figure 3.10: Semantic domains for $F^{\mu!}$.

relating $(f_1, v_2) \in G_{L_1}$ and $(v_2, f_3) \in G_{L_2}$. The second requirement, together with Lemma 4, rules out the “bad” case [CALL](#) above.

Formally, since CTy and CVal are countable sets, there exists an injective function $\mathbf{I} \in \text{CTy} \times \text{CTy} \times \text{CVal} \times \text{CVal} \rightarrow \text{CVal}$ (e.g., returning an integer value that uniquely encodes all the arguments). Using this function, one can decompose $G \in \text{VRelF}$ as follows:

$$\begin{aligned}
G_{(1)}(\tau_1 \rightarrow \tau_2) &:= \{(f_1, \mathbf{I}(\tau_1, \tau_2, f_1, f_3)) \mid (f_1, f_3) \in G(\tau_1 \rightarrow \tau_2)\} \\
G_{(2)}(\tau_1 \rightarrow \tau_2) &:= \{(\mathbf{I}(\tau_1, \tau_2, f_1, f_3), f_3) \mid (f_1, f_3) \in G(\tau_1 \rightarrow \tau_2)\}
\end{aligned}$$

Taking G_{L_1} to be $G_{(1)}$ is, however, incorrect because if L_1 relates any values at function type, then the global knowledge will not be closed w.r.t. it (i.e., $G_{(1)} \notin \text{GK}(L_1)$). To address this problem, we simply close $G_{(1)}$ accordingly, i.e., we take G_{L_1} to be the least solution to the fixed-point equation $G_{L_1} = L_1(G_{L_1}) \cup G_{(1)}$ (and similarly for G_{L_2}).

With these definitions, we can show $\overline{G_{L_1}}(\tau) \circ \overline{G_{L_2}}(\tau) = \overline{G}(\tau)$ if $G \in \text{GK}(L)$, and use it in case [EVAL](#) above.

In the problematic case [CALL](#), we make use of the key Lemma 4, which lets us conclude that the functions being called—say, (f_2, f_3) —are *really* external, i.e., related not merely by G_{L_2} but more precisely by $G_{(2)}$. But by construction, this means that f_2 is an integer and thus e'_2 gets stuck, contradicting the prior assumption that e'_1 and e'_2 reduce to values.

3.6 Parametric Bisimulations for $F^{\mu!}$

In this section, we present the full-blown parametric bisimulations for $F^{\mu!}$. This model generalizes the model from the previous section in a superficially very simple way: whereas previously we proved two terms equivalent by exhibiting a consistent local knowledge L , we now do so by exhibiting a consistent *world* W .

3.6.1 Worlds

Worlds are state transition systems (equipped with “public” and “private” transitions, just as described in Section 2.2) that control how the local knowledge of a module and the properties of its local state may evolve over time. Formally (Figure 3.12), a transition system T consists of a (possibly infinite) state space (\mathbf{S}), the private (or full) transition relation (\sqsubseteq) and a smaller public transition relation (\sqsubseteq_{pub}),

$$\begin{aligned} \overline{R}(\tau) &:= R(\tau) \quad \text{if } \tau \in \mathbf{CTyF} \\ &\vdots \\ \overline{R}(\exists\alpha. \tau) &:= \{(\mathbf{pack} \ v_1, \mathbf{pack} \ v_2) \mid \exists\tau'. (v_1, v_2) \in \overline{R}(\tau[\tau'/\alpha])\} \end{aligned}$$

Figure 3.11: Value closure for $F^{\mu!}$ (if $R \in \mathbf{VRelF}$, then $\overline{R} \in \mathbf{VRel}$).

$$\begin{aligned} \text{step}(e) &:= \begin{cases} e' & \text{if } \forall h. \langle h; e \rangle \hookrightarrow \langle h; e' \rangle \\ \perp & \text{otherwise} \end{cases} \\ \text{FixVal} &:= \{v \in \mathbf{CVal} \mid \forall v'. \text{step}(v \ v') \neq \perp\} \\ \text{GenVal} &:= \{v \in \mathbf{CVal} \mid \text{step}(v \ []) \neq \perp\} \\ T \in \mathbf{TrSys} &:= \{(\mathbf{S}, \sqsupseteq, \sqsupseteq_{\text{pub}}) \in \mathbf{Set} \times \mathcal{P}(\mathbf{S} \times \mathbf{S}) \times \mathcal{P}(\mathbf{S} \times \mathbf{S}) \mid \\ &\quad (\forall s, s'. s' \sqsupseteq_{\text{pub}} s \implies s' \sqsupseteq s) \wedge \\ &\quad (\forall s. s \sqsupseteq_{\text{pub}} s) \wedge \\ &\quad (\forall s_1, s_2, s_3. s_3 \sqsupseteq s_2 \wedge s_2 \sqsupseteq s_1 \implies s_3 \sqsupseteq s_1) \wedge \\ &\quad (\forall s_1, s_2, s_3. s_3 \sqsupseteq_{\text{pub}} s_2 \wedge s_2 \sqsupseteq_{\text{pub}} s_1 \implies s_3 \sqsupseteq_{\text{pub}} s_1)\} \\ \mathbf{VRelF}^S &:= \{f \in S \rightarrow \mathbf{VRelF} \rightarrow \mathbf{VRelF} \mid \\ &\quad \forall s, s', R, R'. s' \sqsupseteq s \wedge R' \supseteq R \implies f(s')(R') \supseteq f(s)(R)\} \\ \mathbf{HRel}^S &:= \{f \in S \rightarrow \mathbf{VRelF} \rightarrow \mathbf{HRel} \mid \\ &\quad \forall s, R, R'. R' \supseteq R \implies f(s)(R') \supseteq f(s)(R)\} \\ \mathbf{World}^S &:= \{(T, L, H, N) \in \mathbf{TrSys} \times \mathbf{VRelF}^{S \times T.S} \times \mathbf{HRel}^{S \times T.S} \times \mathcal{P}(\mathbf{TyNam}) \mid \\ &\quad (\forall s, R. \forall \mathbf{n} \notin \mathcal{P}(\mathbf{TyNam}). L(s)(R)(\mathbf{n}) = \emptyset) \wedge \\ &\quad (\forall s, R, \tau, \tau'. L(s)(R)(\tau \rightarrow \tau') \subseteq \text{FixVal} \times \text{FixVal}) \wedge \\ &\quad (\forall s, R, \alpha, \tau. L(s)(R)(\forall\alpha. \tau) \subseteq \text{GenVal} \times \text{GenVal})\} \\ \mathbf{LWorld} &:= \{w \in \mathbf{World}^{W_{\text{ref}}.S} \mid \forall s, R, \tau. w.L(s)(R)(\text{ref } \tau) = \emptyset\} \\ \mathbf{World} &:= \{W \in \mathbf{World}^1\} \end{aligned}$$

Figure 3.12: Definition of worlds.

both preorders. A world then consists of a transition system (T), a mapping from states to local knowledges (L), a mapping from states to heap relations (H), and a set of type names that are used to represent abstract types (N). For now, ignore the distinction between different kinds of worlds in the figure.

As before, the local knowledge (at each state) is parameterized by—and must be monotone in—the global knowledge (R). The same applies to the heap relation, which describes pairs of subheaps that are “owned” by the world. The parameter R here provides a way of referring to the global equivalence on values when establishing invariants on the contents of local heaps; this is especially critical in dealing with higher-order state.

While the local knowledge mapping must be monotone in its state index (w.r.t. the full transition relation), the heap relation mapping need not be. Since a module’s local state is hidden from the environment, there is no necessity to require that heaps related in one state will continue to be related in future states. The freedom this gives us is critical even in very simple examples: imagine a transition system of two states, one in which a particular heap location contains 0 and the other in which the same location contains 1.

We have seen in the previous section how a local knowledge and its closure relate values at λ^μ types. This carries over to the full setting. But how do we deal with the additional types of $F^{\mu!}$, *i.e.*, with universal, existential, and reference types?

3.6.2 Treatment of Universal and Existential Types

Universal types, like function types, are considered flexible. That is, a world’s local knowledge can relate any values at any closed type $\forall\alpha.\tau$, as long as, when instantiated, those values can run for at least one step.

Existential types, like product types, are considered rigid, and thus their interpretation is given by the value closure (Figure 3.11). Note that the witness of related packages must be the *same* type τ' . How, then, do we support reasoning about parametricity?

The key is that the witness type τ' may be an abstract *type name*. We extend the syntax of $F^{\mu!}$ types in our PBS model with type names \mathbf{n} (Figure 3.10), and the local knowledge of a world can pick a subset of these names and interpret them however it wants. To avoid conflicts with other worlds, the choice of names must be recorded in the N component of the world (the local knowledge must not relate anything at other names). Note that N cannot be inferred from the local knowledge as the set of names whose interpretation is nonempty, because the empty relation is a perfectly valid interpretation of a type name.

3.6.3 Treatment of Reference Types

Reference types are considered flexible, but they really are a special case. Intuitively, the collection of all reference types can be seen as a separate module that is used by all other modules. Accordingly, we construct a designated world W_{ref} (explained below)

$$\begin{aligned}
W_{\text{ref}}.S &:= \{s_{\text{ref}} \in \mathcal{P}_{\text{fin}}(\text{CTy} \times \text{Loc} \times \text{Loc}) \mid \\
&\quad \forall(\tau, l_1, l_2) \in s_{\text{ref}}. \forall(\tau', l'_1, l'_2) \in s_{\text{ref}}. \\
&\quad (l_1 = l'_1 \implies \tau = \tau' \wedge l_2 = l'_2) \wedge \\
&\quad (l_2 = l'_2 \implies \tau = \tau' \wedge l_1 = l'_1)\} \\
\\
W_{\text{ref}}.\sqsubseteq &:= \subseteq \\
W_{\text{ref}}.\sqsubseteq_{\text{pub}} &:= \subseteq \\
\\
W_{\text{ref}}.L(s_{\text{ref}})(G)(\text{ref } \tau) &:= \{(l_1, l_2) \mid (\tau, l_1, l_2) \in s_{\text{ref}}\} \\
\\
W_{\text{ref}}.H(s_{\text{ref}})(G) &:= \{(h_1, h_2) \mid \\
&\quad \text{dom}(h_1) = \{l_1 \mid \exists \tau, l_2. (\tau, l_1, l_2) \in s_{\text{ref}}\} \wedge \\
&\quad \text{dom}(h_2) = \{l_2 \mid \exists \tau, l_1. (\tau, l_1, l_2) \in s_{\text{ref}}\} \wedge \\
&\quad \forall(\tau, l_1, l_2) \in s_{\text{ref}}. (h_1(l_1), h_2(l_2)) \in \overline{G}(\tau)\} \\
\\
W_{\text{ref}}.N &:= \emptyset
\end{aligned}$$

Figure 3.13: W_{ref} provides the meaning of reference types.

that interprets $\text{ref } \tau$, and bar ordinary worlds from relating anything at such types. We therefore distinguish between two kinds of worlds: *local worlds* (LWorld) and *full worlds* (World). For conciseness, both are defined in terms of the same underlying structure of *dependent worlds* (DepWorld). We sometimes call W_{ref} a *shared world* or *global world*, because it is not local to the programs we are reasoning about.

A world depending on state space S is a world as described above, except that its local knowledge and heap relation additionally take a state from S as argument (more precisely, their usual state argument is paired with a state from S). Intuitively, S is the state transition system of some other world. Thus, a full world $W \in \text{World}$ is simply a world depending on nothing (a singleton set $1 := \emptyset$). In contrast, a local world $w \in \text{LWorld}$ is a world that depends on—will later be “linked” with— W_{ref} and does not itself relate any values at reference types.

As a matter of notational convenience: if $W \in \text{World}$ and $s \in W.S$, then we will usually just write $W.L(s)$ for $W.L(\emptyset, s)$, and similarly for the H component. (We use the dot notation to project components out of a world.)

3.6.3.1 The World for Reference Types

Figure 3.13 defines $W_{\text{ref}} \in \text{World}$, the world that provides the meaning of reference types. Its states are finite ternary relations (between a type τ and two heap locations l_1, l_2) that are functional in the location arguments. They associate each allocated location on the left with the corresponding one on the right and the type of values stored. The relations are finite because only a finite number of locations can ever be allocated. And, as dictated by the language, they can only grow over time.

Its local knowledge $W_{\text{ref}}.L(s_{\text{ref}})$ relates precisely the locations related by the current state s_{ref} , at the corresponding reference types. The heap relation $W_{\text{ref}}.H(s_{\text{ref}})$ relates heaps that contain exactly the locations related by the current state s_{ref} and that store (globally) related values at those locations. Note the critical use of the global knowledge parameter G in defining $W_{\text{ref}}.H$.

3.6.4 Lifting and Separating Conjunction of Local Worlds

Now, if we have a local world $w \in \text{LWorld}$, then we can link it with W_{ref} , thereby *lifting* it to a full world $w\uparrow \in \text{World}$. This operation is defined in Figure 3.14. The full world's transition system $w\uparrow.T$ is the synchronous product of W_{ref} 's and w 's. Its local knowledge relates values iff they are related by either component's local knowledge, and its heap relation relates heaps iff they can be split into disjoint parts that are related by $W_{\text{ref}}.H$ and $w.H$, respectively. Note how the state s_{ref} of the reference world W_{ref} is passed to $w.L$ and $w.H$ along with the state of w itself.

Similarly, given two local worlds $w_1, w_2 \in \text{LWorld}$ that own disjoint sets of abstract types (*i.e.*, $w_1.N \cap w_2.N = \emptyset$), we can construct their separating conjunction $w_1 \otimes w_2 \in \text{LWorld}$. The definition is also given in Figure 3.14. Note how the same state s_{ref} is passed to the L and H components of both w_1 and w_2 . While this construction is not needed for defining the model itself, it is critical for composing proofs (*e.g.*, when proving soundness). In fact, separating conjunction of local worlds is a generalization of the union operation on local knowledges, which we have seen in Section 3.5.3.

3.6.5 Program Equivalence

With these constructions in hand, we can now describe the definition of program equivalence in Figures 3.15 and 3.16.

We say that two programs are equivalent (\sim) iff there exists a local world w that (1) does not depend on a particular choice of names to represent its abstract types; (2) is *stable*; and (3) when lifted, relates the expressions. Stability means that the local world's heap relation in some sense tolerates “environmental” changes: whenever the shared world W_{ref} is advanced to a future state s'_{ref} , then w should be able to respond to that change by moving to a public future state s' such that any local heaps that were related previously by $w.H$ are still related at s' . This is a technical condition that is required for compositionality (see Lemma 9 later on) but is satisfied trivially in the common case that $w.H$ does not actually depend on its s_{ref} parameter.

A world W (such as $w\uparrow$) relates two expressions (\sim_W) iff (3a) it is *inhabited*; (3b) it is *consistent*; and (3c) the expressions, when closed using pointwise related substitutions, are related by the expression relation (see below). Inhabitation says there exists a state at which $W.H$ relates the empty heaps. Consistency is essentially the same as for λ^μ , but extended straightforwardly to universal types. In all these definitions, the global knowledge G is drawn from $\text{GK}(W)$. As before, this enforces

$$\begin{aligned}
& (-) \times (-) \in \text{TrSys} \rightarrow \text{TrSys} \rightarrow \text{TrSys} \\
& (T_1 \times T_2).S \quad := \quad T_1.S \times T_2.S \\
& (T_1 \times T_2).\sqsubseteq \quad := \quad \{((s_1, s_2), (s'_1, s'_2)) \mid s_1 \sqsubseteq s'_1 \wedge s_2 \sqsubseteq s'_2\} \\
& (T_1 \times T_2).\sqsubseteq_{\text{pub}} \quad := \quad \{((s_1, s_2), (s'_1, s'_2)) \mid s_1 \sqsubseteq_{\text{pub}} s'_1 \wedge s_2 \sqsubseteq_{\text{pub}} s'_2\} \\
\\
& (-) \otimes (-) \in \text{HRel} \rightarrow \text{HRel} \rightarrow \text{HRel} \\
& H_1 \otimes H_2 := \{(h_1 \sqcup h'_1, h_2 \sqcup h'_2) \mid (h_1, h_2) \in H_1 \wedge (h'_1, h'_2) \in H_2 \wedge \\
& \quad h_1 \sqcup h'_1 \neq \perp \wedge h_2 \sqcup h'_2 \neq \perp\} \\
\\
& (-)^\uparrow \in \text{LWorld} \rightarrow \text{World} \\
& w^\uparrow.\mathsf{T} \quad := \quad W_{\text{ref}}.\mathsf{T} \times w.\mathsf{T} \\
& w^\uparrow.\mathsf{L}(s_{\text{ref}}, s)(G) \quad := \quad W_{\text{ref}}.\mathsf{L}(s_{\text{ref}})(G) \cup w.\mathsf{L}(s_{\text{ref}}, s)(G) \\
& w^\uparrow.\mathsf{H}(s_{\text{ref}}, s)(G) \quad := \quad W_{\text{ref}}.\mathsf{H}(s_{\text{ref}})(G) \otimes w.\mathsf{H}(s_{\text{ref}}, s)(G) \\
& w^\uparrow.\mathsf{N} \quad := \quad w.\mathsf{N} \\
\\
& (-) \otimes (-) \in \text{LWorld} \rightarrow \text{LWorld} \rightarrow \text{LWorld} \\
& (w_1 \otimes w_2).\mathsf{T} \quad := \quad w_1.\mathsf{T} \times w_2.\mathsf{T} \\
& (w_1 \otimes w_2).\mathsf{L}(s_{\text{ref}}, s)(G) \quad := \quad w_1.\mathsf{L}(s_{\text{ref}}, s)(G) \cup w_2.\mathsf{L}(s_{\text{ref}}, s)(G) \\
& (w_1 \otimes w_2).\mathsf{H}(s_{\text{ref}}, s)(G) \quad := \quad w_1.\mathsf{H}(s_{\text{ref}}, s)(G) \otimes w_2.\mathsf{H}(s_{\text{ref}}, s)(G) \\
& (w_1 \otimes w_2).\mathsf{N} \quad := \quad w_1.\mathsf{N} \uplus w_2.\mathsf{N}
\end{aligned}$$

Figure 3.14: Lifting and separating conjunction of local worlds.

$$\begin{aligned}
R' &\supseteq_{\text{ref}}^{\mathcal{N}} R := \\
&\quad R' \supseteq R \wedge (\forall \tau. R'(\text{ref } \tau) = R(\text{ref } \tau)) \wedge (\forall \mathbf{n} \in \mathcal{N}. R'(\mathbf{n}) = R(\mathbf{n})) \\
\text{GK}(W) &:= \{G \in W.S \rightarrow \mathbf{VRelF} \mid \\
&\quad (\forall s, s'. s' \supseteq s \implies G(s') \supseteq G(s)) \wedge (\forall s. G(s) \supseteq_{\text{ref}}^{W.N} W.L(s)(G(s)))\} \\
\mathbf{E}_W(G)(s_0)(s)(\tau) &:= \{(e_1, e_2) \mid \\
&\quad \forall (h_1, h_2) \in W.H(s)(G(s)). \forall h_1^F, h_2^F. h_1 \sqcup h_1^F \neq \perp \wedge h_2 \sqcup h_2^F \neq \perp \implies \\
&\quad \quad (\text{DIV}) \quad \langle h_1 \sqcup h_1^F; e_1 \rangle \uparrow \wedge \langle h_2 \sqcup h_2^F; e_2 \rangle \uparrow \\
&\quad \vee (\text{EVAL}) \quad \exists h'_1, h'_2, v_1, v_2, s'. \\
&\quad \quad \langle h_1 \sqcup h_1^F; e_1 \rangle \hookrightarrow^* \langle h'_1 \sqcup h_1^F; v_1 \rangle \wedge \langle h_2 \sqcup h_2^F; e_2 \rangle \hookrightarrow^* \langle h'_2 \sqcup h_2^F; v_2 \rangle \wedge \\
&\quad \quad s' \supseteq s \wedge s' \supseteq_{\text{pub}} s_0 \wedge (h'_1, h'_2) \in W.H(s')(G(s')) \wedge (v_1, v_2) \in \overline{G(s')}(\tau) \\
&\quad \vee (\text{CALL}) \quad \exists h'_1, h'_2, \tau', K_1, K_2, e'_1, e'_2, s'. \\
&\quad \quad (h_1 \sqcup h_1^F, e_1) \hookrightarrow^* (h'_1 \sqcup h_1^F, K_1[e'_1]) \wedge \\
&\quad \quad (h_2 \sqcup h_2^F, e_2) \hookrightarrow^* (h'_2 \sqcup h_2^F, K_2[e'_2]) \wedge \\
&\quad \quad s' \supseteq s \wedge (h'_1, h'_2) \in W.H(s')(G(s')) \wedge (e'_1, e'_2) \in \mathbf{U}(G(s'), G(s'))(\tau') \wedge \\
&\quad \quad \forall s'' \supseteq_{\text{pub}} s'. \forall G' \supseteq G. (K_1, K_2) \in \mathbf{K}_W(G')(s_0)(s'')(\tau', \tau)\} \\
\mathbf{K}_W(G)(s_0)(s)(\tau)(\tau') &:= \{(K_1, K_2) \mid \\
&\quad \forall (v_1, v_2) \in \overline{G(s)}(\tau). (K_1[v_1], K_2[v_2]) \in \mathbf{E}_W(G)(s_0)(s)(\tau')\} \\
\mathbf{U}(R, R')(\tau) &:= \\
&\quad \{(v_1 \ v'_1, v_2 \ v'_2) \mid \exists \tau'. (v_1, v_2) \in R(\tau' \rightarrow \tau) \wedge (v'_1, v'_2) \in \overline{R'}(\tau')\} \cup \\
&\quad \{(v_1 \ \square, v_2 \ \square) \mid \exists \tau_1, \tau_2. \tau = \tau_1[\tau_2/\alpha] \wedge (v_1, v_2) \in R(\forall \alpha. \tau_1)\}
\end{aligned}$$

Figure 3.15: Mutually coinductive definitions of expression equivalence, $\mathbf{E}_W \in \text{GK}(W) \rightarrow W.S \times W.S \rightarrow \mathbf{ERel}$, and continuation equivalence, $\mathbf{K}_W \in \text{GK}(W) \rightarrow W.S \times W.S \rightarrow \mathbf{KRel}$ for $F^{\mu!}$.

$$\begin{aligned}
\text{inhabited}(W) &:= \\
&\quad \forall G \in \mathbf{GK}(W). \exists s_0. (\emptyset, \emptyset) \in W.H(s_0)(G(s_0)) \\
\\
\text{consistent}(W) &:= \\
&\quad \forall G \in \mathbf{GK}(W). \forall s. \forall \tau. \forall (e_1, e_2) \in \mathbf{U}(W.L(s)(G(s)), G(s))(\tau). \\
&\quad (step(e_1), step(e_2)) \in \mathbf{E}_W(G)(s, s)(\tau) \\
\\
\text{stable}(w) &:= \\
&\quad \forall G \in \mathbf{GK}(w\uparrow). \forall s_{\text{ref}}, s. \forall (h_1, h_2) \in w.H(s_{\text{ref}}, s)(G(s_{\text{ref}}, s)). \forall s'_{\text{ref}} \sqsupseteq s_{\text{ref}}. \\
&\quad \forall (h^1_{\text{ref}}, h^2_{\text{ref}}) \in W_{\text{ref}}.H(s'_{\text{ref}})(G(s'_{\text{ref}}, s)). h^1_{\text{ref}} \sqcup h_1 \neq \perp \wedge h^2_{\text{ref}} \sqcup h_2 \neq \perp \implies \\
&\quad \exists s' \sqsupseteq_{\text{pub}} s. (h_1, h_2) \in w.H(s'_{\text{ref}}, s')(G(s'_{\text{ref}}, s')) \\
\\
\text{Env}(\Gamma, R) &:= \\
&\quad \{(\gamma_1, \gamma_2) \in (\text{dom}(\Gamma) \rightarrow \mathbf{CVal})^2 \mid \forall x:\tau \in \Gamma. (\gamma_1 x, \gamma_2 x) \in \overline{R}(\tau)\} \\
\\
\Delta; \Gamma \vdash e_1 \sim_W e_2 : \tau &:= \\
&\quad \text{inhabited}(W) \wedge \text{consistent}(W) \wedge \\
&\quad \forall G \in \mathbf{GK}(W). \forall s. \forall \delta \in \Delta \rightarrow \mathbf{CTy}. \forall (\gamma_1, \gamma_2) \in \text{Env}(\delta\Gamma, G(s)). \\
&\quad (\gamma_1 e_1, \gamma_2 e_2) \in \mathbf{E}_W(G)(s)(s)(\delta\tau) \\
\\
\Delta; \Gamma \vdash e_1 \sim e_2 : \tau &:= \\
&\quad \forall \mathcal{N} \in \mathcal{P}(\mathbf{TyNam}). \mathcal{N} \text{ countably infinite} \implies \\
&\quad \exists w. w.N \subseteq \mathcal{N} \wedge \text{stable}(w) \wedge \Delta; \Gamma \vdash e_1 \sim_{w\uparrow} e_2 : \tau
\end{aligned}$$

Figure 3.16: Definitions of world consistency and program equivalence for $F^{\mu!}$.

that G must *contain* the local knowledge. At reference types, and at abstract type names owned by W , however, $\mathbf{GK}(W)$ also enforces that G must *not extend* the local knowledge. Intuitively, this is because W should completely control the meaning of those types. Furthermore, since G is state-indexed, it must, like the local knowledge of W , be monotone w.r.t. any state changes.

3.6.6 Expression and Continuation Equivalence

The new definitions of \mathbf{E} and \mathbf{K} are also given in Figure 3.15. Notice that they are now defined relative to a world W (as \mathbf{E}_W and \mathbf{K}_W) and that their types have changed to $\mathbf{GK}(W) \rightarrow W.S \rightarrow W.S \rightarrow \mathbf{ERel}$ and $\mathbf{GK}(W) \rightarrow W.S \times W.S \rightarrow \mathbf{KRel}$, respectively: they take both an “initial” state, s_0 , and a “current” state, s , as arguments.

Given a world W , a global knowledge $G \in \mathbf{GK}(W)$, states $s_0, s \in W.S$, and a type τ , we say that two expressions are “locally” equivalent, written $(e_1, e_2) \in \mathbf{E}_W(G)(s_0)(s)(\tau)$, iff, when executed starting in heaps that satisfy the heap relation of W at the current state s , then (as before) one of three cases holds:

Case DIV. Both expressions diverge (run forever).

Case EVAL. Both expressions run successfully to completion, producing related values. In this case, the values need not be related in the current state s , but rather in some future state, $s' \sqsupseteq s$, which, however, must also be a public future state of the initial state of the expression: $s' \sqsupseteq_{\text{pub}} s_0$. Moreover, this future state must be consistent with the resulting heaps: $(h'_1, h'_2) \in W.H(s')(G(s'))$.

Case CALL. Both expressions reduce after some number of steps to some expressions of the form $K_i[e'_i]$, where e'_i are either both applications or both instantiations that are related at some future state $s' \sqsupseteq s$. This state must be consistent with the corresponding heaps. Finally, the continuations, K_1 and K_2 , must be equivalent under any public future state $s'' \sqsupseteq_{\text{pub}} s'$ and any (pointwise) larger global knowledge $G' \sqsupseteq G$.

We restrict s'' to be a public future state of s' rather than an arbitrary future state because the end-to-end effect of a function call (or universal instantiation) is assumed to always be a public transition (recall the discussion of public vs. private transitions in Section 2.2). For this assumption to be sound, in return we will have to ensure that the end-to-end behaviors of equivalent function bodies indeed change the state only into public future states. This is why we thread the s_0 argument through the mutual recursion and check that the final state in case **EVAL** is a public future state of s_0 .

The intuitive reason for quantifying over a larger global knowledge G' in \mathbf{K} is this: At the point when the continuations are run, not only might W ’s state s' have changed to a future state s'' , but also the states of all *other* “modules”, which is reflected by the growth of the global knowledge (also see the next section). This was not needed in Section 3.5 due to the purity of λ^μ .

In all three cases, the definition quantifies over “frame” heaps h_1^F and h_2^F : the execution of e_1 and e_2 should not modify any part of the heap that they do not own according to the heap relation of the current state. This framing aspect of our definition is a semantic version of the frame rule of separation logic and allows us to concentrate the reasoning about the heaps only on the parts of the heaps accessed by the program. (Baking the frame rule into the semantic model is quite common in more recent models of separation logic [10, 88], essentially because it allows one to avoid proving any “safety monotonicity” or “frame” properties of the operational semantics itself.)

3.6.7 Living in a Different World

In Chapter 2, we reviewed at a high level the idea of STSs and how they are used in KLRs. While PBs build very closely on this technique, there are some big differences in the implementation, which are not obvious due to the informal style of Chapter 2.

In KLRs, when proving the equivalence of two programs, we consider them in an arbitrary initial world. As we step through the programs, we may move to a future world, *i.e.*, extend the current world with new state transition systems governing newly allocated pieces of local memory and advance the states of existing state transition systems that we know about. In return, when establishing the relatedness of functions or continuations, we are always forced to consider their behaviour in future worlds. Intuitively, this means two things:

1. Since we extend the world “on the fly”, the STSs that we add are *instance-specific*. They normally refer to the particular locations that were allocated in one run of the programs that we are reasoning about.
2. However, due to the consideration of arbitrary initial and future worlds, the worlds we work with do not only contain the STSs that we explicitly added (even if we only explicitly care about those in our proof). They also contain (i) those belonging to other instances of our programs, and (ii) those belonging to the environment, *i.e.*, governing other modules that our programs were linked with.

Let us contrast this with how STSs function in PBs. Recall that here, in order to prove two programs equivalent, we need to come up with a suitable world, and thus a single STS, *a priori*.

1. This world must hence describe *all* possible instances of our programs, *i.e.*, it must be *module-specific* rather than *instance-specific* (or, *static* rather than *dynamic*). This typically means that the STS that we construct is an n -ary version of what we would construct in a KLR setting (for unbounded n) and thus more complex. The examples in Section 3.8 will make this clearer. The same section also shows how this extra complexity can be reduced.

2. What was the quantification over future worlds in KLRs, now is (i) the quantification over future states, plus (ii) the quantification over larger global knowledges. The latter is necessary because the world only contains our module's constraints, not those of other modules.

3.7 Metatheory

3.7.1 Basics

In a typical proof using PBS, we are given a global knowledge $G \in \mathbf{GK}(W)$ (recall the definition of \sim). Sometimes, however, it is important that we can *construct* such a global knowledge. We now define the least global knowledge for a given world.

Definition 3. For a monotone function $F \in \mathbf{VRelF} \rightarrow \mathbf{VRelF}$ and $R \in \mathbf{VRelF}$, we define $[F]_R^*$ as the least fixed point of the monotone function $F(-) \cup R$:

$$[F]_R^* := \mu X. F(X) \cup R$$

Definition 4 (Least global knowledge). For $W \in \mathbf{World}$ we define $[W] \in \mathbf{W.T.S} \rightarrow \mathbf{VRelF}$ as follows:

$$[W](s) := [W.L(s)]_\emptyset^*$$

Lemma 7 (Least global knowledge). For any W , we have $[W] \in \mathbf{GK}(W)$. Moreover, $[W] \subseteq G$ whenever $G \in \mathbf{GK}(W)$.

Next are three lemmas concerning product worlds. The first one says that a valid global knowledge for a (lifted) product world is also valid for the (lifted) constituent worlds, after fixing the respectively “missing” state. This generalizes Lemma 1.

Lemma 8. If $w_1, w_2 \in \mathbf{LWorld}$ and $G \in \mathbf{GK}((w_1 \otimes w_2)\uparrow)$, then:

1. $\forall s_2 \in w_2.\mathbf{T.S}. G(-, -, s_2) \in \mathbf{GK}(w_1\uparrow)$.
2. $\forall s_1 \in w_1.\mathbf{T.S}. G(-, s_1, -) \in \mathbf{GK}(w_2\uparrow)$.

(Note that the way the product world $w_1 \otimes w_2$ is used as an argument to $(-)\uparrow$ implicitly assumes that the product is defined, *i.e.*, $w_1.\mathbf{N}$ and $w_2.\mathbf{N}$ are disjoint.)

The second lemma relates a product world's **E** and **K** relations to those of its constituent worlds. Its proof is illustrative because (i) it is the only one that relies on stability, and (ii) its structure is representative of other metatheoretical proofs that work by coinduction on **E** and **K**.

Lemma 9. Suppose $w = w_1 \otimes w_2$ and $G \in \mathbf{GK}(w\uparrow)$ and $s_{\text{ref}}^0, s_{\text{ref}} \in W_{\text{ref}}.\mathbf{S}$ and $s_1^0, s_1 \in w_1.\mathbf{S}$ and $s_2^0, s_2 \in w_2.\mathbf{S}$. If $s_2 \sqsupseteq_{\text{pub}} s_2^0$ and $\text{stable}(w_2)$, then:

1. $\mathbf{E}_{w_1\uparrow}(G(-, -, s_2))(s_{\text{ref}}^0, s_1^0)(s_{\text{ref}}, s_1) \subseteq \mathbf{E}_{w\uparrow}(G)(s_{\text{ref}}^0, s_1^0, s_2^0)(s_{\text{ref}}, s_1, s_2)$

$$2. \mathbf{K}_{w_1\uparrow}(G(-, -, s_2))(s_{\text{ref}}^0, s_1^0)(s_{\text{ref}}, s_1) \subseteq \mathbf{K}_{w\uparrow}(G)(s_{\text{ref}}^0, s_1^0, s_2^0)(s_{\text{ref}}, s_1, s_2)$$

And similarly for w_2 if $s_1 \sqsupseteq_{\text{pub}} s_1^0$ and $\text{stable}(w_1)$.

Proof. In order to be able to apply coinduction, we define auxiliaries E and K as follows:

$$\begin{aligned} E(G)(s_{\text{ref}}^0, s_1^0, s_2^0)(s_{\text{ref}}, s_1, s_2)(\tau) &= \{(e_1, e_2) \mid \\ &\quad s_2 \sqsupseteq_{\text{pub}} s_2^0 \wedge (e_1, e_2) \in \mathbf{E}_{w_1\uparrow}(G(-, -, s_2))(s_{\text{ref}}^0, s_1^0)(s_{\text{ref}}, s_1)(\tau)\} \\ K(G)(s_{\text{ref}}^0, s_1^0, s_2^0)(s_{\text{ref}}, s_1, s_2)(\tau')(\tau) &= \{(K_1, K_2) \mid \\ &\quad s_2 \sqsupseteq_{\text{pub}} s_2^0 \wedge (K_1, K_2) \in \mathbf{K}_{w_1\uparrow}(G(-, -, s_2))(s_{\text{ref}}^0, s_1^0)(s_{\text{ref}}, s_1)(\tau')(\tau)\} \end{aligned}$$

Note that it suffices to prove $E \subseteq \mathbf{E}_{w\uparrow}$ and $K \subseteq \mathbf{K}_{w\uparrow}$ (assuming stability of w_2), which we do by coinduction. Concretely, we have to show:

1. $\forall G, s_{\text{ref}}^0, s_1^0, s_2^0, s_{\text{ref}}, s_1, s_2, \tau. \forall (e_1, e_2) \in E(G)(s_{\text{ref}}^0, s_1^0, s_2^0)(s_{\text{ref}}, s_1, s_2)(\tau). \forall (h_1, h_2) \in w\uparrow.\mathbf{H}(s_{\text{ref}}, s_1, s_2)(G(s_{\text{ref}}, s_1, s_2)). \forall h_1^F, h_2^F. \\ h_1 \sqcup h_1^F \neq \perp \wedge h_2 \sqcup h_2^F \neq \perp \implies \\ (\langle h_1 \sqcup h_1^F; e_1 \rangle \uparrow \wedge \langle h_2 \sqcup h_2^F; e_2 \rangle \uparrow) \\ \vee (\exists h'_1, h'_2, v_1, v_2, s'. \\ \langle h_1 \sqcup h_1^F; e_1 \rangle \hookrightarrow^* \langle h'_1 \sqcup h_1^F; v_1 \rangle \wedge \langle h_2 \sqcup h_2^F; e_2 \rangle \hookrightarrow^* \langle h'_2 \sqcup h_2^F; v_2 \rangle \wedge \\ s' \sqsupseteq (s_{\text{ref}}, s_1, s_2) \wedge s' \sqsupseteq_{\text{pub}} (s_{\text{ref}}^0, s_1^0, s_2^0) \wedge \\ (h'_1, h'_2) \in w\uparrow.\mathbf{H}(s')(G(s')) \wedge (v_1, v_2) \in \overline{G(s')}(\tau)) \\ \vee (\exists h'_1, h'_2, \tau', K_1, K_2, e'_1, e'_2, s'. \\ (h_1 \sqcup h_1^F, e_1) \hookrightarrow^* (h'_1 \sqcup h_1^F, K_1[e'_1]) \wedge (h_2 \sqcup h_2^F, e_2) \hookrightarrow^* (h'_2 \sqcup h_2^F, K_2[e'_2]) \wedge \\ s' \sqsupseteq (s_{\text{ref}}, s_1, s_2) \wedge (h'_1, h'_2) \in w\uparrow.\mathbf{H}(s')(G(s')) \wedge (e'_1, e'_2) \in \mathbf{U}(G(s'), G(s'))(\tau') \wedge \\ (K_1, K_2) \in K(G)(s_{\text{ref}}^0, s_1^0, s_2^0)(s')(\tau', \tau))$
2. $\forall G, s_{\text{ref}}^0, s_1^0, s_2^0, s_{\text{ref}}, s_1, s_2, \tau', \tau. \\ \forall (K_1, K_2) \in K(G)(s_{\text{ref}}^0, s_1^0, s_2^0)(s_{\text{ref}}, s_1, s_2)(\tau', \tau). \forall (v_1, v_2) \in \overline{G(s_{\text{ref}}, s_1, s_2)}(\tau'). \\ (K_1[v_1], K_2[v_2]) \in E(G)(s_{\text{ref}}^0, s_1^0, s_2^0)(s_{\text{ref}}, s_1, s_2)(\tau)$

Part (2) is straightforward. Let us look at part (1). We are given heaps h_1, h_2 related by the lifted product world $w\uparrow$. We know that these are composed as follows:

$$\begin{aligned} h_1 &= h'_1 \uplus h''_1 \\ h_2 &= h'_2 \uplus h''_2 \\ (h_1^a, h_2^a) &\in w_1\uparrow.\mathbf{H}(s_{\text{ref}}, s_1)(G(s_{\text{ref}}, s_1, s_2)) \quad (*) \\ (h_1^b, h_2^b) &\in w_2.\mathbf{H}(s_{\text{ref}}, s_2)(G(s_{\text{ref}}, s_1, s_2)) \end{aligned}$$

Naturally, we use $(*)$ to instantiate the relatedness of (e_1, e_2) by $\mathbf{E}_{w_1\uparrow}$ (in the definition of E). As frame heaps, we choose $h_1^b \uplus h_1^F$ and $h_2^b \uplus h_2^F$, so that the programs cannot violate the constraints imposed by w_2 . We do a case split on the disjunct according to which e_1 and e_2 are related and in each case choose to show the corresponding disjunct of the goal.

Stability is required for the cases **EVAL** and **CALL**. For instance, suppose e_1 and e_2 are related because they both terminate with related values in state $(s'_{\text{ref}}, s'_1) \sqsupseteq (s_{\text{ref}}, s_1)$. We know that the frame heaps have been preserved and that there are $h_1^{a'}, h_2^{a'}$ (disjoint from the frames) and future states s'_{ref}, s'_1 with $s'_{\text{ref}} \sqsupseteq s_{\text{ref}}, s'_{\text{ref}} \sqsupseteq_{\text{pub}} s_{\text{ref}}^0, s'_1 \sqsupseteq s_1, s_1 \sqsupseteq_{\text{pub}} s_1^0$, and

$$(h_1^{a'}, h_2^{a'}) \in w_1 \uparrow. \mathbf{H}(s'_{\text{ref}}, s'_1)(G(s'_{\text{ref}}, s'_1, s_2)).$$

To show the goal, we must now find s'_2 with $s'_2 \sqsupseteq s_2$ and $s'_2 \sqsupseteq_{\text{pub}} s_2^0$ such that

$$(h_1^{a'} \uplus h_1^b, h_2^{a'} \uplus h_2^b) \in w \uparrow. \mathbf{H}(s'_{\text{ref}}, s'_1, s'_2)(G(s'_{\text{ref}}, s'_1, s'_2)),$$

i.e., (using monotonicity of \mathbf{H} and G) such that:

$$(h_1^b, h_2^b) \in w_2. \mathbf{H}(s'_{\text{ref}}, s'_2)(G(s'_{\text{ref}}, s'_1, s'_2))$$

Relying on $s_2 \sqsupseteq_{\text{pub}} s_2^0$ (enforced by the definition of E), stability of w_2 yields exactly what we need. \square

The third property of product worlds generalizes Lemma 2.

Lemma 10. If $w = w_1 \otimes w_2$ with $\text{stable}(w_1)$ and $\text{stable}(w_2)$, then:

1. $\text{stable}(w)$
2. If $\text{inhabited}(w_1 \uparrow)$ and $\text{inhabited}(w_2 \uparrow)$, then $\text{inhabited}(w \uparrow)$.
3. If $\text{consistent}(w_1 \uparrow)$ and $\text{consistent}(w_2 \uparrow)$, then $\text{consistent}(w \uparrow)$.

The next lemma talks about composing continuations with other continuations or expressions and generalizes Lemma 3.

Lemma 11. Given $G, s_0, s'_0, s, \tau, \tau', K_1, K_2$ such that

$$\forall G' \sqsupseteq G. \forall s'. s' \sqsupseteq s \wedge s' \sqsupseteq_{\text{pub}} s'_0 \implies (K_1, K_2) \in \mathbf{K}_W(G')(s_0)(s')(\tau')(\tau),$$

we have:

1. $(e_1, e_2) \in \mathbf{E}_W(G)(s'_0, s)(\tau') \implies (K_1[e_1], K_2[e_2]) \in \mathbf{E}_W(G)(s_0, s)(\tau)$
2. $(K'_1, K'_2) \in \mathbf{K}_W(G)(s'_0, s)(\tau'')(\tau') \implies (K_1[K'_1], K_2[K'_2]) \in \mathbf{K}_W(G)(s_0, s)(\tau'')(\tau)$

Let us now come to the external call lemma, the generalization of Lemma 4. First we parameterize \mathbf{E} over the relation by which functions are related via **CALL**.

Definition 5.

$$\begin{aligned}
\mathbf{E}_W^{\mathcal{R}}(G)(s_0)(s)(\tau) := & \{ (e_1, e_2) \mid \\
& \forall (h_1, h_2) \in W.H(s)(G(s)). \forall h_1^F, h_2^F. h_1 \sqcup h_1^F \neq \perp \wedge h_2 \sqcup h_2^F \neq \perp \implies \\
& \quad (\langle h_1 \sqcup h_1^F; e_1 \rangle \uparrow \wedge \langle h_2 \sqcup h_2^F; e_2 \rangle \uparrow) \\
& \vee (\exists h'_1, h'_2, v_1, v_2, s'. \\
& \quad \langle h_1 \sqcup h_1^F; e_1 \rangle \hookrightarrow^* \langle h'_1 \sqcup h_1^F; v_1 \rangle \wedge \langle h_2 \sqcup h_2^F; e_2 \rangle \hookrightarrow^* \langle h'_2 \sqcup h_2^F; v_2 \rangle \wedge \\
& \quad s' \sqsupseteq s \wedge s' \sqsupseteq_{\text{pub}} s_0 \wedge (h'_1, h'_2) \in W.H(s')(G(s')) \wedge (v_1, v_2) \in \overline{G(s')}(\tau)) \\
& \vee (\exists h'_1, h'_2, \tau', K_1, K_2, e'_1, e'_2, s'. \\
& \quad (h_1 \sqcup h_1^F, e_1) \hookrightarrow^* (h'_1 \sqcup h_1^F, K_1[e'_1]) \wedge (h_2 \sqcup h_2^F, e_2) \hookrightarrow^* (h'_2 \sqcup h_2^F, K_2[e'_2]) \wedge \\
& \quad s' \sqsupseteq s \wedge (h'_1, h'_2) \in W.H(s')(G(s')) \wedge (e'_1, e'_2) \in \mathbf{U}(\mathcal{R}(s'), G(s'))(\tau') \wedge \\
& \quad (K_1, K_2) \in \mathbf{K}_W(G)(s_0)(s')(\tau')(\tau)) \\
& \}
\end{aligned}$$

Note how we use $\mathcal{R}(s')$ rather than $G(s')$ as first argument to \mathbf{U} . Other than that, $\mathbf{E}^{\mathcal{R}}$ is identical to \mathbf{E} . It is actually defined in terms of \mathbf{E} , and not recursive itself.

Lemma 12 (External call). Suppose $\text{consistent}(W)$, $G \in \text{GK}(W)$ and $\mathcal{R} \in W.T.S \rightarrow \text{VRelF}$. If \mathcal{R} is the “external” part of G , i.e.,

$$\forall s. G(s) = W.L(s)(G(s)) \cup \mathcal{R}(s),$$

then $\mathbf{E}_W(G) = \mathbf{E}_W^{\mathcal{R}}(G)$.

Proof. The \sqsupseteq -direction is trivial due to monotonicity of \mathbf{U} . For the other direction, the high-level proof idea is the same as in the simple setting: we eliminate all uses of the **CALL** disjunct of \mathbf{E} where the invoked functions are related by W by relying on W ’s consistency.

More formally, we prove by induction on n that

$$\forall n. \mathbf{E}_W(G) \subseteq X_n(G),$$

where X_n is defined like $\mathbf{E}_W^{\mathcal{R}}$ except that divergence (\uparrow) in the case **DIV** is replaced with \uparrow^n , meaning that the two programs can take at least n steps. Notice how this universally quantified property is therefore equivalent to $\mathbf{E}_W(G) \subseteq \mathbf{E}_W^{\mathcal{R}}(G)$.

In the base case ($n = 0$), we are done immediately by choosing the **DIV** disjunct of X_0 (clearly any program can take at least 0 steps).

The inductive case is the interesting one. We assume $\mathbf{E}_W(G) \subseteq X_{n-1}(G)$ and must show $\mathbf{E}_W(G) \subseteq X_n(G)$. So suppose $(e_1, e_2) \in \mathbf{E}_W(G)(s^0)(s)(\tau)$. We are now given related heaps and frame heaps as dictated by the definition of X_n and use these to instantiate $(e_1, e_2) \in \mathbf{E}_W(G)(s^0)(s)(\tau)$. This yields three cases, one per disjunct. In case **DIV**, we are done because \uparrow implies \uparrow^n . In case **EVAL**, we are done trivially. In case **CALL** we have $e_1 \hookrightarrow^* K_1[e'_1]$ and $e_2 \hookrightarrow^* K_2[e'_2]$ where e'_1 and e'_2 are calls to G -related functions. Now, if these functions are related by \mathcal{R} , the external part of G , then we are done trivially again. Otherwise they are related by the local

knowledge $W.L$. Since the world is consistent, we can exploit this fact to learn that the beta-reduced calls $step(e'_1)$ and $step(e'_2)$ are related by $\mathbf{E}_W(G)$. By Lemma 11, $K_1[step(e'_1)]$ and $K_2[step(e'_2)]$ are related by $\mathbf{E}_W(G)$ as well, and hence by $X_{n-1}(G)$. Since the initial programs take at least one step each to reach these intermediate forms (*i.e.*, $e_i \hookrightarrow^+ K_i[step(e'_i)]$), we can conclude that e_1 and e_2 are in fact related by $X_n(G)$. \square

3.7.2 Symmetry

We now show that PBs are symmetric (see the final Theorem 3). The proof is not complicated, but it does require a few definitions and lemmas.

Definition 6. Given $R \in \mathbf{VRel}$ (or \mathbf{VRelF}), we define its inverse $R^{-1} \in \mathbf{VRel}$ (or \mathbf{VRelF}) pointwise:

$$R^{-1} := \lambda\tau. R(\tau)^{-1}$$

Lemma 13. $\forall R \in \mathbf{VRelF}. (\overline{R})^{-1} = \overline{R^{-1}}$

Proof. By an easy induction. \square

Lemma 14. $\forall R, R' \in \mathbf{VRel}. \mathbf{U}(R^{-1}, R'^{-1}) = (\mathbf{U}(R, R'))^{-1}$

Definition 7. Given $w \in \mathbf{LWorld}$, we define $w^{-1} \in \mathbf{LWorld}$ as follows:

$$\begin{aligned} w^{-1}.T &:= w.T \\ w^{-1}.L(s_{\text{ref}}, s)(R) &:= (w.L(s_{\text{ref}}^{-1}, s)(R^{-1}))^{-1} \\ w^{-1}.H(s_{\text{ref}}, s)(R) &:= (w.H(s_{\text{ref}}^{-1}, s)(R^{-1}))^{-1} \\ w^{-1}.N &:= w.N \end{aligned}$$

Here, s_{ref}^{-1} is short for $\{(\tau, l_2, l_1) \mid (\tau, l_1, l_2) \in s_{\text{ref}}\}$.

Lemma 15. $\forall w, s_{\text{ref}}, s, R. w^{-1}\uparrow.H(s_{\text{ref}}, s)(R) = (w\uparrow.H(s_{\text{ref}}^{-1}, s)(R^{-1}))^{-1}$

Proof.

$$\begin{aligned} &w^{-1}\uparrow.H(s_{\text{ref}}, s)(R) \\ = &W_{\text{ref}}.H(s_{\text{ref}})(R) \otimes w^{-1}.H(s_{\text{ref}}, s)(R) \\ = &\left((W_{\text{ref}}.H(s_{\text{ref}})(R))^{-1} \otimes (w^{-1}.H(s_{\text{ref}}, s)(R))^{-1} \right)^{-1} \\ = &(W_{\text{ref}}.H(s_{\text{ref}}^{-1})(R^{-1}) \otimes w.H(s_{\text{ref}}^{-1}, s)(R^{-1}))^{-1} \\ = &(w\uparrow.H(s_{\text{ref}}^{-1}, s)(R^{-1}))^{-1} \end{aligned}$$

\square

Lemma 16. $\forall w, s_{\text{ref}}, s, R. w^{-1}\uparrow.L(s_{\text{ref}}, s)(R) = (w\uparrow.L(s_{\text{ref}}^{-1}, s)(R^{-1}))^{-1}$

Proof. Analogous to Lemma 15. \square

Definition 8. Given $G \in \mathbf{GK}(w^{-1}\uparrow)$, we define its inverse G^{-1} as follows:

$$G^{-1}(s_{\text{ref}}, s) := G(s_{\text{ref}}^{-1}, s)^{-1}$$

Lemma 17. If $G \in \mathbf{GK}(w^{-1}\uparrow)$, then $G^{-1} \in \mathbf{GK}(w\uparrow)$.

Proof. Monotonicity of G^{-1} follows immediately from monotonicity of G . It remains to show

$$G^{-1}(s_{\text{ref}}, s) \supseteq_{\text{ref}}^{w\uparrow.\mathbf{N}} w\uparrow.\mathbf{L}(s_{\text{ref}}, s)(G^{-1}(s_{\text{ref}}, s))$$

for any s_{ref}, s . This is easy using Lemma 16:

$$\begin{aligned} & G^{-1}(s_{\text{ref}}, s) \\ = & G(s_{\text{ref}}^{-1}, s)^{-1} \\ \supseteq_{\text{ref}}^{w\uparrow.\mathbf{N}} & (w^{-1}\uparrow.\mathbf{L}(s_{\text{ref}}^{-1}, s)(G(s_{\text{ref}}^{-1}, s)))^{-1} \\ = & w\uparrow.\mathbf{L}(s_{\text{ref}}, s)(G(s_{\text{ref}}^{-1}, s)^{-1}) \\ = & w\uparrow.\mathbf{L}(s_{\text{ref}}, s)(G^{-1}(s_{\text{ref}}, s)) \end{aligned}$$

□

Lemma 18. If $\text{stable}(w)$, then $\text{stable}(w^{-1})$.

Lemma 19. If $\text{inhabited}(w\uparrow)$, then $\text{inhabited}(w^{-1}\uparrow)$.

Lemma 20. If $G \in \mathbf{GK}(w^{-1}\uparrow)$, then:

$$(\mathbf{E}_{w\uparrow}(G^{-1})(s_{\text{ref}0}^{-1}, s_0)(s_{\text{ref}}^{-1}, s))^{-1} \subseteq \mathbf{E}_{w^{-1}\uparrow}(G)(s_{\text{ref}0}, s_0)(s_{\text{ref}}, s)$$

Proof. By a straightforward coinduction.

□

Lemma 21. If $\text{consistent}(w\uparrow)$, then $\text{consistent}(w^{-1}\uparrow)$.

Theorem 3.

$$\frac{\Delta; \Gamma \vdash e_1 \sim e_2 : \sigma}{\Delta; \Gamma \vdash e_2 \sim e_1 : \sigma}$$

Proof. Given $\mathcal{N} \in \mathbf{TyNam}$, we get a stable w with $\Delta; \Gamma \vdash e_1 \sim_{w\uparrow} e_2 : \sigma$ as well as $w.\mathbf{N} \subseteq \mathcal{N}$. Using $w^{-1}.\mathbf{N} = w.\mathbf{N}$ and Lemma 18, it remains to show:

$$\Delta; \Gamma \vdash e_2 \sim_{w^{-1}\uparrow} e_1 : \sigma$$

With Lemmas 19 and 21, this in turn reduces to showing:

$$(\gamma_1 e_2, \gamma_2 e_1) \in \mathbf{E}_{w^{-1}\uparrow}(G)(s_{\text{ref}}, s)(s_{\text{ref}}, s)(\delta\sigma)$$

for any $s_{\text{ref}}, s, \delta$ and $G \in \mathbf{GK}(w^{-1}\uparrow)$ and $(\gamma_1, \gamma_2) \in \mathbf{Env}(\delta\Gamma, G(s_{\text{ref}}, s))$.

Note that $(\gamma_2, \gamma_1) \in \mathbf{Env}(\delta\Gamma, G(s_{\text{ref}}, s)^{-1}) = \mathbf{Env}(\delta\Gamma, G^{-1}(s_{\text{ref}}^{-1}, s))$ with the help of Lemma 13 and the definition of G^{-1} . Using Lemma 17, $\Delta; \Gamma \vdash e_1 \sim_{w\uparrow} e_2 : \sigma$ thus yields:

$$(\gamma_2 e_1, \gamma_1 e_2) \in \mathbf{E}_{w\uparrow}(G^{-1})(s_{\text{ref}}^{-1}, s)(s_{\text{ref}}^{-1}, s)(\delta\sigma)$$

We are done by Lemma 20.

□

3.7.3 Compatibilities

Compatibilities are very useful lemmas stating that the language's constructs preserve our equivalence. For each typing rule we have a corresponding compatibility. Figure 3.17 lists them all. We have seen two of them—and their proof sketches—earlier in the simple setting (FIX and APP in Figure 3.9).

These proofs are somewhat tedious and, in general, similar to their counterparts in a KLR setting. One of the key differences lies in the proofs concerning functions, which we already sketched earlier. Here we show these proofs in detail, generalized to the full $F^{\mu l}$ setting of this section.

Lemma 22 (Compatibility: App).

$$\frac{\Delta; \Gamma \vdash p_1 \sim p'_1 : \sigma_1 \rightarrow \sigma_2 \quad \Delta; \Gamma \vdash p_2 \sim p'_2 : \sigma_1}{\Delta; \Gamma \vdash p_1 p_2 \sim p'_1 p'_2 : \sigma_2}$$

Proof. Given a countably infinite set $\mathcal{N} \in \mathcal{P}(\text{TyNam})$, we can split it into two countably infinite parts: $\mathcal{N} = \mathcal{N}_1 \uplus \mathcal{N}_2$. Using these to instantiate the first and second premise, respectively, we obtain two stable local worlds w_1 and w_2 whose product $w := w_1 \otimes w_2$ is defined and satisfies $w.N \subseteq \mathcal{N}$. We use Lemma 10 to infer stability of w and consistency and inhabitation of $w \uparrow$. It thus remains to show

$$(\gamma(p_1 p_2), \gamma'(p'_1 p'_2)) \in \mathbf{E}_{w \uparrow}(G)(s_{\text{ref}}, s_1, s_2)(s_{\text{ref}}, s_1, s_2)(\delta\sigma_2)$$

for substitutions $(\gamma, \gamma') \in \text{Env}(\Gamma, G(s_{\text{ref}}, s_1, s_2))$. The first instantiated premise lets us assume

$$(\gamma p_1, \gamma' p'_1) \in \mathbf{E}_{w_1 \uparrow}(G(-, -, s_2))(s_{\text{ref}}, s_1)(s_{\text{ref}}, s_1)(\delta(\sigma_1 \rightarrow \sigma_2)),$$

which Lemma 9 turns into

$$(\gamma p_1, \gamma' p'_1) \in \mathbf{E}_{w \uparrow}(G)(s_{\text{ref}}, s_1, s_2)(s_{\text{ref}}, s_1, s_2)(\delta(\sigma_1 \rightarrow \sigma_2)).$$

By Lemma 11 it therefore suffices to show

$$(\bullet \gamma p_2, \bullet \gamma' p'_2) \in \mathbf{K}_{w \uparrow}(G')(s_{\text{ref}}, s_1, s_2)(s'_{\text{ref}}, s'_1, s'_2)(\delta\sigma_1 \rightarrow \delta\sigma_2)(\delta\sigma_2)$$

for any $G' \supseteq G$ and any $(s'_{\text{ref}}, s'_1, s'_2) \sqsupseteq_{\text{pub}} (s_{\text{ref}}, s_1, s_2)$. Given related values $(v_1, v'_1) \in \overline{G'(s'_{\text{ref}}, s'_1, s'_2)}(\delta\sigma_1 \rightarrow \delta\sigma_2)$, we therefore must show:

$$(v_1 \gamma p_2, v'_1 \gamma' p'_2) \in \mathbf{E}_{w \uparrow}(G')(s_{\text{ref}}, s_1, s_2)(s'_{\text{ref}}, s'_1, s'_2)(\delta\sigma_2)$$

Now, the second instantiated premise lets us assume

$$(\gamma p_2, \gamma' p'_2) \in \mathbf{E}_{w_2 \uparrow}(G'(-, -, s'_2))(s_{\text{ref}}, s_1)(s'_{\text{ref}}, s'_1)(\delta\sigma_1),$$

which Lemma 9 turns into

$$(\gamma p_2, \gamma' p'_2) \in \mathbf{E}_{w \uparrow}(G')(s_{\text{ref}}, s_1, s_2)(s'_{\text{ref}}, s'_1, s'_2)(\delta\sigma_1).$$

$\frac{\Delta \vdash \Gamma \quad x:\sigma \in \Gamma}{\Delta; \Gamma \vdash x \sim x : \sigma} \quad \frac{\Delta \vdash \Gamma}{\Delta; \Gamma \vdash n \sim n : \text{nat}}$
$\frac{\Delta; \Gamma \vdash p_0 \sim p'_0 : \text{nat} \quad \Delta; \Gamma \vdash p_1 \sim p'_1 : \sigma \quad \Delta; \Gamma \vdash p_2 \sim p'_2 : \sigma}{\Delta; \Gamma \vdash \text{ifz } p_0 \text{ then } p_1 \text{ else } p_2 \sim \text{ifz } p'_0 \text{ then } p'_1 \text{ else } p'_2 : \sigma}$
$\frac{\Delta; \Gamma \vdash p_1 \sim p'_1 : \sigma_1 \quad \Delta; \Gamma \vdash p_2 \sim p'_2 : \sigma_2}{\Delta; \Gamma \vdash \langle p_1, p_2 \rangle \sim \langle p'_1, p'_2 \rangle : \sigma_1 \times \sigma_2}$
$\frac{\Delta; \Gamma \vdash p \sim p' : \sigma_1 \times \sigma_2}{\Delta; \Gamma \vdash p.1 \sim p'.1 : \sigma_1} \quad \frac{\Delta; \Gamma \vdash p \sim p' : \sigma_1 \times \sigma_2}{\Delta; \Gamma \vdash p.2 \sim p'.2 : \sigma_2}$
$\frac{\Delta; \Gamma \vdash p \sim p' : \sigma_1 \quad \Delta; \Gamma \vdash \sigma_2}{\Delta; \Gamma \vdash \text{inl}_{\sigma_2} p \sim \text{inl}_{\sigma_2} p' : \sigma_1 + \sigma_2} \quad \frac{\Delta; \Gamma \vdash \sigma_1 \quad \Delta; \Gamma \vdash p \sim p' : \sigma_2}{\Delta; \Gamma \vdash \text{inr}_{\sigma_1} p \sim \text{inr}_{\sigma_2} p' : \sigma_1 + \sigma_2}$
$\frac{\Delta; \Gamma \vdash p_0 \sim p'_0 : \sigma_1 + \sigma_2 \quad \Delta; \Gamma, x:\sigma_1 \vdash p_1 \sim p'_1 : \sigma \quad \Delta; \Gamma, x:\sigma_2 \vdash p_2 \sim p'_2 : \sigma}{\Delta; \Gamma \vdash \text{case } p_0 (x. p_1) (x. p_2) \sim \text{case } p'_0 (x. p'_1) (x. p'_2) : \sigma}$
$\frac{\Delta; \Gamma, f:(\sigma_1 \rightarrow \sigma_2), x:\sigma_1 \vdash p \sim p' : \sigma_2}{\Delta; \Gamma \vdash \text{fix } f(x:\sigma_1):\sigma_2. p \sim \text{fix } f(x:\sigma_1):\sigma_2. p' : \sigma_1 \rightarrow \sigma_2}$
$\frac{\Delta; \Gamma \vdash p_1 \sim p'_1 : \sigma_1 \rightarrow \sigma_2 \quad \Delta; \Gamma \vdash p_2 \sim p'_2 : \sigma_1}{\Delta; \Gamma \vdash p_1 p_2 \sim p'_1 p'_2 : \sigma_2}$
$\frac{\Delta, \alpha; \Gamma \vdash p \sim p' : \sigma}{\Delta; \Gamma \vdash \Lambda \alpha. p \sim \Lambda \alpha. p' : \forall \alpha. \sigma} \quad \frac{\Delta; \Gamma \vdash p \sim p' : \forall \alpha. \sigma_1 \quad \Delta; \Gamma \vdash \sigma_2}{\Delta; \Gamma \vdash p [\sigma_2] \sim p' [\sigma_2] : \sigma_1 [\sigma_2 / \alpha]}$
$\frac{\Delta; \Gamma \vdash \sigma_1 \quad \Delta; \Gamma \vdash p \sim p' : \sigma_2 [\sigma_1 / \alpha]}{\Delta; \Gamma \vdash \text{pack } \langle \sigma_1, p \rangle \text{ as } \exists \alpha. \sigma_2 \sim \text{pack } \langle \sigma_1, p' \rangle \text{ as } \exists \alpha. \sigma_2 : \exists \alpha. \sigma_2}$
$\frac{\Delta; \Gamma \vdash p_1 \sim p'_1 : \exists \alpha. \sigma_1 \quad \Delta, \alpha; \Gamma, x:\sigma_1 \vdash p_2 \sim p'_2 : \sigma_2 \quad \Delta; \Gamma \vdash \sigma_2}{\Delta; \Gamma \vdash \text{unpack } p_1 \text{ as } \langle \alpha, x \rangle \text{ in } p_2 \sim \text{unpack } p'_1 \text{ as } \langle \alpha, x \rangle \text{ in } p'_2 : \sigma_2}$
$\frac{\Delta; \Gamma \vdash p \sim p' : \sigma [\mu \alpha. \sigma / \alpha]}{\Delta; \Gamma \vdash \text{roll}_{\mu \alpha. \sigma} p \sim \text{roll}_{\mu \alpha. \sigma} p' : \mu \alpha. \sigma} \quad \frac{\Delta; \Gamma \vdash p' : \mu \alpha. \sigma}{\Delta; \Gamma \vdash \text{unroll } p \sim \text{unroll } p' : \sigma [\mu \alpha. \sigma / \alpha]}$
$\frac{\Delta; \Gamma \vdash p \sim p' : \sigma}{\Delta; \Gamma \vdash \text{ref } p \sim \text{ref } p' : \text{ref } \sigma} \quad \frac{\Delta; \Gamma \vdash p_1 \sim p'_1 : \text{ref } \sigma \quad \Delta; \Gamma \vdash p_2 \sim p'_2 : \sigma}{\Delta; \Gamma \vdash p_1 := p_2 \sim p'_1 := p'_2 : \text{unit}}$
$\frac{\Delta; \Gamma \vdash p \sim p' : \text{ref } \sigma}{\Delta; \Gamma \vdash !p \sim !p' : \sigma} \quad \frac{\Delta; \Gamma \vdash p_1 \sim p'_1 : \text{ref } \sigma \quad \Delta; \Gamma \vdash p_2 \sim p'_2 : \text{ref } \sigma}{\Delta; \Gamma \vdash p_1 == p_2 \sim p'_1 == p'_2 : \text{nat}}$

Figure 3.17: Compatibilities

Again by Lemma 11 it suffices to show

$$(v_1 \bullet, v'_1 \bullet) \in \mathbf{K}_{w\uparrow}(G'')(s_{\text{ref}}, s_1, s_2)(s''_{\text{ref}}, s''_1, s''_2)(\delta\sigma_1)(\delta\sigma_2)$$

for any $G'' \supseteq G'$ and any $(s''_{\text{ref}}, s''_1, s''_2) \sqsupseteq (s'_{\text{ref}}, s'_1, s'_2)$ with $(s''_{\text{ref}}, s''_1, s''_2) \sqsupseteq_{\text{pub}} (s_{\text{ref}}, s_1, s_2)$. Given related values $(v_2, v'_2) \in \overline{G''(s''_{\text{ref}}, s''_1, s''_2)}(\delta\sigma_1)$, we therefore must show

$$(v_1 \ v_2, v'_1 \ v'_2) \in \mathbf{E}_{w\uparrow}(G'')(s_{\text{ref}}, s_1, s_2)(s''_{\text{ref}}, s''_1, s''_2)(\delta\sigma_2),$$

which is the interesting part of the proof. (In fact, the reasoning up to here is not specific to function application but common to many other program forms involving subterms. In our formal Coq development, some of the boilerplate is moved into shared lemmas.)

To show this last goal we appeal to the **CALL** disjunct of **E**. Accordingly, it suffices to show three things:

- $(v_1, v'_1) \in \overline{G''(s''_{\text{ref}}, s''_1, s''_2)}(\delta\sigma_1 \rightarrow \delta\sigma_2)$
- $(v_2, v'_2) \in \overline{G''(s''_{\text{ref}}, s''_1, s''_2)}(\delta\sigma_1)$
- $(\bullet, \bullet) \in \mathbf{K}_{w\uparrow}(G''')(s_{\text{ref}}, s_1, s_2)(s''')(\delta\sigma_2)(\delta\sigma_2)$ for any $G''' \supseteq G''$ and any $s''' \sqsupseteq (s''_{\text{ref}}, s''_1, s''_2)$ with $s''' \sqsupseteq_{\text{pub}} (s_{\text{ref}}, s_1, s_2)$

The second was an assumption, the first follows from the assumptions by monotonicity, and the third, which talks about empty continuations, is easy to verify. \square

Lemma 23 (Compatibility: Fix).

$$\frac{\Delta; \Gamma, f:\sigma_1 \rightarrow \sigma_2, x:\sigma_1 \vdash e_1 \sim e_2 : \sigma_2}{\Delta; \Gamma \vdash \text{fix } f(x). e_1 \sim \text{fix } f(x). e_2 : \sigma_1 \rightarrow \sigma_2}$$

Proof. Given $\mathcal{N} \in \mathcal{P}(\text{TyNam})$, the premise yields a stable w_1 such that $w_1.\mathbf{N} \subseteq \mathcal{N}$ and $\Delta; \Gamma, f:\sigma_1 \rightarrow \sigma_2, x:\sigma_1 \vdash e_1 \sim_{w_1\uparrow} e_2 : \sigma_2$. We choose the local world in which we prove the two functions related to be $w := w_1 \otimes w_2$, where w_2 is defined as the following trivially stable single-state world:

$$\begin{aligned} w_2.\mathbf{T.S} &:= \{1\} \\ w_2.\mathbf{L}(\cdot)(\cdot)(R)(\tau) &:= \{(\gamma_1 \text{fix } f(x). e_1, \gamma_2 \text{fix } f(x). e_2) \mid \exists \delta \in \Delta \rightarrow \mathbf{CTy}. \\ &\quad \tau = \delta\sigma_1 \rightarrow \delta\sigma_2 \wedge (\gamma_1, \gamma_2) \in \mathbf{Env}(\delta\Gamma, R)\} \\ w_2.\mathbf{H}(\cdot)(\cdot)(\cdot) &:= \{(\emptyset, \emptyset)\} \\ w_2.\mathbf{N} &:= \emptyset \end{aligned}$$

With Lemma 10 and $w.\mathbf{N} = (w_1 \otimes w_2).\mathbf{N} = w_1.\mathbf{N} \subseteq \mathcal{N}$, it suffices to show:

$$\Delta; \Gamma \vdash \text{fix } f(x). e_1 \sim_{w\uparrow} \text{fix } f(x). e_2 : \sigma_1 \rightarrow \sigma_2$$

We need to show $\text{inhabited}(w\uparrow)$ and $\text{consistent}(w\uparrow)$ —the rest follows by construction of w_2 . Regarding inhabitation, by Lemma 10 it suffices to show $\text{inhabited}(w_2\uparrow)$, which

is witnessed by state $(\emptyset, 1)$. Regarding consistency, using the same lemma, we could simply show $\text{consistent}(w_2\uparrow)$. However, since $w_2.\mathbf{L}$ refers to e_2 and e'_2 , there is no hope to show this. Instead, we have to prove $\text{consistent}(w\uparrow)$ directly.

The part of $\text{consistent}(w\uparrow)$ concerning universal types follows easily by Lemma 9 from $\text{consistent}(w_1\uparrow)$, because $w_2.\mathbf{L}$ is empty at universal types.

Concerning arrow types, we are given $G, s_{\text{ref}}, s_1, s_2, \sigma'_1, \sigma'_2$, related functions

$$(v_1, v_2) \in w\uparrow.\mathbf{L}(s_{\text{ref}}, s_1, s_2)(G(s_{\text{ref}}, s_1, s_2))(\sigma'_1 \rightarrow \sigma'_2)$$

and related arguments $(v'_1, v'_2) \in \overline{G(s_{\text{ref}}, s_1, s_2)}(\sigma'_1)$, and must show:

$$(\text{step}(v_1 \ v'_1), \text{step}(v_2 \ v'_2)) \in \mathbf{E}_{w\uparrow}(G)(s_{\text{ref}}, s_1, s_2)(s_{\text{ref}}, s_1, s_2)(\sigma'_2)$$

If the functions are related by $w_1\uparrow$,

$$(v_1, v_2) \in w_1\uparrow.\mathbf{L}(s_{\text{ref}}, s_1)(G(s_{\text{ref}}, s_1, s_2))(\sigma'_1 \rightarrow \sigma'_2)$$

then the claim follows from $\text{consistent}(w_1\uparrow)$ with Lemma 9.

Otherwise, the functions are necessarily from w_2 :

$$(v_1, v_2) \in w_2.\mathbf{L}(s_{\text{ref}})(s_2)(G(s_{\text{ref}}, s_1, s_2))(\sigma'_1 \rightarrow \sigma'_2)$$

Hence $\sigma'_1 \rightarrow \sigma'_2 = \delta\sigma_1 \rightarrow \delta\sigma_2$ and $v_1 = \gamma_1 \text{fix } f(x). e_1$ and $v_2 = \gamma_2 \text{fix } f(x). e_2$ for some δ and $(\gamma_1, \gamma_2) \in \mathbf{Env}(\delta\Gamma, G(s_{\text{ref}}, s_1, s_2))$. Writing γ'_i for $\gamma_i, f \mapsto v_i, x \mapsto v'_i$, it remains to show

$$(\gamma'_1 e_1, \gamma'_2 e_2) \in \mathbf{E}_{w\uparrow}(G)(s_{\text{ref}}, s_1, s_2)(s_{\text{ref}}, s_1, s_2)(\delta\sigma_2),$$

which reduces by Lemma 9 to

$$(\gamma'_1 e_1, \gamma'_2 e_2) \in \mathbf{E}_{w_1\uparrow}(G(-, -, s_2))(s_{\text{ref}}, s_1)(s_{\text{ref}}, s_1)(\delta\sigma_2).$$

This follows from the instantiated premise if we can show:

$$(\gamma'_1, \gamma'_2) \in \mathbf{Env}(\delta(\Gamma, f:\sigma_1 \rightarrow \sigma_2, x:\sigma_1), G(s_{\text{ref}}, s_1, s_2))$$

Given what we know about γ_1, γ_2 , it suffices to focus on their extensions. We already know that $(v'_1, v'_2) \in \overline{G(s_{\text{ref}}, s_1, s_2)}(\delta\sigma_1)$. So it remains to show

$$(\gamma_1 \text{fix } f(x). e_1, \gamma_2 \text{fix } f(x). e_2) \in \overline{G(s_{\text{ref}}, s_1, s_2)}(\delta\sigma_1 \rightarrow \delta\sigma_2),$$

which, exploiting that $G(-, s_1, -) \in \mathbf{GK}(w_2\uparrow)$, we do by showing

$$(\gamma_1 \text{fix } f(x). e_1, \gamma_2 \text{fix } f(x). e_2) \in w_2\uparrow.\mathbf{L}(s_{\text{ref}}, s_2)(G(s_{\text{ref}}, s_1, s_2))(\delta\sigma_1 \rightarrow \delta\sigma_2).$$

Indeed, this holds by construction of w_2 . □

3.7.4 Congruency

By an easy induction, the compatibility lemmas from Figure 3.17 yield reflexivity for well-typed programs.

Theorem 4 (Reflexivity).

$$\frac{\Delta; \Gamma \vdash p : \sigma}{\Delta; \Gamma \vdash p \sim p : \sigma}$$

We also have transitivity of \sim , hence it is a proper equivalence relation (we already proved symmetry in Section 3.7.2).

Theorem 5 (Transitivity).

$$\frac{\Delta; \Gamma \vdash p_1 \sim p_2 : \sigma \quad \Delta; \Gamma \vdash p_2 \sim p_3 : \sigma}{\Delta; \Gamma \vdash p_1 \sim p_3 : \sigma}$$

Proof. The subject of Section 3.9. □

Lemma 24 (Context closure).

$$\frac{\Delta; \Gamma \vdash p_1 \sim p_2 : \sigma \quad \vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\Delta'; \Gamma'; \sigma')}{\Delta'; \Gamma' \vdash C[p_1] \sim C[p_2] : \sigma'}$$

Proof. By induction on the derivation of the context typing, in each case using the corresponding compatibility lemma. For contexts containing subterms we also need Theorem 4. □

Lemma 25 (Substitutivity).

$$\frac{\Delta; \Gamma, x:\tau' \vdash e_1 \sim e_2 : \tau \quad \Delta; \Gamma \vdash v_1 \sim v_2 : \tau'}{\Delta; \Gamma \vdash e_1[v_1/x] \sim e_2[v_2/x] : \tau}$$

3.7.5 Soundness

Lemma 26 (Adequacy). If $\vdash p_1 \sim p_2 : \tau$, then for any heaps h_1, h_2 :

1. either $\langle h_1; |p_1| \rangle \uparrow \wedge \langle h_2; |p_2| \rangle \uparrow$,
2. or $\exists h'_1, h'_2, v_1, v_2. \langle h_1; |p_1| \rangle \hookrightarrow^* \langle h'_1; v_1 \rangle \wedge \langle h_2; |p_2| \rangle \hookrightarrow^* \langle h'_2; v_2 \rangle$.

Proof. We instantiate the premise with **TyNam** and get w with $\vdash p_1 \sim_{w\uparrow} p_2 : \tau$. Using the least global knowledge (Lemma 7), inhabitation of $w\uparrow$ tells us that there is s_0 with $(\emptyset, \emptyset) \in w\uparrow.H(s_0)([w\uparrow](s_0))$. Now we use this state to obtain

$$(e_1, e_2) \in \mathbf{E}_{w\uparrow}([w\uparrow])(s_0)(s_0)(\tau),$$

which by key Lemma 12 implies

$$(e_1, e_2) \in \mathbf{E}_{w\uparrow}^{\lambda s, \emptyset}([w\uparrow])(s_0)(s_0)(\tau).$$

Instantiating this with the empty heaps $(\emptyset, \emptyset) \in w \uparrow. \mathbf{H}(s_0)([w \uparrow](s_0))$ as well as with empty frame heaps, yields three cases. Cases **DIV** and **EVAL** imply the goal. Case **CALL** cannot possibly hold, because the called functions would be related by the empty relation. \square

Finally, combining adequacy and congruency, we can show our main soundness theorem: for well-typed programs, our equivalence relation is included in contextual equivalence.

Theorem 6 (Soundness). If $\Delta; \Gamma \vdash p_1 : \sigma$ and $\Delta; \Gamma \vdash p_2 : \sigma$, then:

$$\Delta; \Gamma \vdash p_1 \sim p_2 : \sigma \implies \Delta; \Gamma \vdash p_1 \sim_{\text{ctx}} p_2 : \sigma$$

Proof. Suppose $\Delta; \Gamma \vdash p_1 \sim p_2 : \sigma$ as well as $\vdash C : (\Delta; \Gamma; \sigma) \rightsquigarrow (\epsilon; \epsilon; \tau)$. By Lemma 24 we have $\vdash C[p_1] \sim C[p_2] : \tau$. By adequacy (Lemma 26) we have $\langle h; |C[p_1]| \rangle \uparrow \iff \langle h; |C[p_2]| \rangle \uparrow$ for any h , so we are done. \square

3.8 Examples

In this section, we present several example PB equivalence proofs. Here we will not actually rely on the ability of local knowledges and heap relations to depend on the state of the reference world W_{ref} . Consequently, the stability property in the definition of program equivalence will be trivially satisfied (by choosing $s' := s$). However, dependent worlds are of critical importance in the transitivity proof discussed in Section 3.9.

3.8.1 Well-Bracketed State Change

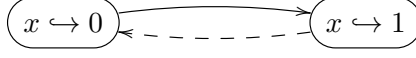
Recall the example from Section 2.2.3 and its high-level proof-sketch using an STS. We will now show in some formal detail how this proof is done using our method. Concretely, we prove $\vdash v_1 \sim e_2 : \tau$, where:

$$\begin{aligned} \tau &:= (\text{unit} \rightarrow \text{unit}) \rightarrow \text{nat} \\ v_1 &:= \lambda f. (f \langle \rangle; f \langle \rangle; 1) \\ e_2 &:= \text{let } x = \text{ref } 0 \text{ in } v_2 \\ v_2 &:= \lambda f. (x := 0; f \langle \rangle; x := 1; f \langle \rangle; !x) \end{aligned}$$

Constructing a Suitable World We construct a local world w that we will then show to be consistent and to relate v_1 and e_2 .

$$\begin{aligned} w.\mathbf{T.S} &:= \text{Loc} \xrightarrow{\text{fin}} \{0, 1\} \subset \text{Heap} \\ w.\mathbf{T.}\sqsubseteq &:= \{(s, s') \mid \text{dom}(s) \subseteq \text{dom}(s')\} \\ w.\mathbf{T.}\sqsubseteq_{\text{pub}} &:= \{(s, s') \in w.\sqsubseteq \mid \forall (l, 1) \in s. (l, 1) \in s'\} \\ w.\mathbf{L}(_, s)(G)(\tau) &:= \{(v_1, v_2[l/x]) \mid l \in \text{dom}(s)\} \\ w.\mathbf{H}(_, s)(G) &:= \{(\emptyset, s)\} \\ w.\mathbf{N} &:= \emptyset \end{aligned}$$

The STS that we build into w is essentially the one from Section 2.2.3:



A state $s \in w.S$ is to be understood as follows: for each running instance of e_2 , identified by the location l_x that that instance initially allocated, $s(l_x)$ says whether the instance is in the (pictorially) left state (l_x points to 0) or in the right one (l_x points to 1). Accordingly, the heap relation $w.H$ at state s is just $\{(\emptyset, s)\}$. Finally, the local knowledge $w.L$ at state s relates v_1 with $v_2[l_x/x]$ for any location l_x belonging to an instance. Finally, since the programs don't involve abstract types, we can define $w.N$ to be empty.

It is easy to see that $w \in \text{LWorld}$. In particular, $w.L$ and $w.H$ are monotone as required. Note that $\text{stable}(w)$ (the dependency is vacuous) and that $\text{inhabited}(w \uparrow)$ for $s_0 = (\emptyset, \emptyset)$. To show $\vdash v_1 \sim_{w \uparrow} e_2 : \tau$, two parts remain.

Proving Its Consistency Establishing $\text{consistent}(w \uparrow)$ is the real meat of the proof. Consider two functions related by $w \uparrow.L$ at a state (s_{ref}^1, s_1) . Clearly, one is v_1 and the other is $v_2[l/x]$ for some $l \in \text{dom}(s_1)$. Now suppose we are given related arguments $(v'_1, v'_2) \in G_1(s_{\text{ref}}^1, s_1)(\text{unit} \rightarrow \text{unit})$. We need to show:

$$((v'_1 \langle \rangle; v'_1 \langle \rangle; 1), (l := 0; v'_2 \langle \rangle; l := 1; v'_2 \langle \rangle; !l)) \in \mathbf{E}_{w \uparrow}(G_1)(s_{\text{ref}}^1, s_1)(s_{\text{ref}}^1, s_1)(\text{nat})$$

Note that for $(h_1, h_2) \in w.H(s_{\text{ref}}^1, s_1)(G(s_{\text{ref}}^1, s_1))$ we know by construction that $h_1 = \emptyset$ and $h_2 = s_1$. Consequently, for any frame heaps h_1^F, h_2^F , we have

$$\begin{aligned} & \langle h_1 \sqcup h_1^F; (v'_1 \langle \rangle; v'_1 \langle \rangle; 1) \rangle \\ \hookrightarrow^* & \langle h_1 \sqcup h_1^F; (v'_1 \langle \rangle; v'_1 \langle \rangle; 1) \rangle \\ & \langle h_2 \sqcup h_2^F; (l := 0; v'_2 \langle \rangle; l := 1; v'_2 \langle \rangle; !l) \rangle \\ \hookrightarrow^* & \langle (s_1 \setminus l) \sqcup [l \mapsto 0] \sqcup h_2^F; (v'_2 \langle \rangle; l := 1; v'_2 \langle \rangle; !l) \rangle \end{aligned}$$

where $s_1 \setminus l$ denotes the restriction of s_1 to domain $\text{dom}(s_1) \setminus \{l\}$.

By the **CALL** disjunct in the definition of $\mathbf{E}_{w \uparrow}$ it now suffices to find $s'_1 \sqsupseteq s_1$ such that:

1. $(\emptyset, (s_1 \setminus l) \sqcup [l \mapsto 0]) \in w.H(s_{\text{ref}}^1, s'_1)(G_1(s_{\text{ref}}^1, s'_1))$
2. $((\bullet; v'_1 \langle \rangle; 1), (\bullet; l := 1; v'_2 \langle \rangle; !l)) \in \mathbf{K}_{w \uparrow}(G_1)(s_{\text{ref}}^1, s_1)(s_{\text{ref}}^1, s'_1)(\text{unit}, \text{nat})$

Naturally, we pick $s'_1 = (s_1 \setminus l) \sqcup [l \mapsto 0]$. Note that $s'_1 \sqsupseteq s_1$ because their domains coincide. Then (1) holds by construction of w and it remains to show (2).

After unfolding the definition of \mathbf{K} and being given $(s_{\text{ref}}^2, s_2) \sqsupseteq_{\text{pub}} (s_{\text{ref}}^1, s'_1)$ as well as $G_2 \sqsupseteq G_1$, we repeat the previous chain of arguments one more time and arrive at the goal of finding $s'_2 \sqsupseteq s_2$ such that:

3. $(\emptyset, (s_2 \setminus l) \sqcup [l \mapsto 1]) \in w.H(s_{\text{ref}}^2, s'_2)(G_2(s_{\text{ref}}^2, s'_2))$

$$4. ((\bullet; 1), (\bullet; !l)) \in \mathbf{K}_{w\uparrow}(G_2)(s_{\text{ref}}^1, s_1)(s_{\text{ref}}^2, s'_2)(\text{unit}, \text{nat})$$

Naturally, we pick $s'_2 = (s_2 \setminus l) \sqcup [l \mapsto 1] \sqsupseteq s_2$. Then (3) holds by construction of w and it remains to show (4).

We observe that, for any $(s_{\text{ref}}^3, s_3) \sqsupseteq_{\text{pub}} (s_{\text{ref}}^2, s'_2)$, it must be that $s_3(l) = 1$ since $s'_2(l) = 1$. Hence for $(h'_1, h'_2) \in w.\mathbf{H}(s_{\text{ref}}^3, s_3)(G_3(s_{\text{ref}}^3, s_3))$ we know by construction $h'_2(l) = 1$. Consequently, for any frame heaps $h_1^{\text{F}'}, h_2^{\text{F}'}$, we have:

$$\langle h'_1 \sqcup h_1^{\text{F}'}; \langle \rangle; 1 \rangle \hookrightarrow^* \langle h'_1 \sqcup h_1^{\text{F}'}; 1 \rangle$$

$$\langle h'_2 \sqcup h_2^{\text{F}'}; \langle \rangle; !l \rangle \hookrightarrow^* \langle h'_2 \sqcup h_2^{\text{F}'}; 1 \rangle$$

Hence we are done if we can show $(s_{\text{ref}}^3, s_3) \sqsupseteq_{\text{pub}} (s_{\text{ref}}^1, s_1)$. Indeed, this is easy to verify.

Showing the Programs Related By It Given how we constructed our local world, this final goal is fairly easy to accomplish. Formally, we must show

$$(v_1, e_2) \in \mathbf{E}_{w\uparrow}(G)(s_{\text{ref}}, s)(s_{\text{ref}}, s)(\tau)$$

for any G, s_{ref}, s . Note that if $(h_1, h_2) \in w.\mathbf{H}(s_{\text{ref}}, s)(G(s_{\text{ref}}, s))$, then

$$\langle h_1; v_1 \rangle \hookrightarrow^* \langle h_1; v_1 \rangle$$

$$\langle h_2; e_2 \rangle \hookrightarrow \langle h_2 \sqcup [l \mapsto 0]; v_2[l/x] \rangle$$

for some fresh l . It therefore suffices to find $s' \sqsupseteq_{\text{pub}} s$ such that the following hold:

$$5. (h_1, h_2 \sqcup [l \mapsto 0]) \in w.\mathbf{H}(s_{\text{ref}}, s')(G(s_{\text{ref}}, s'))$$

$$6. (v_1, v_2[l/x]) \in \overline{G(s_{\text{ref}}, s')}(\tau)$$

We pick $s' = s \sqcup [l \mapsto 0]$. Note that s' is well-defined because l is fresh (so $l \notin \text{dom}(s)$), and that $s' \sqsupseteq_{\text{pub}} s$ as required. To show (6), it suffices by definition of **GK** to show $(v_1, v_2[l/x]) \in w.\mathbf{L}(s_{\text{ref}}, s')(G(s_{\text{ref}}, s'))(\tau)$. This holds by construction of w and s' , and so does (5).

Notice that our world in this proof was very simple—we didn't actually use the global knowledge parameter. An easy adaptation of the example programs, however, would necessitate that—for instance, if we were to replace the constants 0 and 1 with free variables a and b , respectively.

3.8.2 Twin Abstraction

This final example (originally due to Ahmed *et al.* [4]) demonstrates the interaction of local state with abstract types.

$$\begin{aligned}
\tau &:= \exists \alpha. \exists \beta. (\text{unit} \rightarrow \alpha) \times (\text{unit} \rightarrow \beta) \times (\alpha \times \beta \rightarrow \text{nat}) \\
e_1 &:= \text{let } x = \text{ref } 0 \text{ in} \\
&\quad \text{pack } \langle \text{nat}, \text{pack } \langle \text{nat}, \lambda_{-}. x := !x + 1; !x, \\
&\quad \quad \lambda_{-}. x := !x + 1; !x, \\
&\quad \quad \lambda p. p.1 = p.2 \rangle \rangle \\
e_2 &:= \text{let } x = \text{ref } 0 \text{ in} \\
&\quad \text{pack } \langle \text{nat}, \text{pack } \langle \text{nat}, \lambda_{-}. x := !x + 1; !x, \\
&\quad \quad \lambda_{-}. x := !x + 1; !x, \\
&\quad \quad \lambda p. 0 \rangle \rangle
\end{aligned}$$

Both e_1 and e_2 return a name generator ADT consisting of two abstract types α and β , together with a function for generating a fresh name of type α , a function for generating a fresh name of type β , and a function for comparing an α name and a β name for equality. Both implementations represent names as integers, and in e_1 , the comparison operation really tests the names for equality. In e_2 , however, the comparison just always returns false right away. Nevertheless, the two programs are contextually equivalent because the α names and the β names are generated by the *same* underlying integer counter, and thus no value can be both an α name and a β name at the same time.

We now show the construction of a local world w that can be used to prove $\vdash e_1 \sim e_2 : \tau$. Let a countably infinite $\mathcal{N} \in \mathcal{P}(\text{TyNam})$ be given. Since Loc is also countably infinite, we can think of \mathcal{N} as being split into $\{\mathbf{n}_l^\alpha \mid l \in \text{Loc}\} \uplus \{\mathbf{n}_l^\beta \mid l \in \text{Loc}\}$.

We define w as follows:

$$\begin{aligned}
w.\mathbf{T}.S &:= \{s \in \mathbf{Loc} \times \mathbf{Loc} \xrightarrow{\text{fin}} \mathcal{P}(\mathbb{N}_{>0}) \times \mathcal{P}(\mathbb{N}_{>0}) \mid \\
&\quad \text{dom}(s) \text{ partial bijection} \wedge \\
&\quad \forall (l_1, l_2, S_1, S_2) \in s. S_1 \cap S_2 = \emptyset\} \\
w.\mathbf{T}.\sqsubseteq &:= \{(s, s') \mid \forall (l_1, l_2, S_1, S_2) \in s. \exists S'_1 \supseteq S_1, S'_2 \supseteq S_2. (l_1, l_2, S'_1, S'_2) \in s'\} \\
w.\mathbf{T}.\sqsubseteq_{\text{pub}} &:= w.\mathbf{T}.\sqsubseteq \\
w.\mathbf{L}(s_{\text{ref}}, s)(G) &:= \{(\mathbf{n}_{l_1}^\alpha, n, n) \mid \exists l_2, S_1, S_2. (l_1, l_2, S_1, S_2) \in s \wedge n \in S_1\} \\
&\quad \uplus \{(\mathbf{n}_{l_1}^\beta, n, n) \mid \exists l_2, S_1, S_2. (l_1, l_2, S_1, S_2) \in s \wedge n \in S_2\} \\
&\quad \uplus \{((\text{unit} \rightarrow \mathbf{n}_{l_1}^\alpha), ++l_1, ++l_2) \mid (l_1, l_2) \in \text{dom}(s)\} \\
&\quad \uplus \{((\text{unit} \rightarrow \mathbf{n}_{l_1}^\beta), ++l_1, ++l_2) \mid (l_1, l_2) \in \text{dom}(s)\} \\
&\quad \uplus \{((\mathbf{n}_{l_1}^\alpha \times \mathbf{n}_{l_1}^\beta \rightarrow \text{nat}), (\lambda p. p.1 = p.2), (\lambda p. 0)) \mid \exists l_2. (l_1, l_2) \in \text{dom}(s)\} \\
w.\mathbf{H}(s_{\text{ref}}, s)(G) &:= \{(h_1, h_2) \mid \\
&\quad \text{dom}(h_1) = \{l_1 \mid \exists l_2. (l_1, l_2) \in \text{dom}(s)\} \wedge \\
&\quad \text{dom}(h_2) = \{l_2 \mid \exists l_1. (l_1, l_2) \in \text{dom}(s)\} \wedge \\
&\quad \forall (l_1, l_2, S_1, S_2) \in s. h_1(l_1) = h_2(l_2) = \max(\{0\} \uplus S_1 \uplus S_2)\} \\
w.\mathbf{N} &:= \mathcal{N}
\end{aligned}$$

Here, $++l$ is short for $\lambda_. l := !l + 1; !l$.

Similar to the world construction in Section 3.8.1, states $s \in w.S$ are functions defined for those locations (l_1, l_2) that, intuitively, were allocated in an instance of e_1 and e_2 , respectively. They are mapped to sets S_1 and S_2 of positive integers, representing⁵ the current inhabitants of the abstract types α and β , respectively, for that instance. The crucial invariant here is that S_1 and S_2 are always disjoint. The local knowledge $w.L$ declares e_1 's functions equivalent to those of e_2 ; it also defines the meaning of type $\mathbf{n}_{l_1}^\alpha$ as the identity relation restricted to those numbers that inhabit α in the instance pair identified by l_1 (and similarly for $\mathbf{n}_{l_1}^\beta$ and β).

According to this interpretation, the transition relation only allows S_1 and S_2 to grow (the distinction between public and private transitions is not needed for this example). Finally, $w.H$ says that the related heaps at state s are any (h_1, h_2) where h_i contains exactly those locations allocated in instances of e_i , and each such location stores the largest value handed out so far (no matter if at α or β). This latter condition is critical to ensure that S_1 and S_2 stay disjoint in each instance.

Using this world, it is straightforward to finish the proof.

⁵To keep the definitions as simple as possible, the state space includes some states that actual program behaviour cannot result in (but that nevertheless are consistent with the property we want to prove).

3.8.3 World Generator

As explained in Section 3.6.7, and as one may observe from the previous examples, worlds must often describe “ n -ary” state spaces, where each state consists of n copies of states drawn from some simpler state space, one copy for each dynamic instance of the object or ADT. Thus, it would be convenient if one were able to reason about program equivalence under the degenerate case of a single copy (*i.e.*, $n = 1$). Fortunately, it is not hard to (i) define a *world generator* that, given a single-instance world, automatically performs the multiplexing; and (ii) show that consistency of that single-instance world implies equivalence consistency of the multiplexed world.

Here we present the definition and properties of such a world generator. In the next section, we use it to simplify the proof of (a variant of) the example from Section 3.8.1

$$\text{NLWorld} := \{\mathcal{W} \in \text{Names} \rightarrow \text{LWorld} \mid \forall \mathcal{N}. \mathcal{W}(\mathcal{N}).\mathbf{N} \subseteq \mathcal{N}\}$$

Definition 9 (World generator). We define $\mathbf{G} \in \text{NLWorld} \rightarrow \text{NLWorld}$ as follows.

$$\begin{aligned} \mathbf{G}(\mathcal{W})(\mathcal{N}).\mathbf{T.S} &:= \{(s_1, \dots, s_n) \mid n \in \mathbb{N} \wedge \forall i \in \{1, \dots, n\}. s_i \in \mathcal{W}(\mathcal{N}_i).\mathbf{S}\} \\ \mathbf{G}(\mathcal{W})(\mathcal{N}).\mathbf{T}.\sqsubseteq &:= \text{see below} \\ \mathbf{G}(\mathcal{W})(\mathcal{N}).\mathbf{T}.\sqsubseteq_{\text{pub}} &:= \text{see below} \\ \mathbf{G}(\mathcal{W})(\mathcal{N}).\mathbf{L}(s_{\text{ref}}, (s_1, \dots, s_n))(R) &:= \bigcup_{i \in \{1, \dots, n\}} \mathcal{W}(\mathcal{N}_i).\mathbf{L}(s_{\text{ref}}, s_i)(R) \\ \mathbf{G}(\mathcal{W})(\mathcal{N}).\mathbf{H}(s_{\text{ref}}, (s_1, \dots, s_n))(R) &:= \bigotimes_{i \in \{1, \dots, n\}} \mathcal{W}(\mathcal{N}_i).\mathbf{H}(s_{\text{ref}}, s_i)(R) \\ \mathbf{G}(\mathcal{W})(\mathcal{N}).\mathbf{N} &:= \mathcal{N} \end{aligned}$$

Here $\{\mathcal{N}_i\}$ is a countably infinite splitting of \mathcal{N} into countably infinite sets, *i.e.*, $\mathcal{N} = \mathcal{N}_1 \uplus \mathcal{N}_2 \uplus \mathcal{N}_3 \uplus \dots$ and each \mathcal{N}_i is countably infinite. Many such splittings exists, and it does not matter which one we choose.

The transition relations $\mathbf{G}(\mathcal{W})(\mathcal{N}).\sqsubseteq$ and $\mathbf{G}(\mathcal{W})(\mathcal{N}).\sqsubseteq_{\text{pub}}$ are defined as the least preorders closed under the corresponding following rules.

$$\begin{array}{c} \frac{s'_1 \sqsubseteq_{\text{pub}} s_1 \quad \dots \quad s'_k \sqsubseteq_{\text{pub}} s_k}{(s'_1, \dots, s'_k) \sqsubseteq_{\text{pub}} (s_1, \dots, s_k)} \quad \frac{}{(s_1, \dots, s_k, s_{k+1}) \sqsubseteq_{\text{pub}} (s_1, \dots, s_k)} \\[10pt] \frac{s'_1 \sqsubseteq s_1 \quad \dots \quad s'_k \sqsubseteq s_k}{(s'_1, \dots, s'_k) \sqsubseteq (s_1, \dots, s_k)} \quad \frac{s' \sqsubseteq_{\text{pub}} s}{s' \sqsubseteq s} \end{array}$$

Hence in addition to pointwise transitions, they also allow extensions (moving from a state of arity k to one of arity $k + n$).

For the next lemmas, we define the following notation.

$$G[s_k]_{k \in \{n \setminus i\}} := \lambda s_{\text{ref}}, s'_i. G(s_{\text{ref}}, (s_1, \dots, s_{i-1}, s'_i, s_{i+1}, \dots, s_n))$$

Lemma 27.

$$\forall G \in \mathbf{GK}(\mathbf{G}(\mathcal{W})(\mathcal{N})\uparrow). \forall s_1 \dots s_{i-1}, s_{i+1} \dots s_n. G[s_k]_{k \in \{n \setminus i\}} \in \mathbf{GK}(\mathcal{W}(\mathcal{N}_i)\uparrow)$$

Proof. The fact that $G[s_k]_{k \in \{n \setminus i\}}$ is monotone w.r.t. \sqsubseteq follows directly from the definition of \sqsubseteq and monotonicity of G .

It remains to show

$$G[s_k]_{k \in \{n \setminus i\}}(s_{\text{ref}}, s_i) \supseteq_{\text{ref}}^{\mathcal{N}_i} \mathcal{W}(\mathcal{N}_i) \uparrow \cdot \mathbf{L}(s_{\text{ref}}, s_i)(G[s_k]_{k \in \{n \setminus i\}}(s_{\text{ref}}, s_i))$$

for any s_{ref}, s_i , which we do as follows.

$$\begin{aligned} & G[s_k]_{k \in \{n \setminus i\}}(s_{\text{ref}}, s_i) \\ = & G(s_{\text{ref}}, s_1, \dots, s_n) \\ \supseteq_{\text{ref}}^{\mathcal{N}} & \mathbf{G}(\mathcal{W})(\mathcal{N}) \uparrow \cdot \mathbf{L}(s_{\text{ref}}, s_1, \dots, s_n)(G(s_{\text{ref}}, s_1, \dots, s_n)) \quad (G \in \mathbf{GK}(\mathbf{G}(\mathcal{W})(\mathcal{N}) \uparrow)) \\ \supseteq_{\text{ref}}^{\mathcal{N}_i} & \mathbf{G}(\mathcal{W})(\mathcal{N}) \uparrow \cdot \mathbf{L}(s_{\text{ref}}, s_1, \dots, s_n)(G(s_{\text{ref}}, s_1, \dots, s_n)) \quad (\mathcal{N}_i \subseteq \mathcal{N}) \\ \supseteq_{\text{ref}}^{\mathcal{N}_i} & \mathcal{W}(\mathcal{N}_i) \uparrow \cdot \mathbf{L}(s_{\text{ref}}, s_i)(G(s_{\text{ref}}, s_1, \dots, s_n)) \\ = & \mathcal{W}(\mathcal{N}_i) \uparrow \cdot \mathbf{L}(s_{\text{ref}}, s_i)(G[s_k]_{k \in \{n \setminus i\}}(s_{\text{ref}}, s_i)) \end{aligned}$$

The last inequality relies on $\mathcal{W}(\mathcal{N}_i) \in \mathbf{LWorld}$ and the disjointness of $\mathcal{N}_1, \dots, \mathcal{N}_n$. \square

Lemma 28. If $W = \mathbf{G}(\mathcal{W})(\mathcal{N}) \uparrow$ and $\forall \mathcal{N}'. \text{stable}(\mathcal{W}(\mathcal{N}'))$ and $G \in \mathbf{GK}(W)$, then:

1. $\mathbf{E}_{\mathcal{W}(\mathcal{N}_i) \uparrow}(G[s_k]_{k \in \{n \setminus i\}})(s_{\text{ref}}^0, s_i^0)(s_{\text{ref}}, s_i) \subseteq \mathbf{E}_W(G)(s_{\text{ref}}^0, s_1, \dots, s_{i-1}, s_i^0, s_{i+1}, \dots, s_n)(s_{\text{ref}}, s_1, \dots, s_n)$
2. $\mathbf{K}_{\mathcal{W}(\mathcal{N}_i) \uparrow}(G[s_k]_{k \in \{n \setminus i\}})(s_{\text{ref}}^0, s_i^0)(s_{\text{ref}}, s_i) \subseteq \mathbf{K}_W(G)(s_{\text{ref}}^0, s_1, \dots, s_{i-1}, s_i^0, s_{i+1}, \dots, s_n)(s_{\text{ref}}, s_1, \dots, s_n)$

Proof. By coinduction. \square

Lemma 29. Suppose $\forall \mathcal{N}. \text{stable}(\mathcal{W}(\mathcal{N}))$.

1. $\forall \mathcal{N}. \text{stable}(\mathbf{G}(\mathcal{W})(\mathcal{N}))$
2. If $\forall \mathcal{N}. \text{consistent}(\mathcal{W}(\mathcal{N}) \uparrow)$, then $\forall \mathcal{N}. \text{consistent}(\mathbf{G}(\mathcal{W})(\mathcal{N}) \uparrow)$.

Lemma 30. $\text{inhabited}(\mathbf{G}(\mathcal{W})(\mathcal{N}) \uparrow)$

Proof. It is easy to check that $(\emptyset, \emptyset) \in \mathbf{G}(\mathcal{W})(\mathcal{N}) \uparrow \cdot \mathbf{H}(\emptyset, ())(R)$ for any R . \square

In the next section, we illustrate the use of \mathbf{G} and these lemmas.

3.8.4 Twin Abstraction, Alternate Proof

Recall the example from Section 3.8.2. We now sketch how to prove it with the help of the world generator from the previous section.

Constructing the world.

$$\begin{aligned}
w.\mathbf{T}.S &:= \{s \in \mathbf{Loc} \times \mathbf{Loc} \stackrel{\text{fin}}{\rightarrow} \mathcal{P}(\mathbb{N}_{>0}) \times \mathcal{P}(\mathbb{N}_{>0}) \mid \\
&\quad \text{dom}(s) \text{ partial bijection} \wedge \\
&\quad \forall (l_1, l_2, S_1, S_2) \in s. S_1 \cap S_2 = \emptyset\} \\
w.\mathbf{T}.\sqsubseteq &:= \{(s, s') \mid \forall (l_1, l_2, S_1, S_2) \in s. \exists S'_1 \supseteq S_1, S'_2 \supseteq S_2. (l_1, l_2, S'_1, S'_2) \in s'\} \\
w.\mathbf{T}.\sqsubseteq_{\text{pub}} &:= w.\mathbf{T}.\sqsubseteq \\
w.\mathbf{L}(s_{\text{ref}}, s)(G) &:= \{(\mathbf{n}_{l_1}^\alpha, n, n) \mid \exists l_2, S_1, S_2. (l_1, l_2, S_1, S_2) \in s \wedge n \in S_1\} \\
&\quad \uplus \{(\mathbf{n}_{l_1}^\beta, n, n) \mid \exists l_2, S_1, S_2. (l_1, l_2, S_1, S_2) \in s \wedge n \in S_2\} \\
&\quad \uplus \{((\text{unit} \rightarrow \mathbf{n}_{l_1}^\alpha), ++l_1, ++l_2) \mid (l_1, l_2) \in \text{dom}(s)\} \\
&\quad \uplus \{((\text{unit} \rightarrow \mathbf{n}_{l_1}^\beta), ++l_1, ++l_2) \mid (l_1, l_2) \in \text{dom}(s)\} \\
&\quad \uplus \{((\mathbf{n}_{l_1}^\alpha \times \mathbf{n}_{l_1}^\beta \rightarrow \text{nat}), (\lambda p. p.1 = p.2), (\lambda p. 0)) \mid \exists l_2. (l_1, l_2) \in \text{dom}(s)\} \\
w.\mathbf{H}(s_{\text{ref}}, s)(G) &:= \{(h_1, h_2) \mid \\
&\quad \text{dom}(h_1) = \{l_1 \mid \exists l_2. (l_1, l_2) \in \text{dom}(s)\} \wedge \\
&\quad \text{dom}(h_2) = \{l_2 \mid \exists l_1. (l_1, l_2) \in \text{dom}(s)\} \wedge \\
&\quad \forall (l_1, l_2, S_1, S_2) \in s. h_1(l_1) = h_2(l_2) = \max(\{0\} \uplus S_1 \uplus S_2)\} \\
w.\mathbf{N} &:= \mathcal{N}
\end{aligned}$$

First, we define $\mathcal{W} \in \text{NLWorld}$ as follows:

$$\begin{aligned}
\mathcal{W}(\mathcal{N}').\mathbf{T}.S &:= \{(l_1, l_2, S_1, S_2) \in \mathbf{Loc} \times \mathbf{Loc} \times \mathcal{P}(\mathbb{N}_{>0}) \times \mathcal{P}(\mathbb{N}_{>0}) \mid S_1 \cap S_2 = \emptyset\} \\
\mathcal{W}(\mathcal{N}').\mathbf{T}.\sqsubseteq &:= \{((l_1, l_2, S_1, S_2), (l'_1, l'_2, S'_1, S'_2)) \mid l_1 = l'_1 \wedge l_2 = l'_2 \wedge S_1 \subseteq S'_1 \wedge S_2 \subseteq S'_2\} \\
\mathcal{W}(\mathcal{N}').\mathbf{T}.\sqsubseteq_{\text{pub}} &:= \mathcal{W}(\mathcal{N}').\mathbf{T}.\sqsubseteq \\
\mathcal{W}(\mathcal{N}').\mathbf{L}(s_{\text{ref}}, (l_1, l_2, S_1, S_2))(R) &:= \\
&\quad \{(\mathcal{N}'(1), n, n) \mid n \in S_1\} \uplus \{(\mathcal{N}'(2), n, n) \mid n \in S_2\} \uplus \\
&\quad \{((\text{unit} \rightarrow \mathcal{N}'(1)), ++l_1, ++l_2)\} \uplus \{((\text{unit} \rightarrow \mathcal{N}'(2)), ++l_1, ++l_2)\} \uplus \\
&\quad \{((\mathcal{N}'(1) \times \mathcal{N}'(2) \rightarrow \text{nat}), (\lambda p. p.1 = p.2), (\lambda p. 0))\} \\
\mathcal{W}(\mathcal{N}').\mathbf{H} &:= \lambda(l_1, l_2, S_1, S_2). R. \{([l_1 \mapsto n], [l_2 \mapsto n]) \mid n = \max(\{0\} \uplus S_1 \uplus S_2)\} \\
\mathcal{W}(\mathcal{N}').\mathbf{N} &:= \{\mathcal{N}'(1), \mathcal{N}'(2)\}
\end{aligned}$$

where $\mathcal{N}'(1)$ and $\mathcal{N}'(2)$ denote two distinct elements of \mathcal{N}' .

3.8.5 A Free Theorem

The last example demonstrates the treatment of universal types, and the fact that our method may be used to prove at least some simple so-called “free theorems” [90]. Suppose $\vdash p : \forall \alpha. \alpha$ and $|p| = v$. We want to prove that $\langle h; v [] \rangle \uparrow$ for any h , which implies that there are no closed values of type $\forall \alpha. \alpha$.

We start by applying reflexivity to obtain $\vdash v \sim v : \forall \alpha. \alpha$. This gives us w with $\vdash v \sim_{w\uparrow} v : \forall \alpha. \alpha$ and $w.\mathbf{N} \subseteq \text{TyNam} \setminus \{\mathbf{n}\}$ for some arbitrary \mathbf{n} of our choosing. We

now pick $G := [w\uparrow] \in \mathbf{GK}(w\uparrow)$ (recall Definition 4). From $\vdash v \sim_{w\uparrow} v : \forall\alpha. \alpha$ we then get

$$(v, v) \in \mathbf{E}_{w\uparrow}(G)(s_0)(s_0)(\forall\alpha. \alpha)$$

where s_0 is the state witnessing $\text{inhabited}(w\uparrow)$. Notice that this can only hold due to the second disjunct in \mathbf{E} because v is a value. Consequently, we get $s \sqsupseteq_{\text{pub}} s_0$ such that:

1. $(v, v) \in \overline{G(s)}(\forall\alpha. \alpha) = G(s)(\forall\alpha. \alpha)$
2. $(\emptyset, \emptyset) \in w\uparrow.\mathbf{H}(s)(G(s))$

By construction we have $G(s) = w\uparrow.\mathbf{L}(s)(G(s))$. Hence from (1) and $\text{consistent}(w\uparrow)$, we can derive

$$(\text{step}(v \square), \text{step}(v \square)) \in \mathbf{E}_{w\uparrow}(G)(s)(s)(\tau)$$

for any τ . So, in particular, we have:

$$(\text{step}(v \square), \text{step}(v \square)) \in \mathbf{E}_{w\uparrow}(G)(s)(s)(\mathbf{n})$$

In fact, due to the construction of G , Lemma 12 tells us

$$(\text{step}(v \square), \text{step}(v \square)) \in \mathbf{E}_{w\uparrow}^{\mathcal{R}}(G)(s)(s)(\mathbf{n})$$

for $\mathcal{R} = \lambda s. \emptyset$. Together with (2), this allows only two cases for any heap h : either $\langle h; \text{step}(v \square) \rangle$ diverges (then we are done), or it terminates and the resulting values are related by $\overline{G(s')}(n)$ for some $s' \sqsupseteq_{\text{pub}} s$. However, because $W_{\text{ref}}.\mathbf{N} = \emptyset$ and $w.\mathbf{N} \subseteq \text{TyNam} \setminus \{\mathbf{n}\}$, we know $\overline{G(s')}(n) = w\uparrow.\mathbf{L}(s')(G(s'))(n) = \emptyset$, which rules out that second case.

3.9 Transitivity

In Section 3.5, we presented parametric bisimulations for λ^μ , a very simple fragment of $F^{\mu l}$ lacking quantified types and state. In particular, we sketched the proof of transitivity in Section 3.5.4.

In this section, we generalize that proof of transitivity to account for the full PB model of the full $F^{\mu l}$ language, which turns out (unsurprisingly) to be highly non-trivial. We first describe the high-level structure of our extended transitivity proof. The proof divides into two major parts, presented in Sections 3.9.2 and 3.9.3, respectively. It is highly intricate, requiring several tricky auxiliary constructions. In these sections, we highlight the key technical challenges, explain carefully the central constructions and the intuitions behind them, give detailed proofs of the main lemmas, and give proof sketches of other lemmas. The complete proof is of course fully machine-checked in Coq. To give a rough sense of the complexity, it required approximately 3200 lines of Coq, versus 1500 lines to formalize the language, 400 lines to formalize the model, and 2000 lines to prove soundness of the model w.r.t. contextual equivalence.

3.9.1 Structure of the Transitivity Proof

Given that $\Delta; \Gamma \vdash e_1 \sim e_2 : \tau$ and $\Delta; \Gamma \vdash e_2 \sim e_3 : \tau$, we must show $\Delta; \Gamma \vdash e_1 \sim e_3 : \tau$. Unfolding the definition of this goal, we are given a countably infinite set of type names \mathcal{N} and must construct a stable local world w such that (a) $\Delta; \Gamma \vdash e_1 \sim_{w\uparrow} e_3 : \tau$ and (b) $w.N \subseteq \mathcal{N}$ (i.e., $w.L$ defines no names outside of \mathcal{N}). To do so, we split \mathcal{N} into three disjoint (and also countably infinite) pieces: \mathcal{N}_1 , \mathcal{N}_2 , and \mathcal{N}_\exists . The first two pieces are used to instantiate the assumptions regarding $e_1 \sim e_2$ and $e_2 \sim e_3$, respectively, thus yielding two stable local worlds w_1 and w_2 such that

$$\Delta; \Gamma \vdash e_1 \sim_{w_1\uparrow} e_2 : \tau \quad (3.1)$$

$$\text{and } \Delta; \Gamma \vdash e_2 \sim_{w_2\uparrow} e_3 : \tau \quad (3.2)$$

as well as $w_1.N \subseteq \mathcal{N}_1$ and $w_2.N \subseteq \mathcal{N}_2$. Keeping \mathcal{N}_1 and \mathcal{N}_2 disjoint is a matter of basic hygiene: it ensures that w_1 and w_2 , which we will be using in the construction of w , do not step on each other's toes by defining the same type name in incompatible ways. As for the names in \mathcal{N}_\exists , we reserve them for a special purpose to be explained later.

At this point, the proof divides into two separate parts. In the first part, we use w_1 and w_2 to directly construct a full world W such that $\Delta; \Gamma \vdash e_1 \sim_W e_3 : \tau$ (and $W.N \subseteq \mathcal{N}$). While this construction and the proof of the required properties is quite subtle, they are essentially an extension of the transitivity proof for λ^μ that we sketched in detail in Section 3.5.4 and which we present here in much greater detail. The main novelty over that previous proof is that we now (in this first part) deal with quantified types; reference types do not cause much of a problem.

However, the second part of the proof has all to do with references. Specifically, the world W that we create in the first part does not have the required shape of a lifted local world $w\uparrow$. Thus, in the second part, we (i) develop a theory of *weak isomorphisms* between worlds and prove that they preserve term equivalence, and (ii) construct a stable local world w such that $w\uparrow$ is weakly isomorphic to W .

3.9.2 First Part: Constructing the Full World W

3.9.2.1 High-Level Explanation

As mentioned above, this first part of the proof is essentially agnostic as to whether the language/model supports mutable state. To ease the presentation, we therefore gloss over any state-related details at first; we will be more precise in Section 3.9.2.2. We also write W_i as shorthand for $w_i\uparrow$.

We want to construct W such that $\Delta; \Gamma \vdash e_1 \sim_W e_3 : \tau$. The proofs of consistency of W and **E**-relatedness of closed instances of (e_1, e_3) turn out to be very similar, so let us focus on the latter here. We are given a global knowledge $G \in \text{GK}(W)$ and related substitutions γ_1 and γ_3 . These substitutions map each variable bound in Γ to related values $(v_1, v_3) \in \overline{G}(\tau')$. In order to make use of (3.1) and (3.2), we want to be able to: (i) “decompose” G into global knowledges $G_{(1)} \in \text{GK}(W_1)$ and

$G_{(2)} \in \mathbf{GK}(W_2)$ that would be suitable for instantiating (3.1) and (3.2), and (ii) find a mediating substitution γ_2 s.t. for each $x \in \text{dom}(\Gamma)$, it is the case that $\overline{G_{(1)}}$ relates $(\gamma_1 x, \gamma_2 x)$ and $\overline{G_{(2)}}$ relates $(\gamma_2 x, \gamma_3 x)$. Formally, we want:

$$\overline{G}(\tau) \subseteq \overline{G_{(1)}}(\tau_{(1)}) \circ \overline{G_{(2)}}(\tau_{(2)}) \quad (3.3)$$

where \circ is ordinary relational composition. (Pretend *for now* that $\tau_{(i)} = \tau$. We will soon see why that's not good enough.)

If we have this, we can instantiate (3.1) and (3.2), thus obtaining $(\gamma_1 e_1, \gamma_2 e_2) \in \mathbf{E}(G_{(1)})(\tau_{(1)})$ and $(\gamma_2 e_2, \gamma_3 e_3) \in \mathbf{E}(G_{(2)})(\tau_{(2)})$. It thus remains for us to show $(\gamma_1 e_1, \gamma_3 e_3) \in \mathbf{E}(G)(\tau)$. In other words, we must (unsurprisingly) show some kind of transitivity property for \mathbf{E} :

$$\mathbf{E}_{W_1}(G_{(1)})(\tau_{(1)}) \circ \mathbf{E}_{W_2}(G_{(2)})(\tau_{(2)}) \subseteq \mathbf{E}_W(G)(\tau) \quad (3.4)$$

Proving this will certainly require us to prove the analogous transitivity property for values, which is the inverse of (3.3):

$$\overline{G}(\tau) \supseteq \overline{G_{(1)}}(\tau_{(1)}) \circ \overline{G_{(2)}}(\tau_{(2)}) \quad (3.5)$$

It is therefore only natural that we will define $W.L$ to at least include the composition of $W_1.L$ and $W_2.L$:

$$W.L(G)(\tau) \supseteq W_1.L(G_{(1)})(\tau_{(1)}) \circ W_2.L(G_{(2)})(\tau_{(2)}) \quad (3.6)$$

In order to see what else we want to put into W , and how to define $G_{(i)}$ and $\tau_{(i)}$, let us consider proving (3.3) and (3.5). For a flexible type τ' the conjunction of (3.3) and (3.5) is:

$$G(\tau') = G_{(1)}(\tau'_{(1)}) \circ G_{(2)}(\tau'_{(2)}).$$

One simple (but inadequate) choice for $G_{(i)}$ would be to define it as the minimal global knowledge in $\mathbf{GK}(W_i)$ (introduced in Definition 4). But there is no reason to believe that for any $(v_1, v_3) \in G(\tau')$, there magically happens to be a “mediating” value v_2 such that $W_1.L$ relates (v_1, v_2) and $W_2.L$ relates (v_2, v_3) , as required by the \subseteq part of the above equation. To work this magic, we will explicitly *add* such mediating values to $G_{(1)}$ and $G_{(2)}$! Concretely, we define $G_{(1)}$ as the smallest global knowledge in $\mathbf{GK}(W_1)$ that relates $(v_1, \mathbf{I}(\tau, v_1, v_3))$ at $\tau_{(1)}$ whenever G relates (v_1, v_3) at τ , and similarly $G_{(2)}$ as the smallest global knowledge in $\mathbf{GK}(W_2)$ that relates $(\mathbf{I}(\tau, v_1, v_3), v_3)$ at $\tau_{(2)}$ whenever G relates (v_1, v_3) at τ .

Now, what is this magic \mathbf{I} ? For proving (3.3), it could be anything that maps to \mathbf{CVal} . But for (3.5), it is crucial that each mediating value *uniquely* encodes the corresponding value pair (v_1, v_3) . We therefore require \mathbf{I} to be *injective*. Since all involved sets are countably infinite, such an encoding function exists and we do not care about the particular choice—except that we will choose its range to be of *rigid* type, specifically \mathbf{nat} , for reasons we will explain shortly.

The proofs of (3.3) and (3.5) are by induction on the value closure (recall that it is constructed as a least fixed point). The reason why we must add more to $W.L$ than just the composition in (3.6), and why $\tau_{(i)}$ cannot just be τ in general, has to do with abstract types. We illustrate the issue here for existential types, but the same problem arises for universals (although in a different place, namely in the proof of (3.4)).

Suppose $\tau_{(i)}$ were the identity and consider (3.5) at some type $\exists\alpha.\tau$: We would have to show that if $(\text{pack } v_1, \text{pack } v_2) \in \overline{G_{(1)}}(\exists\alpha.\tau)$ and $(\text{pack } v_2, \text{pack } v_3) \in \overline{G_{(2)}}(\exists\alpha.\tau)$, then $(\text{pack } v_1, \text{pack } v_3) \in \overline{G}(\exists\alpha.\tau)$. Unfolding the value closure, this means that given some τ'_1, τ'_2 with $(v_1, v_2) \in \overline{G_{(1)}}(\tau[\tau'_1/\alpha])$ and $(v_2, v_3) \in \overline{G_{(2)}}(\tau[\tau'_2/\alpha])$, we must come up with τ' such that $(v_1, v_3) \in \overline{G}(\tau[\tau'/\alpha])$. Now, if the two given representation types happen to be the same ($\tau'_1 = \tau'_2$), then we could proceed by just picking $\tau' := \tau'_1$. But of course in general τ'_1 and τ'_2 will be different!

The intuition behind our solution is quite simple: we pick τ' to be a fresh *type name*, which we use to represent the semantic composition of τ'_1 and τ'_2 . More concretely, we use a type name \mathbf{n} from \mathcal{N}_\exists (which we reserved for exactly this purpose) to uniquely encode τ'_1 and τ'_2 , then define \mathbf{n} 's meaning in W to be precisely $\overline{G_{(1)}}(\tau'_1) \circ \overline{G_{(2)}}(\tau'_2)$, and finally choose τ' to be \mathbf{n} . Since we don't know what τ'_1 and τ'_2 are, we simply have to encode all pairs of types this way. To pick the names, we use an injective function

$$\mathbf{A} \in \text{CTy} \times \text{CTy} \rightarrow \mathcal{N}_\exists$$

which, like \mathbf{I} , exists because all involved sets are countably infinite (and, as with \mathbf{I} , we do not care about its concrete definition). It should be clear by now what $\tau'_{(i)}$ does and that it is crucial for making the induction go through: it *decodes* τ' by traversing its structure and replacing each type name \mathbf{n} that equals $\mathbf{A}(\tau_1, \tau_2)$ —for some τ_1, τ_2 —with τ_i .

That's the intuition; the reality is a bit more complex. It turns out that, in order to prove (3.3) and (3.5), the decoding must in fact be *bijective*, but the one sketched above is not injective. For instance, $\mathbf{A}(\text{nat}, \text{nat})$ and nat are obviously distinct types but both decode to nat (by either projection). Fortunately, there is an easy way to obtain the desired bijectivity: we only encode a type pair (τ_1, τ_2) directly as a name $\mathbf{A}(\tau_1, \tau_2)$ if τ_1 and τ_2 are “sufficiently different”. If they share some structure, however, we keep the parts that are the same and only apply \mathbf{A} to the parts that are different. To take a simple example: to encode $(\text{nat}, \text{unit})$, we would use the type name $\mathbf{A}(\text{nat}, \text{unit})$, but to encode $(\text{nat} \rightarrow \text{nat}, \text{nat} \rightarrow \text{unit})$, we would pick $\text{nat} \rightarrow \mathbf{A}(\text{nat}, \text{unit})$ instead of $\mathbf{A}(\text{nat} \rightarrow \text{nat}, \text{nat} \rightarrow \text{unit})$.

Finally, property (3.4) is shown by coinduction. Recalling that \mathbf{E} comprises three cases (**DIV**, **EVAL**, **CALL**), the key here is that the case of \mathbf{E} by which the first and second terms—call them e'_1 and e'_2 —are related should match the case by which the second and third terms— e'_2 and e'_3 —are related. In other words, out of the 3×3 possible cases, 6 should never arise. As a representative example, consider the situation where e'_1 and e'_2 are related because they reduce to related values (case

EVAL), and e'_2 and e'_3 because they reduce to related calls with related continuations (case **CALL**). By Lemma 12 we can assume that these calls are to *external* functions related by $G_{(2)}$. But this means that the function called in e'_2 must be of the form $\mathbf{I}(\tau, v_1, v_3)$, *i.e.*, not a function at all (recall that we required \mathbf{I} to map to *rigid* values)! Hence, e'_2 gets stuck, contradicting the assumption that it reduced to a value.

The remaining possibilities (where the cases of relatedness for (e'_1, e'_2) and (e'_2, e'_3) match) are handled as follows:

Case DIV. Then both e'_1 and e'_3 diverge, so we are done.

Case EVAL. Then e'_1 , e'_2 , and e'_3 reduce to values v_1 , v_2 , and v_3 , such that $\overline{G_{(1)}}$ relates (v_1, v_2) and $\overline{G_{(2)}}$ relates (v_2, v_3) . Thus, by (3.5), \overline{G} relates (v_1, v_3) and we are done.

Case CALL. We know that e'_1 and e'_2 reduce to related function calls in related continuations, and that the same applies to e'_2 and e'_3 . Using Lemma 12 as above, we know that all four function calls are actually stuck. Since, by determinacy, e'_2 cannot get stuck in two different ways, we have a unique function call and a unique continuation in the middle. So, it suffices to show a corresponding transitivity property for \mathbf{U} and for continuations. The one for continuations follows from (3.3) and the coinductive hypothesis. The one for \mathbf{U} follows from (3.5) and injectivity of \mathbf{I} ; when reasoning about type instantiations $(f_i \square)$, we must also make use of our type encoding \mathbf{A} —dually to how we handle **pack** in proving (3.5).

3.9.2.2 The Gory Details

We first formalize the syntactic encoding of types. The previously motivated notion of two types being “sufficiently different” is defined as the negation of *similarity*.

Definition 10. Similarity, written \approx , is defined inductively via the following rules:

$$\begin{array}{c}
\frac{\mathbf{n} \notin \mathcal{N}_{\exists}}{\mathbf{n} \approx \mathbf{n}} \quad \frac{}{\alpha \approx \alpha} \quad \frac{}{\text{unit} \approx \text{unit}} \quad \frac{}{\text{nat} \approx \text{nat}} \quad \frac{\tau \sim \sigma}{\text{ref } \tau \approx \text{ref } \sigma} \\
\\
\frac{\tau \sim \sigma}{\mu\alpha. \tau \approx \mu\alpha. \sigma} \quad \frac{\tau \sim \sigma}{\forall\alpha. \tau \approx \forall\alpha. \sigma} \quad \frac{\tau \sim \sigma}{\exists\alpha. \tau \approx \exists\alpha. \sigma} \\
\\
\frac{\tau \sim \sigma \quad \tau' \sim \sigma'}{\tau \rightarrow \tau' \approx \sigma \rightarrow \sigma'} \quad \frac{\tau \sim \sigma \quad \tau' \sim \sigma'}{\tau \times \tau' \approx \sigma \times \sigma'} \quad \frac{\tau \sim \sigma \quad \tau' \sim \sigma'}{\tau + \tau' \approx \sigma + \sigma'} \\
\\
\frac{\tau, \sigma \in \mathbf{CTy}}{\tau \sim \sigma} \quad \frac{\tau \approx \sigma}{\tau \sim \sigma}
\end{array}$$

Two closed dissimilar types are encoded as a type name from \mathcal{N}_{\exists} using a bijection

$$\mathbf{A} \in \{(\tau_1, \tau_2) \in \mathbf{CTy} \times \mathbf{CTy} \mid \tau_1 \not\approx \tau_2\} \rightarrow \mathcal{N}_{\exists}.$$

The encoding of two arbitrary closed types is a natural lifting of \mathbf{A} . Instead of defining it explicitly, we find it simpler to define the decoding and then state the existence of a corresponding encoding.

Definition 11. We recursively define *decoding projections* $(-)_i \in \mathbf{Ty} \rightarrow \mathbf{Ty}$ (for $i = 1, 2$) as follows. Note that τ is closed iff $\tau_{(i)}$ is closed.

$$\begin{aligned} \mathbf{n}_{(i)} &:= \begin{cases} \tau_i & \text{if } \mathbf{n} = \mathbf{A}(\tau_1, \tau_2) \text{ for some } \tau_1, \tau_2 \\ \mathbf{n} & \text{otherwise, i.e., } \mathbf{n} \notin \mathcal{N}_\exists \end{cases} \\ \alpha_{(i)} &:= \alpha \\ \mathbf{unit}_{(i)} &:= \mathbf{unit} \\ \mathbf{nat}_{(i)} &:= \mathbf{nat} \\ (\mathbf{ref} \tau)_{(i)} &:= \mathbf{ref} \tau_{(i)} \\ (\mu\alpha. \tau)_{(i)} &:= \mu\alpha. \tau_{(i)} \\ (\forall\alpha. \tau)_{(i)} &:= \forall\alpha. \tau_{(i)} \\ (\exists\alpha. \tau)_{(i)} &:= \exists\alpha. \tau_{(i)} \\ (\tau \rightarrow \tau')_{(i)} &:= \tau_{(i)} \rightarrow \tau'_{(i)} \\ (\tau \times \tau')_{(i)} &:= \tau_{(i)} \times \tau'_{(i)} \\ (\tau + \tau')_{(i)} &:= \tau_{(i)} + \tau'_{(i)} \end{aligned}$$

While we necessarily defined similarity and the decoding projections for arbitrary types (\mathbf{Ty}), ultimately we are only interested in closed types (\mathbf{CTy}). The encoding of two closed types is now implicit in the surjectivity part of the next lemma.

Lemma 31. $(\lambda\tau. \langle \tau_{(1)}, \tau_{(2)} \rangle) \in \mathbf{CTy} \rightarrow \mathbf{CTy} \times \mathbf{CTy}$ is bijective.

Proof. Injectivity (generalized to open types) can be easily shown by induction on types using two sub-lemmas:

- (a) $\forall\tau. \tau_{(1)} \sim \tau_{(2)}$, which is shown by straightforward induction on τ .
- (b) $\forall\tau. \tau_{(1)} \approx \tau_{(2)} \iff \tau \notin \mathcal{N}_\exists$, which is shown by case analysis on τ using (a).

Surjectivity is generalized as follows:

$$\forall\tau_1, \tau_2. \tau_1 \sim \tau_2 \implies \exists\tau. \tau_{(1)} = \tau_1 \wedge \tau_{(2)} = \tau_2$$

(Note that if τ_1 and τ_2 are closed, the premise holds trivially.) This property can be proven by induction on τ_1 . In each case we ask if $\tau_1 \approx \tau_2$ holds. If it does, we use the inductive hypothesis; otherwise, the premise implies that τ_1, τ_2 are closed and thus we can pick τ to be $\mathbf{A}(\tau_1, \tau_2)$. \square

Now, let us define the *decomposition* of value relations. This will eventually be used to decompose a global knowledge for the yet-to-be-defined W into one for W_1 and one for W_2 .

Definition 12. Given $R \in \mathbf{VRelF}$, we define $R_{\{i\}} \in \mathbf{VRelF}$ and $R_{(i)}^* \in W_i.\mathbf{T.S} \rightarrow \mathbf{VRelF}$ (for $i = 1, 2$) as follows.

$$\begin{aligned} R_{\{1\}}(\tau_1) &:= \{(v_1, \mathbf{I}(\tau, v_1, v_3)) \mid \exists \tau \in \mathbf{CTyF}_{\backslash \text{ref}}^{\mathcal{N}}. \tau_1 = \tau_{(1)} \wedge (v_1, v_3) \in R(\tau)\} \\ R_{\{2\}}(\tau_2) &:= \{(\mathbf{I}(\tau, v_1, v_3), v_3) \mid \exists \tau \in \mathbf{CTyF}_{\backslash \text{ref}}^{\mathcal{N}}. \tau_2 = \tau_{(2)} \wedge (v_1, v_3) \in R(\tau)\} \\ R_{(i)}^*(s) &:= [W_i.\mathbf{L}(s)]_{(R_{\{i\}})}^* \end{aligned}$$

Here, $\mathbf{CTyF}_{\backslash \text{ref}}^{\mathcal{N}}$ stands for $\mathbf{CTyF} \setminus (\mathcal{N} \cup \{\text{ref } \tau \in \mathbf{CTy}\})$ and \mathbf{I} is an injective function in $\mathbf{CTy} \times \mathbf{CVal} \times \mathbf{CVal} \rightarrow \mathbf{CVal}$. By construction, $R_{(i)}^*$ is the least fixed point of $W_i.\mathbf{L}$ that contains $R_{\{i\}}$ (recall Definition 3).

With the help of this, we can now finally construct W .

$$\begin{aligned} W.\mathbf{T} &:= W_1.\mathbf{T} \times W_2.\mathbf{T} \\ W.\mathbf{L}(s_1, s_2)(R)(\tau) &:= \begin{cases} \overline{R_{(1)}^*(s_1)(\tau_{(1)}) \circ R_{(2)}^*(s_2)(\tau_{(2)})} & \text{if } \tau \in \mathcal{N}_{\exists} \\ W_1.\mathbf{L}(s_1)(R_{(1)}^*(s_1))(\tau_{(1)}) & \text{if } \tau \notin \mathcal{N}_{\exists} \\ \quad \circ W_2.\mathbf{L}(s_2)(R_{(2)}^*(s_2))(\tau_{(2)}) & \end{cases} \\ W.\mathbf{H}(s_1, s_2)(R) &:= W_1.\mathbf{H}(s_1)(R_{(1)}^*(s_1)) \circ W_2.\mathbf{H}(s_2)(R_{(2)}^*(s_2)) \\ W.\mathbf{N} &:= \mathcal{N} \end{aligned}$$

Its well-formedness is easy to check. Note that, although W only actually defines names from \mathcal{N}_{\exists} , we declare that it owns the larger set \mathcal{N} in order to reduce the set of global knowledges that we have to worry about.

Next, some notation for decomposing a global knowledge.

Definition 13. Given $G \in \mathbf{GK}(W)$, we define $G_{(1)}^{s_2} \in W_1.\mathbf{S} \rightarrow \mathbf{VRelF}$ for $s_2 \in W_2.\mathbf{S}$, and $G_{(2)}^{s_1} \in W_2.\mathbf{S} \rightarrow \mathbf{VRelF}$ for $s_1 \in W_1.\mathbf{S}$.

$$G_{(1)}^{s_2}(s) := G(s, s_2)_{(1)}^*(s) \quad G_{(2)}^{s_1}(s) := G(s_1, s)_{(2)}^*(s)$$

Note that if the global knowledge argument R in the definition of $W.\mathbf{L}$ and $W.\mathbf{H}$ is of the form $G(s_1, s_2)$ for $G \in \mathbf{GK}(W)$, then the corresponding argument passed to $W_1.\mathbf{L}$ and $W_1.\mathbf{H}$ is exactly $G_{(1)}^{s_2}(s_1)$ (and similarly for W_2). It remains to show that $G_{(1)}^{s_2}(s_1)$ and $G_{(2)}^{s_1}(s_2)$ are in fact valid global knowledges for W_1 and W_2 , respectively:

Lemma 32. For any $G \in \mathbf{GK}(W)$:

- $\forall s_2. G_{(1)}^{s_2} \in \mathbf{GK}(W_1)$
- $\forall s_1. G_{(2)}^{s_1} \in \mathbf{GK}(W_2)$

Proof. We show the first part (the second is analogous). Monotonicity of $G_{(1)}^{s_2}$ is proven by induction (recall the construction) using the fact that G , $(-)_{\{1\}}$ and $W_1.\mathbf{L}$ are monotone. $G_{(1)}^{s_2}(s_1) \supseteq_{\text{ref}}^{W_1.\mathbf{N}} W_1.\mathbf{L}(s_1)(G_{(1)}^{s_2}(s_1))$ follows from $G_{(1)}^{s_2}(s_1)(\tau) = W_1.\mathbf{L}(s_1)(G_{(1)}^{s_2}(s_1))(\tau) \cup G(s_1, s_2)_{\{1\}}(\tau)$. (Note that $G(s_1, s_2)_{\{1\}}(\tau) = \emptyset$ if τ is a reference type or a name in $W_1.\mathbf{N} = \mathcal{N}_1 \subseteq \mathcal{N}$.) \square

We now come to the main lemma of this part, namely the conjunction of properties (3.3) and (3.5) from Section 3.9.2.1:

Lemma 33. $\forall G \in \mathbf{GK}(W). \forall \tau \in \mathbf{CTy}. \forall s_1, s_2.$

$$\overline{G_{(1)}^{s_2}(s_1)}(\tau_{(1)}) \circ \overline{G_{(2)}^{s_1}(s_2)}(\tau_{(2)}) = \overline{G(s_1, s_2)}(\tau)$$

Proof. The \supseteq part is proven by induction on $\overline{G(s_1, s_2)}$ (recall that $\overline{(-)}$ is defined as a least fixed point). The other part is equivalent to $\overline{G_{(1)}^{s_2}(s_1)}(\tau) \subseteq X$, where X is

$$\{(v_1, v_2) \mid \forall \sigma, v_3. \tau = \sigma_{(1)} \wedge (v_2, v_3) \in \overline{G_{(2)}^{s_1}(s_2)}(\sigma_{(2)}) \implies (v_1, v_3) \in \overline{G(s_1, s_2)}(\sigma)\}.$$

This is proven by (generalized) induction on $\overline{G_{(1)}^{s_2}(s_1)}$.

The base cases of both inductive proofs follow from

$$\overline{G_{(1)}^{s_2}(s_1)}(\tau_{(1)}) \circ \overline{G_{(2)}^{s_1}(s_2)}(\tau_{(2)}) = G(s_1, s_2)(\tau)$$

for *flexible* τ (i.e., $\tau \in \mathbf{CTyF}$), which we now show by case analysis on τ . If $\tau \in \mathcal{N}_\exists$, then

$$\begin{aligned} G(s_1, s_2)(\tau) &= W.L(s_1, s_2)(G(s_1, s_2))(\tau) && (G \in \mathbf{GK}(W)) \\ &= \overline{G_{(1)}^{s_2}(s_1)}(\tau_{(1)}) \circ \overline{G_{(2)}^{s_1}(s_2)}(\tau_{(2)}) && (\text{def. of } W.L) \end{aligned}$$

and we are done. Otherwise ($\tau \in \mathbf{CTyF} \setminus \mathcal{N}_\exists$), we have:

$$\begin{aligned} &\overline{G_{(1)}^{s_2}(s_1)}(\tau_{(1)}) \circ \overline{G_{(2)}^{s_1}(s_2)}(\tau_{(2)}) \\ &= \overline{G_{(1)}^{s_2}(s_1)}(\tau_{(1)}) \circ \overline{G_{(2)}^{s_1}(s_2)}(\tau_{(2)}) \\ &= (W_1.L(s_1)(\overline{G_{(1)}^{s_2}(s_1)}(\tau_{(1)}) \cup G(s_1, s_2)_{\{1\}}(\tau_{(1)})) \circ \\ &\quad (W_2.L(s_2)(\overline{G_{(2)}^{s_1}(s_2)}(\tau_{(2)}) \cup G(s_1, s_2)_{\{2\}}(\tau_{(2)})) \end{aligned}$$

Now, if τ is a reference type or a name from $\mathcal{N}_1 \uplus \mathcal{N}_2$, we finish by rewriting as follows:

$$\begin{aligned} &(W_1.L(s_1)(\overline{G_{(1)}^{s_2}(s_1)}(\tau_{(1)}) \cup G(s_1, s_2)_{\{1\}}(\tau_{(1)})) \circ \\ &\quad (W_2.L(s_2)(\overline{G_{(2)}^{s_1}(s_2)}(\tau_{(2)}) \cup G(s_1, s_2)_{\{2\}}(\tau_{(2)})) \\ &= (W_1.L(s_1)(\overline{G_{(1)}^{s_2}(s_1)}(\tau_{(1)}) \cup \emptyset) \circ \\ &\quad (W_2.L(s_2)(\overline{G_{(2)}^{s_1}(s_2)}(\tau_{(2)}) \cup \emptyset) && (\text{def. of } (-)_{\{i\}}) \\ &= W.L(s_1, s_2)(G(s_1, s_2))(\tau) && (\text{def. of } W.L) \\ &= G(s_1, s_2)(\tau) && (G \in \mathbf{GK}(W)) \end{aligned}$$

Otherwise ($\tau \in \mathbf{CTyF}_{\setminus \text{ref}}^{\mathcal{N}}$), we continue by distributing \circ over \cup :

$$\begin{aligned} &(W_1.L(s_1)(\overline{G_{(1)}^{s_2}(s_1)}(\tau_{(1)}) \cup G(s_1, s_2)_{\{1\}}(\tau_{(1)})) \circ \\ &\quad (W_2.L(s_2)(\overline{G_{(2)}^{s_1}(s_2)}(\tau_{(2)}) \cup G(s_1, s_2)_{\{2\}}(\tau_{(2)})) \\ &= (W_1.L(s_1)(\overline{G_{(1)}^{s_2}(s_1)}(\tau_{(1)}) \circ W_2.L(s_2)(\overline{G_{(2)}^{s_1}(s_2)}(\tau_{(2)})) \cup \\ &\quad (G(s_1, s_2)_{\{1\}}(\tau_{(1)}) \circ G(s_1, s_2)_{\{2\}}(\tau_{(2)})) \cup \\ &\quad (W_1.L(s_1)(\overline{G_{(1)}^{s_2}(s_1)}(\tau_{(1)}) \circ G(s_1, s_2)_{\{2\}}(\tau_{(2)})) \cup \\ &\quad (G(s_1, s_2)_{\{1\}}(\tau_{(1)}) \circ W_2.L(s_2)(\overline{G_{(2)}^{s_1}(s_2)}(\tau_{(2)})) \end{aligned}$$

The first disjunct equals $W.L(s_1, s_2)(G(s_1, s_2))(\tau)$ by construction of W ; the second equals $G(s_1, s_2)(\tau)$ by construction of $(-)\{i\}$, injectivity of \mathbf{I} and the injectivity part of Lemma 31; and the third and the fourth are empty by construction of $(-)\{i\}$: thanks to the FixVal and GenVal restriction (cf. Figure 3.12), no rigid value (such as the \mathbf{I} -encoded values in $G(s_1, s_2)\{i\}$) can be related by a local knowledge. So, the big union above simply becomes $W.L(s_1, s_2)(G(s_1, s_2))(\tau) \cup G(s_1, s_2)(\tau)$. Since $G \in \mathbf{GK}(W)$, the second disjunct already contains the first and we are done.

The inductive cases of both inductive proofs boil down to showing that for any $R_1, R_2, R \in \mathbf{VRelF}$ and $S_1, S_2 \in \mathbf{VRel}$ and $\tau \notin \mathbf{CTyF}$ the equation

$$F_{R_1}(S_1)(\tau_{(1)}) \circ F_{R_2}(S_2)(\tau_{(2)}) = F_R(S_1 \bullet S_2)(\tau)$$

holds, where $F_R \in \mathbf{VRel} \rightarrow \mathbf{VRel}$ denotes the monotone generating function of \bar{R} (i.e., the function of which \bar{R} is the least fixed-point, and which can be easily inferred from Figure 3.11), and $(S_1 \bullet S_2)(\tau) := S_1(\tau_{(1)}) \circ S_2(\tau_{(2)})$. This is straightforward to show by case analysis on τ . The only really interesting case is for existential types, where (in one direction) we are given two witness types τ_1 and τ_2 and then apply Lemma 31 to find a witness type τ satisfying $\tau_{(1)} = \tau_1$ and $\tau_{(2)} = \tau_2$. \square

Finally, we can prove transitive compositionality of \mathbf{E} and then the original goal of this first part.

Lemma 34. $\forall G \in \mathbf{GK}(W). \forall \tau \in \mathbf{CTy}. \forall s_1, s_2.$

$$\mathbf{E}_{W_1}(G_{(1)}^{s_2})(s_1)(\tau_{(1)}) \circ \mathbf{E}_{W_2}(G_{(2)}^{s_1})(s_2)(\tau_{(2)}) \subseteq \mathbf{E}_W(G)(s_1, s_2)(\tau)$$

Proof. By coinduction, following the sketch in Section 3.9.2.1 (and choosing the middle frame heap, h_2^F , to be empty). \square

Lemma 35. $\Delta; \Gamma \vdash e_1 \sim_W e_3 : \tau$

Proof. Inhabitation of W follows easily from that of W_1 and W_2 and the construction of $W.H$. The proofs of consistency and relatedness of (e_1, e_3) by \mathbf{E}_W are very similar and straightforward, using Lemmas 33 and 34. \square

3.9.3 Second Part: Constructing the Corresponding Local World w

We now come to the second and final part of our transitivity proof. Conceptually it is quite simple, but the formal details are very subtle. Recall that we basically want to create a world that relates the same things as W from the previous section, but has the shape of a lifted world, i.e., has the form $w\uparrow$ for some local world w . By definition, the state space of a lifted world is of the form $W_{\text{ref}}.\mathbf{S} \times \dots$, and thus cannot be the same as W 's state space $(W_{\text{ref}}.\mathbf{S} \times w_1.\mathbf{S}) \times (W_{\text{ref}}.\mathbf{S} \times w_2.\mathbf{S})$.

Instead of trying to define w directly and then trying to prove that W and $w\uparrow$ give rise to the same equivalences, we take a more principled approach. We first develop a simple notion of *world isomorphism* in order to characterize what it means for two

worlds with different state spaces to “relate the same things”. Subsequently, we find an isomorphism between W and a lifted local world. Although somewhat tailored to its application in the transitivity proof, the isomorphism theory can be useful in other cases.

3.9.3.1 World Isomorphisms

Roughly, two (full) worlds W_a, W_b are isomorphic iff they declare the same type names, and each state of W_a corresponds to a state of W_b (and vice versa) such that the same values and heaps are related at corresponding states. Different kinds of isomorphism arise depending on what counts as a correspondence. For our purpose, a one-to-one correspondence is too strong. If, say, W_b contains an “inconsistent” state (*i.e.*, a state at which W_b ’s heap relation is empty), then we should not have to worry about finding a similarly irrelevant state in W_a . So, instead of a full one-to-one correspondence, we use a *partial* one, wherein a state s in one world is permitted to have no correspondent in the other if s is inconsistent. This plays a crucial role in our transitivity proof, as we will see in a moment.

Definition 14. For any $W_a, W_b \in \text{World}$, a function $\phi \in W_a.S \rightarrow \mathcal{P}(W_b.S)$ is a *weak morphism* from W_a to W_b iff:

- (1) $\forall s_a, s'_a. \forall s_b \in \phi(s_a). \forall s'_b \in \phi(s'_a). s_a \sqsubseteq s'_a \implies s_b \sqsubseteq s'_b$
- (2) $\forall s_a, s'_a. \forall s_b \in \phi(s_a). \forall s'_b \in \phi(s'_a). s_a \sqsubseteq_{\text{pub}} s'_a \implies s_b \sqsubseteq_{\text{pub}} s'_b$
- (3) $\forall s_a. \forall s_b \in \phi(s_a). W_a.L(s_a) = W_b.L(s_b)$
- (4) $\forall s_a. \forall G \in \text{GK}(W_a). W_a.H(s_a)(G(s_a)) \subseteq \bigcup_{s_b \in \phi(s_a)} W_b.H(s_b)(G(s_a))$
- (5) $W_a.N = W_b.N$

It is easy to see that worlds and weak morphisms form a category.

Although we actually allow a state from one world to be mapped to a whole set of states from the other world, for the purpose of the transitivity proof this set will always be a singleton or empty.

Definition 15. For any $W_a, W_b \in \text{World}$, functions $\phi \in W_a.S \rightarrow \mathcal{P}(W_b.S)$ and $\psi \in W_b.S \rightarrow \mathcal{P}(W_a.S)$ form a *weak isomorphism*, written $\phi : W_a \cong W_b : \psi$, iff they are weak morphisms and satisfy the following:

- $\forall s_a. \forall s_b \in \phi(s_a). \forall s'_a \in \psi(s_b). s_a \sqsubseteq_{\text{pub}} s'_a$
- $\forall s_b. \forall s_a \in \psi(s_b). \forall s'_b \in \phi(s_a). s_b \sqsubseteq_{\text{pub}} s'_b$

Theorem 7 (Weak isomorphisms preserve equivalences). If $\phi : W_a \cong W_b : \psi$, then for any $\Delta, \Gamma, e_1, e_2, \tau$:

$$\Delta; \Gamma \vdash e_1 \sim_{W_a} e_2 : \tau \iff \Delta; \Gamma \vdash e_1 \sim_{W_b} e_2 : \tau$$

Proof. The proof is both somewhat tricky and tedious, and omitted here. At the heart are the following two global knowledge constructions. Recall that for a monotone function $F \in \mathbf{VRelF} \rightarrow \mathbf{VRelF}$ and $R \in \mathbf{VRelF}$, we write $[F]_R^*$ for the least fixed point of the monotone function $F(-) \cup R$.

For $\phi : W_1 \rightarrow W_2$ and $G \in \mathbf{GK}(W_2)$:

$$\begin{aligned} \overleftarrow{G}_\phi &\in \mathbf{GK}(W_1) \\ \overleftarrow{G}_\phi(s_1) &:= [W_1.\mathbf{L}(s_1)]^*_{\bigcup\{G(s'_2) \mid \exists s'_1 \sqsubseteq s_1. s'_2 \in \phi(s'_1)\}} \end{aligned}$$

For $\phi : W_1 \rightarrow W_2$, $G \in \mathbf{GK}(W_1)$, and $s_1 \in W_1.\mathbf{S}$:

$$\begin{aligned} \overrightarrow{G}_\phi^{s_1} &\in \mathbf{GK}(W_2) \\ \overrightarrow{G}_\phi^{s_1}(s_2) &= [W_2.\mathbf{L}(s_2)]^*_{\bigcup\{G(s'_1) \mid s'_1 \sqsubseteq s_1 \wedge \exists s'_2. s'_2 \sqsubseteq s_2 \wedge s'_2 \in \phi(s'_1)\}} \end{aligned}$$

□

Before moving on, we show three simple examples of weak isomorphisms.

Theorem 8 (The *pruning* isomorphism). Given $W \in \mathbf{World}$, we write $\mathit{prune}(W)$ for the world obtained by removing all invalid states from W 's state space, where “invalid” means that no heaps are related at these states. Intuitively, these states simply do not matter, as they can never be meaningfully used in a proof. It is easy to show that W and $\mathit{prune}(W)$ are weakly isomorphic.

Formally, the transformation is defined as follows.

$$\begin{aligned} \mathit{prune}(W).\mathbf{T.S} &:= \{s \in W.\mathbf{T.S} \mid \exists G \in \mathbf{GK}(W). \exists (h_1, h_2) \in W.\mathbf{H}(s)(G(s))\} \\ \mathit{prune}(W).\mathbf{T} \sqsubseteq &:= W.\mathbf{T} \sqsubseteq \\ \mathit{prune}(W).\mathbf{T} \sqsubseteq_{\text{pub}} &:= W.\mathbf{T} \sqsubseteq_{\text{pub}} \\ \mathit{prune}(W).\mathbf{L} &:= W.\mathbf{L} \\ \mathit{prune}(W).\mathbf{H} &:= W.\mathbf{H} \\ \mathit{prune}(W).\mathbf{N} &:= W.\mathbf{N} \end{aligned}$$

We have $\phi : W \cong \mathit{prune}(W) : \psi$, where:

$$\begin{aligned} \phi(s) &:= \{s \mid \exists G \in \mathbf{GK}(W). \exists (h_1, h_2) \in W.\mathbf{H}(s)(G(s))\} \\ \psi(s) &:= \{s\} \end{aligned}$$

(Note that $\phi(s)$ is either empty or a singleton set.)

Theorem 9 (The *flattening* isomorphism). Given $W \in \mathbf{World}$, we write $\mathit{flatten}(W)$ for the world obtained by *flattening* W 's state space, in the sense that at each new state at most one pair of heaps is related. This is done by pairing each state with

two heaps and defining the heap relation to be empty if these heaps are not related at the state. Formally, flattening is defined as follows.

$$\begin{aligned}
\text{flatten}(W).T.S &:= W.T.S \times \text{Heap} \times \text{Heap} \\
\text{flatten}(W).T.\sqsubseteq &:= \{((s', h'_1, h'_2), (s, h_1, h_2)) \mid s' \sqsubseteq s\} \\
\text{flatten}(W).T.\sqsubseteq_{\text{pub}} &:= \{((s', h'_1, h'_2), (s, h_1, h_2)) \mid s' \sqsubseteq_{\text{pub}} s\} \\
\text{flatten}(W).L(s, h_1, h_2) &:= W.L(s) \\
\text{flatten}(W).H(s, h_1, h_2)(R) &:= \{(h_1, h_2)\} \cap W.H(s)(R) \\
\text{flatten}(W).N &:= W.N
\end{aligned}$$

We have $\phi : W \cong \text{flatten}(W) : \psi$, where:

$$\begin{aligned}
\phi(s) &:= \{(s, h_1, h_2) \mid h_1, h_2 \in \text{Heap}\} \\
\psi(s, h_1, h_2) &:= \{s\}
\end{aligned}$$

Theorem 10 (The *swapping* isomorphism). Given $w \in \text{LWorld}$ with $w.S = S_1 \times S_2$, we write $\text{swap}(w)$ for the local world obtained by *swapping* w 's state space. Formally, swapping is defined as follows.

$$\begin{aligned}
\text{swap}(w).T.S &:= S_2 \times S_1 \\
\text{swap}(w).T.\sqsubseteq &:= \{((s'_2, s'_1), (s_2, s_1)) \mid (s'_1, s'_2) \sqsubseteq (s_1, s_2)\} \\
\text{swap}(w).T.\sqsubseteq_{\text{pub}} &:= \{((s'_2, s'_1), (s_2, s_1)) \mid (s'_1, s'_2) \sqsubseteq_{\text{pub}} (s_1, s_2)\} \\
\text{swap}(w).L(s_{\text{ref}}, (s_2, s_1)) &:= w.L(s_{\text{ref}}, (s_1, s_2)) \\
\text{swap}(w).H(s_{\text{ref}}, (s_2, s_1)) &:= w.H(s_{\text{ref}}, (s_1, s_2)) \\
\text{swap}(w).N &:= w.N
\end{aligned}$$

We have $\phi : w \uparrow \cong \text{swap}(w) \uparrow : \psi$, where:

$$\begin{aligned}
\phi(s_{\text{ref}}, (s_1, s_2)) &:= \{(s_{\text{ref}}, (s_2, s_1))\} \\
\psi(s_{\text{ref}}, (s_2, s_1)) &:= \{(s_{\text{ref}}, (s_1, s_2))\}
\end{aligned}$$

Given local worlds $w_1, w_2 \in \text{LWorld}$, note that $\text{swap}(w_1 \otimes w_2) = w_2 \otimes w_1$. Hence we particularly have $\phi : (w_1 \otimes w_2) \uparrow \cong (w_2 \otimes w_1) \uparrow : \psi$. This can be used, for instance, to immediately obtain the two symmetric cases of Lemma 9.

Remark. Our notions of morphism and isomorphism are somewhat simple-minded and there ought to be more general versions. In particular, our morphisms by themselves seem not particularly useful. It is not the case, for instance, that a morphism between two worlds implies one direction of Theorem 7. We have played with the idea of morphisms as *lenses*[15] where a morphism from world W_1 to W_2 would mean that W_2 can be viewed as W_1 , *e.g.*, by ignoring some of W_2 's constraints. Two particular instances of such lenses would be the morphisms from $w_i \uparrow$ to $(w_1 \otimes w_2) \uparrow$, which should then automatically yield Lemma 9. This remains future work, however.

3.9.3.2 Defining w

Recall that lifting a local world means linking it with the shared world W_{ref} , which provides the meaning of reference types. Accordingly, the to-be-constructed local world w 's knowledge must not relate anything at reference types, and, in order for $w\uparrow$ to be isomorphic to W , must correspond to $W.L$ at all other types. This is easy to achieve by just choosing w 's state space to be the same as W 's and then defining $w.L(s_{\text{ref}})(s)$ to be $W.L(s)$ for non-reference types. Regarding reference types, we have to satisfy (by the definition of lifting):

$$W_{\text{ref}}.L(s^{\text{ref}})(R)(\text{ref } \tau) = W.L(s)(R)(\text{ref } \tau)$$

This is problematic. Recall that $W.T = W_1.T \times W_2.T$, so s really has the form $((s_1^{\text{ref}}, s_1^{\text{loc}}), (s_2^{\text{ref}}, s_2^{\text{loc}}))$ (we will later omit the inner parentheses for convenience), with $s_1^{\text{ref}}, s_2^{\text{ref}}$ being states of W_{ref} , and $s_1^{\text{loc}}, s_2^{\text{loc}}$ being states of w_1, w_2 , respectively. Unfolding the definition of $W.L$, the above equation is equivalent to

$$\begin{aligned} W_{\text{ref}}.L(s^{\text{ref}})(R)(\text{ref } \tau) &= W_{\text{ref}}.L(s_1^{\text{ref}})(R_{(1)}^*(s_1^{\text{ref}}, s_1^{\text{loc}}))(\text{ref } \tau_{(1)}) \circ \\ &\quad W_{\text{ref}}.L(s_2^{\text{ref}})(R_{(2)}^*(s_2^{\text{ref}}, s_2^{\text{loc}}))(\text{ref } \tau_{(2)}), \end{aligned}$$

which in turn simplifies to $s^{\text{ref}} = s_1^{\text{ref}} \bullet s_2^{\text{ref}}$, where

$$s_1^{\text{ref}} \bullet s_2^{\text{ref}} := \{(\tau, \ell_1, \ell_3) \mid \exists \ell_2. (\tau_{(1)}, \ell_1, \ell_2) \in s_1^{\text{ref}} \wedge (\tau_{(2)}, \ell_2, \ell_3) \in s_2^{\text{ref}}\}.$$

This clearly cannot be true in general as all three states may be arbitrary. Remember, however, that we do not have to worry about inconsistent states! So the solution is easy: in the states where the equation holds—*i.e.*, where s^{ref} and s are *coherent*—we are fine; we just need to make sure that in any other case $w\uparrow$'s heap relation is empty. And since w 's heap relation may depend on the shared state s_{ref} , this can be easily achieved. Accordingly, the $w\uparrow$ state corresponding to s will be $(s_1^{\text{ref}} \bullet s_2^{\text{ref}}, s)$, and the W state corresponding to (s^{ref}, s) will be s —but only if s^{ref} happens to be $s_1^{\text{ref}} \bullet s_2^{\text{ref}}$.

What should $w.H$ relate when $s^{\text{ref}} = s_1^{\text{ref}} \bullet s_2^{\text{ref}}$ *does* hold? For such states, we want the following equation to be true (ignoring the global knowledge parameter to avoid clutter):

$$w\uparrow.H(s^{\text{ref}}, s) = W.H(s) \tag{3.7}$$

Let us look at what we know about heaps (h_1, h_3) related by $W.H(s)$. First, by construction of W , there is some h_2 mediating between $w_1\uparrow$ and $w_2\uparrow$. By definition of lifting, h_1 and h_2 can be split between $W_{\text{ref}}.H(s_1^{\text{ref}})$ and $w_1.H(s_1^{\text{ref}})(s_1^{\text{loc}})$, and similarly h_2 and h_3 can be split between $W_{\text{ref}}.H(s_2^{\text{ref}})$ and $w_2.H(s_2^{\text{ref}})(s_2^{\text{loc}})$. Of course, in general the two splits of h_2 may be arbitrarily different. This situation is depicted in the first diagram of Figure 3.18.

Note that $w\uparrow.H(s^{\text{ref}}, s)$ in (3.7) unfolds to $W_{\text{ref}}.H(s^{\text{ref}}) \otimes w.H(s^{\text{ref}}, s)$. So basically all we have to do is define $w.H(s^{\text{ref}}, s)$ to be the “septraction” [89] of $W_{\text{ref}}.H(s^{\text{ref}})$ from $W.H(s)$. The way we do this is essentially by describing, in the definition of $w.H$, the situation from the figure but leaving out the pieces related by $W_{\text{ref}}.H(s^{\text{ref}})$.

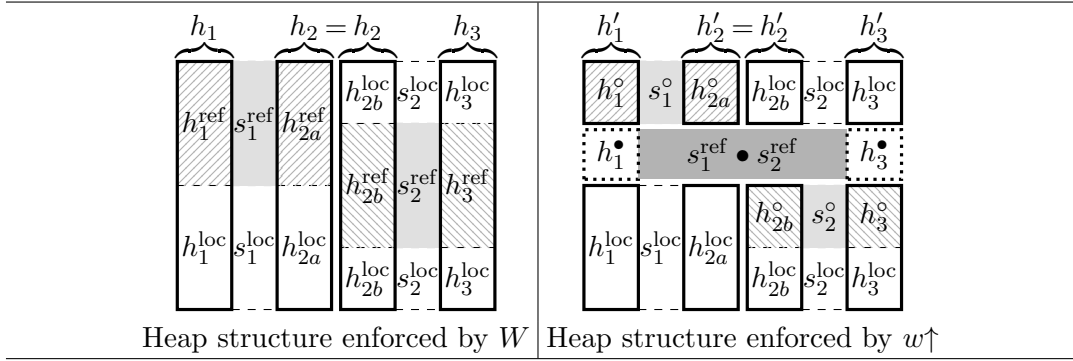


Figure 3.18: Construction of a stable local world $w \in \text{LWorld}$ such that $w \uparrow$ is weakly isomorphic to W .

This is shown in the second diagram of Figure 3.18: $w.H$ relates heaps h'_1, h'_3 iff $h'_i = h_i^{\text{loc}} \sqcup h_i^{\circ} (\neq \perp)$, where h_i° is the sub-heap of h_i^{ref} not covered by $s^{\text{ref}} = s_1^{\text{ref}} \bullet s_2^{\text{ref}}$. The missing pieces, h_1^{\bullet} and h_3^{\bullet} , are then going to be related by $W_{\text{ref}}.H(s^{\text{ref}})$ when w is lifted.

Formally, w is defined as follows.

$$\begin{aligned}
 w.T &:= W.T \\
 w.L(s^{\text{ref}}, s)(R)(\tau) &:= \{(v_1, v_3) \in W.L(s)(R)(\tau) \mid \nexists \tau'. \tau = \text{ref } \tau'\} \\
 w.H(s^{\text{ref}}, s)(R) &:= \{(h'_1, h'_3) \mid \\
 &\quad \exists s_1^{\text{ref}}, s_1^{\text{loc}}, s_2^{\text{ref}}, s_2^{\text{loc}}, h_1^{\circ}, h_1^{\text{loc}}, h_{2a}^{\circ}, h_{2a}^{\text{loc}}, h_{2b}^{\circ}, h_{2b}^{\text{loc}}, h_3^{\circ}, h_3^{\text{loc}}. \\
 &\quad s = (s_1^{\text{ref}}, s_1^{\text{loc}}, s_2^{\text{ref}}, s_2^{\text{loc}}) \wedge s^{\text{ref}} = s_1^{\text{ref}} \bullet s_2^{\text{ref}} \wedge \\
 &\quad h'_1 = h_1^{\circ} \sqcup h_1^{\text{loc}} \wedge h_{2a}^{\circ} \sqcup h_{2a}^{\text{loc}} = h_{2b}^{\circ} \sqcup h_{2b}^{\text{loc}} \wedge h'_3 = h_3^{\circ} \sqcup h_3^{\text{loc}} \wedge \\
 &\quad \text{dom}(h_{2a}^{\text{loc}}) \cap \text{dom}_{[2]}(s_1^{\text{ref}}) = \text{dom}(h_{2b}^{\text{loc}}) \cap \text{dom}_{[1]}(s_2^{\text{ref}}) = \emptyset \wedge \\
 &\quad (h_1^{\circ}, h_{2a}^{\circ}) \in W_{\text{ref}}.H(s_1^{\circ})(R_{(1)}^*(s_1^{\text{ref}}, s_1^{\text{loc}})) \wedge \\
 &\quad (h_{2b}^{\circ}, h_3^{\circ}) \in W_{\text{ref}}.H(s_2^{\circ})(R_{(2)}^*(s_2^{\text{ref}}, s_2^{\text{loc}})) \wedge \\
 &\quad (h_1^{\text{loc}}, h_{2a}^{\text{loc}}) \in w_1.H(s_1^{\text{ref}})(s_1^{\text{loc}})(R_{(1)}^*(s_1^{\text{ref}}, s_1^{\text{loc}})) \wedge \\
 &\quad (h_{2b}^{\text{loc}}, h_3^{\text{loc}}) \in w_2.H(s_2^{\text{ref}})(s_2^{\text{loc}})(R_{(2)}^*(s_2^{\text{ref}}, s_2^{\text{loc}}))\} \\
 w.N &:= W.N
 \end{aligned}$$

As explained, its heap relation $w.H$ is empty whenever s^{ref} is not compatible with s_1^{ref} and s_2^{ref} . The sub-heap h_1° (and similarly h_3°) is characterized by saying that it is related by W_{ref} to a sub-heap h_{2a}° of h_{2a}^{ref} at the state obtained by essentially subtracting those parts from s_1^{ref} that are involved in $s_1^{\text{ref}} \bullet s_2^{\text{ref}}$.

3.9.3.3 Showing w 's Stability

The well-formedness of w (*i.e.*, $w \in \text{LWorld}$) is fairly easy to check, but proving $\text{stable}(w)$ is tricky (because $w.H$'s dependency on s_{ref} is tricky)⁶. Recall that stability

⁶In fact, this proof of transitivity was the whole motivation for allowing a local world's heap relation to depend on the shared state in the first place.

is crucial for soundness, as it ensures that a local world's dependency on the shared state is compatible with any changes to that state.

Preliminary to that proof, we define the conversion of a global knowledge for $w \uparrow$ to one for W .

Lemma 36. If $s_1^{\text{ref}}, s_2^{\text{ref}} \in W_{\text{ref}}.S$, then $s_1^{\text{ref}} \bullet s_2^{\text{ref}} \in W_{\text{ref}}.S$.

Proof. This relies on the injectivity part of Lemma 31. \square

Definition 16. For $G \in \text{GK}(w \uparrow)$, we define $\overleftarrow{G} \in W.S \rightarrow \text{VRelF}$ as follows:

$$\overleftarrow{G}(s_1^{\text{ref}}, s_1^{\text{loc}}, s_2^{\text{ref}}, s_2^{\text{loc}}) := G(s_1^{\text{ref}} \bullet s_2^{\text{ref}}, (s_1^{\text{ref}}, s_1^{\text{loc}}, s_2^{\text{ref}}, s_2^{\text{loc}}))$$

Lemma 37. $\forall G \in \text{GK}(w \uparrow). \overleftarrow{G} \in \text{GK}(W)$

Lemma 38. $\text{stable}(w)$

Proof. Suppose that $G \in \text{GK}(w \uparrow)$, $s = (s_1^{\text{ref}}, s_1^{\text{loc}}, s_2^{\text{ref}}, s_2^{\text{loc}})$, $\hat{s}^{\text{ref}} \sqsupseteq s^{\text{ref}}$, $h'_1 \sqcup \hat{h}_1^\bullet \neq \perp$, $h'_3 \sqcup \hat{h}_3^\bullet \neq \perp$,

$$(h'_1, h'_3) \in w.H(s^{\text{ref}}, s)(G(s^{\text{ref}}, s)) \quad (3.8)$$

$$\text{and } (\hat{h}_1^\bullet, \hat{h}_3^\bullet) \in W_{\text{ref}}.H(\hat{s}^{\text{ref}})(G(\hat{s}^{\text{ref}}, s)). \quad (3.9)$$

Our goal is to find $\hat{s} = (\hat{s}_1^{\text{ref}}, \hat{s}_1^{\text{loc}}, \hat{s}_2^{\text{ref}}, \hat{s}_2^{\text{loc}}) \sqsupseteq s$ such that

$$(h'_1, h'_3) \in w.H(\hat{s}^{\text{ref}}, \hat{s})(G(\hat{s}^{\text{ref}}, \hat{s})).$$

The main idea is to use $\text{stable}(w_i)$ to obtain \hat{s}_i^{loc} (for $i = 1, 2$).

In order to do so, we must first construct suitable states and heaps needed for instantiating those stabilities. From (3.8) we know $s^{\text{ref}} = s_1^{\text{ref}} \bullet s_2^{\text{ref}}$ and that h'_1, h'_3 are structured as depicted in the diagram of Figure 3.19. From $\hat{s}^{\text{ref}} \sqsupseteq s^{\text{ref}}$ we know $\hat{s}^{\text{ref}} = s^{\text{ref}} \uplus s^+$ for some s^+ . Thus by (3.9), $\hat{h}_1^\bullet, \hat{h}_3^\bullet$ can be split as $\hat{h}_i^\bullet = h_i^\bullet \sqcup h_i^+$ such that

$$(h_1^\bullet, h_3^\bullet) \in W_{\text{ref}}.H(s^{\text{ref}})(G(\hat{s}^{\text{ref}}, s)) \quad (3.10)$$

$$\text{and } (h_1^+, h_3^+) \in W_{\text{ref}}.H(s^+)(G(\hat{s}^{\text{ref}}, s)), \quad (3.11)$$

as depicted in the left part of the diagram in Figure 3.19.

We will now “horizontally” decompose s^+ and s^{ref} , as shown in the right part of Figure 3.19. Since $s_i^{\text{ref}} \sqsupseteq s_i^\circ$, there is s_i^\bullet such that $s_i^{\text{ref}} = s_i^\circ \uplus s_i^\bullet$; consequently we have $\text{dom}_{[2]}(s_1^\bullet) = \text{dom}_{[1]}(s_2^\bullet)$ and $s_1^\bullet \bullet s_2^\bullet = s_1^{\text{ref}} \bullet s_2^{\text{ref}} = s^{\text{ref}}$. To decompose s^+ , we choose a set of fresh locations L (of appropriate size) for the middle, *i.e.*, we define s_i^+ such that $s_1^+ \bullet s_2^+ = s^+$ and $\text{dom}_{[2]}(s_1^+) = \text{dom}_{[1]}(s_2^+) = L$. We can then define \hat{s}_i^{ref} (used to instantiate $\text{stable}(w_i)$) as $\hat{s}_i^{\text{ref}} := s_i^\circ \uplus s_i^\bullet \uplus s_i^+$, so we have $\hat{s}_1^{\text{ref}} \bullet \hat{s}_2^{\text{ref}} = s^{\text{ref}} \uplus s^+ = \hat{s}^{\text{ref}}$.

The mediating heaps h_2^+ (with domain L) and h_2^\bullet are constructed as follows. Since $\hat{s}_i^{\text{ref}} \sqsupseteq s_i^{\text{ref}}$, we know

$$\overleftarrow{G}(\hat{s}_1^{\text{ref}}, s_1^{\text{loc}}, \hat{s}_2^{\text{ref}}, s_2^{\text{loc}}) = G(\hat{s}^{\text{ref}}, (\hat{s}_1^{\text{ref}}, s_1^{\text{loc}}, \hat{s}_2^{\text{ref}}, s_2^{\text{loc}})) \sqsupseteq G(\hat{s}^{\text{ref}}, s) \quad (3.12)$$

by monotonicity of G . From (3.10), (3.11), (3.12), monotonicity of $W_{\text{ref}}.\mathbf{H}$, and Lemmas 37 and 33, we can then find mediating values to construct h_2^\bullet and h_2^+ satisfying (for $\star \in \{\bullet, +\}$)

$$(h_1^\star, h_2^\star) \in W_{\text{ref}}.\mathbf{H}(s_1^\star)(\overleftarrow{G}_{(1)}^{(\hat{s}_2^{\text{ref}}, s_2^{\text{loc}})}(\hat{s}_1^{\text{ref}}, s_1^{\text{loc}})) \quad (3.13)$$

$$\text{and } (h_2^\star, h_3^\star) \in W_{\text{ref}}.\mathbf{H}(s_2^\star)(\overleftarrow{G}_{(2)}^{(\hat{s}_1^{\text{ref}}, s_1^{\text{loc}})}(\hat{s}_2^{\text{ref}}, s_2^{\text{loc}})). \quad (3.14)$$

We will now prepare to instantiate $\text{stable}(w_i)$, starting with $\text{stable}(w_1)$. First, observe that $\overleftarrow{G}_{(1)}^{(\hat{s}_2^{\text{ref}}, s_2^{\text{loc}})} \in \mathbf{GK}(w_1\uparrow)$, thanks to Lemmas 37 and 32. Next, by monotonicity of G , we have that (for $\star \in \{\hat{s}_1^{\text{ref}}, s_1^{\text{ref}}\}$)

$$\overleftarrow{G}_{(1)}^{(\hat{s}_2^{\text{ref}}, s_2^{\text{loc}})}(\star, s_1^{\text{loc}}) = \overleftarrow{G}(\star, s_1^{\text{loc}}, \hat{s}_2^{\text{ref}}, s_2^{\text{loc}})_{(1)}^* \supseteq G(s^{\text{ref}}, s)_{(1)}^*$$

and thus (3.8), along with monotonicity of $W_{\text{ref}}.\mathbf{H}$ and $w_1.\mathbf{H}$ and the definition of $w.\mathbf{H}$ in Figure 3.18, gives us:

$$(h_1^\circ, h_{2a}^\circ) \in W_{\text{ref}}.\mathbf{H}(s_1^\circ)(\overleftarrow{G}_{(1)}^{(\hat{s}_2^{\text{ref}}, s_2^{\text{loc}})}(\hat{s}_1^{\text{ref}}, s_1^{\text{loc}})) \quad (3.15)$$

$$\text{and } (h_1^{\text{loc}}, h_{2a}^{\text{loc}}) \in w_1.\mathbf{H}(s_1^{\text{ref}}, s_1^{\text{loc}})(\overleftarrow{G}_{(1)}^{(\hat{s}_2^{\text{ref}}, s_2^{\text{loc}})}(\hat{s}_1^{\text{ref}}, s_1^{\text{loc}})). \quad (3.16)$$

Thus, by (3.13), (3.15), (3.16), and $\hat{s}_1^{\text{ref}} \supseteq s_1^{\text{ref}}$, we can instantiate $\text{stable}(w_1)$, yielding $\hat{s}_1^{\text{loc}} \supseteq s_1^{\text{loc}}$ such that

$$(h_1^{\text{loc}}, h_{2a}^{\text{loc}}) \in w_1.\mathbf{H}(\hat{s}_1^{\text{ref}}, \hat{s}_1^{\text{loc}})(\overleftarrow{G}_{(1)}^{(\hat{s}_2^{\text{ref}}, s_2^{\text{loc}})}(\hat{s}_1^{\text{ref}}, \hat{s}_1^{\text{loc}})). \quad (3.17)$$

In a similar manner, $\text{stable}(w_2)$ yields $\hat{s}_2^{\text{loc}} \supseteq s_2^{\text{loc}}$ such that

$$(h_{2b}^{\text{loc}}, h_3^{\text{loc}}) \in w_2.\mathbf{H}(\hat{s}_2^{\text{ref}}, \hat{s}_2^{\text{loc}})(\overleftarrow{G}_{(2)}^{(\hat{s}_1^{\text{ref}}, s_1^{\text{loc}})}(\hat{s}_2^{\text{ref}}, \hat{s}_2^{\text{loc}})). \quad (3.18)$$

Finally, let $\hat{s} := (\hat{s}_1^{\text{ref}}, \hat{s}_1^{\text{loc}}, \hat{s}_2^{\text{ref}}, \hat{s}_2^{\text{loc}})$. By monotonicity of G , $W_{\text{ref}}.\mathbf{H}$, and $w_i.\mathbf{H}$, and by definition of $w.\mathbf{H}$, we get

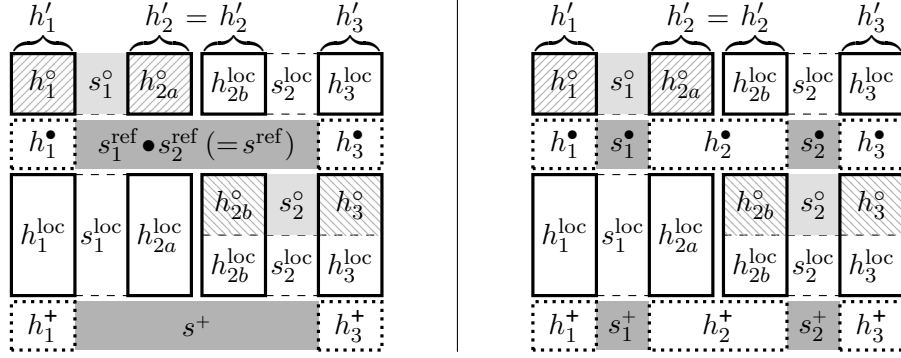
$$(h_1', h_3') \in w.\mathbf{H}(\hat{s}^{\text{ref}}, \hat{s})(\overleftarrow{G}(\hat{s})) = w.\mathbf{H}(\hat{s}^{\text{ref}}, \hat{s})(G(\hat{s}^{\text{ref}}, \hat{s}))$$

from (3.8), (3.17), and (3.18), as desired. \square

3.9.3.4 Proving W and $w\uparrow$ Isomorphic

We now show that W and $w\uparrow$ are weakly isomorphic. Then we can put all the pieces together, arrive at our goal and thereby finishing the proof of transitivity.

Lemma 39. $\exists \phi, \psi. \phi : W \cong w\uparrow : \psi$

Figure 3.19: “Horizontal” decomposition of \hat{s}^{ref} in proof of $\text{stable}(w)$.

Proof. We define $\phi \in W.S \rightarrow \mathcal{P}(w\uparrow.S)$ and $\psi \in w\uparrow.S \rightarrow \mathcal{P}(W.S)$ as discussed earlier.

$$\begin{aligned} \phi(s_1^{\text{ref}}, s_1^{\text{loc}}, s_2^{\text{ref}}, s_2^{\text{loc}}) &:= \{(s_1^{\text{ref}} \bullet s_2^{\text{ref}}, (s_1^{\text{ref}}, s_1^{\text{loc}}, s_2^{\text{ref}}, s_2^{\text{loc}}))\} \\ \psi(s^{\text{ref}}, (s_1^{\text{ref}}, s_1^{\text{loc}}, s_2^{\text{ref}}, s_2^{\text{loc}})) &:= \{(s_1^{\text{ref}}, s_1^{\text{loc}}, s_2^{\text{ref}}, s_2^{\text{loc}}) \mid s^{\text{ref}} = s_1^{\text{ref}} \bullet s_2^{\text{ref}}\} \end{aligned}$$

Showing that ϕ and ψ form a weak isomorphism is mostly straightforward, but for condition (4) of each morphism quite tedious. The key idea behind the proofs of these is the same as that behind Lemma 38: splitting the given heaps as depicted in the diagrams of Figure 3.18. In particular, the proof of morphism condition (4) for ψ is very similar to that of Lemma 38 in the way it uses the construction \overleftarrow{G} and Lemma 37. \square

Theorem 11 (Transitivity). $\Delta; \Gamma \vdash e_1 \sim e_3 : \tau$

Proof. We have $\Delta; \Gamma \vdash e_1 \sim_{w\uparrow} e_3 : \tau$ by Lemmas 35 and 39 and Theorem 7. The result then follows from Lemma 38. \square

3.10 Stuttering Parametric Bisimulations

PBs, as described earlier, suffer from a limitation: they fail to validate the *eta law* for function values,

$$f : (\sigma \rightarrow \tau) \vdash f \sim (\lambda x. f x) : \sigma \rightarrow \tau,$$

as well as more complex equivalences (e.g., the *syntactic minimal invariance* property [12]) which depend on it.

In this section, we sketch *stuttering parametric bisimulations* (SPBs), a variant of PBs that overcomes this limitation. The full details are beyond the scope of this thesis and can be found in a technical report [35].

The key behind SPBs is the use of a *logical reduction semantics* that permits finite but unbounded stuttering steps in between actual “physical” steps. Such stuttering steps effectively enable proofs of equivalence of programs to engage in a game of

“hot potato”, whereby the burden of proof may be tossed back and forth between different parts of the programs, until *eventually* some part makes a physical step of computation. This technique is inspired by the key idea in *well-founded* [58] and *stuttering* [17] bisimulations, adapted in the cited report for the first time to reasoning about open, higher-order programs.

Interestingly, this mechanism also helps overcome a second limitation, which appears at first glance to be completely unrelated: PBs bake in the assumption that the language in question has a “uniform” reduction semantics, *i.e.*, that the reduction relation is parametric in the evaluation context. This assumption is of course not valid in the presence of control flow constructs like `callcc`, so one may naturally wonder how to adapt PBs to a language supporting such features.

3.10.1 The Problem with Eta

Let us examine the inherent problem with eta in the PB model, for simplicity in the setting of λ^μ (Section 3.5). The eta law for an arbitrary function type $\tau' \rightarrow \tau$ corresponds to the following equivalence:

$$f : (\tau' \rightarrow \tau) \vdash f \sim (\lambda x. f\ x) : \tau' \rightarrow \tau$$

This equivalence does *not* hold. Here we prove as much for the case of $\tau' = \tau = \text{nat}$.

Proof. By definition the equivalence holds iff there exists a consistent local knowledge L such that for any global knowledge $G \in \text{GK}(L)$ and any related values $(v_1, v_2) \in G(\text{nat} \rightarrow \text{nat})$ we have $(v_1, \lambda x. v_2\ x) \in \mathbf{E}(G)(\text{nat} \rightarrow \text{nat})$. Being values of function type, such v_1 and $\lambda x. v_2\ x$ are related by $\mathbf{E}(G)$ iff they are related by G . Thus, in order to disprove the eta law, it suffices to construct a “bad” global knowledge $G^\sharp \in \text{GK}(L)$ that relates v_1 and v_2 but does not relate v_1 and $\lambda x. v_2\ x$.

This is an easy task. Let G^\sharp be the least global knowledge that subsumes L and also relates $\lambda y. 0$ and $\lambda y. 1$ at $\text{nat} \rightarrow \text{nat}$. Then we have:

$$G^\sharp(\text{nat} \rightarrow \text{nat}) = L(G^\sharp)(\text{nat} \rightarrow \text{nat}) \cup \{(\lambda y. 0, \lambda y. 1)\}$$

It remains to prove that G^\sharp does not relate $\lambda y. 0$ and $\lambda x. (\lambda y. 1)\ x$. Arguing by contradiction, assume it does. Then from $\text{consistent}(L)$ and, say, $(42, 42) \in \overline{G^\sharp}(\text{nat})$ we know $(\text{step}((\lambda y. 0)\ 42), \text{step}((\lambda x. (\lambda y. 1)\ x)\ 42)) \in \mathbf{E}(G^\sharp)(\text{nat})$, *i.e.*, $(0, (\lambda y. 1)\ 42) \in \mathbf{E}(G^\sharp)(\text{nat})$. Since 0 is a value, this can only mean that $(\lambda y. 1)\ 42$ reduces to a related value, *i.e.*, $(0, 1) \in \overline{G^\sharp}(\text{nat})$, which by definition of G^\sharp is false. \square

To understand better what is going on here, let us now try to *prove* the eta law and see what goes wrong. As is evident from the reasoning at the beginning of the above disproof, we would have to construct a consistent local knowledge L such that any G subsuming it relates v_1 and $\lambda x. v_2\ x$ whenever it relates v_1 and v_2 . Since the only leverage we have over G is what we put in L , the only way to force G to relate certain things is to choose L so that it relates them. Luckily, our definition of L

may depend on G as a parameter, so in order to obtain the aforementioned closure property, we can attempt to define the local knowledge such that:

$$L_\eta(R)(\tau' \rightarrow \tau) = \{(v_1, \lambda x. v_2 x) \mid (v_1, v_2) \in R(\tau' \rightarrow \tau)\}$$

Intuitively, this local knowledge corresponds exactly to what we want to claim: if our context provides us with values v_1 and v_2 that are equivalent at $\tau' \rightarrow \tau$, then we are prepared to claim that v_1 and $\lambda x. v_2 x$ are equivalent at the very same type. Unfortunately, this L_η may be inconsistent! Specifically, suppose we are given $G \in \text{GK}(L_\eta)$ and related arguments $(v'_1, v'_2) \in \overline{G}(\tau')$. For any $(v_1, v_2) \in G(\tau' \rightarrow \tau)$, we must show $(\text{step}(v_1 v'_1), \text{step}((\lambda x. v_2 x) v'_2)) \in \mathbf{E}(G)(\tau)$. The trouble is that, while we know that $\text{step}((\lambda x. v_2 x) v'_2) = v_2 v'_2$, we have no idea what $\text{step}(v_1 v'_1)$ is.⁷

The problem here essentially is that the global knowledge G is under no obligation to be sound w.r.t. contextual equivalence. As a result, if we define a local knowledge like L_η that “re-exports” function values (like v_1) obtained from G , there is no way to know whether applications of such values reduce to well-behaved terms.

3.10.2 Guardedness Revisited

Recall that the requirement of “taking a step” in the definition of PB’s consistency is crucial in ensuring soundness because it guarantees that coinductive reasoning is suitably *guarded*. As the failed proof attempt above shows, however, the guardedness condition appears to be a little too strict. Note that if we were not forced to take that step, then we could easily finish the proof of $\text{consistent}(L_\eta)$ by appealing to $(v_1 v'_1, v_2 v'_2) \in \mathbf{E}(G)(\tau)$, which follows from $(v_1, v_2) \in G(\tau' \rightarrow \tau)$ and $(v'_1, v'_2) \in \overline{G}(\tau')$ (both given), and $(\bullet, \bullet) \in \mathbf{K}(G)(\tau)(\tau)$, using the third disjunct of \mathbf{E} .

Of course, we cannot simply drop the stepping requirement, since this would immediately result in unsoundness—we must have some way of ensuring “productivity” of proofs. What we want to do then, in order to obtain a model that validates the eta law, is to find a slightly weaker guardedness condition (leading to a weaker notion of consistency) that enables the sketched proof of the eta law to go through but is nevertheless strong enough to guarantee soundness of the model.

We achieve this by generalizing the physical notion of “taking a step” to a *logical* one.

3.10.3 Logical Reduction and the Stutter Budget

The idea is very simple. We introduce into the model what we call a *stutter budget* (or just *budget*, for short): two natural numbers, one for each program, that specify how many times one may “stutter”—*i.e.*, avoid taking a reduction step (thus seemingly making no progress)—before eventually taking a step. More precisely, a local knowledge will contain items of the form (n_1, v_1, n_2, v_2) rather than just (v_1, v_2) . When proving consistency for such an item, *i.e.*, when showing that the applications

⁷In fact, it may not even exist: for instance, if v_1 is the integer 5, then $\text{step}(v_1 v'_1)$ is undefined.

of (v_1, v_2) to related arguments (v'_1, v'_2) are related, one then has to make a choice for each application before continuing the reasoning: either one reduces it by one physical step (as before), *or* one leaves it untouched but decreases the corresponding budget instead (n_1 for the application of v_1 , and n_2 for the application of v_2). When one chooses the latter option, one temporarily shirks one's responsibility to make physical progress, passing the proof burden—or “hot potato” as we called it earlier—to the subgoal of showing that the applications are in the **E** relation. Using the third disjunct (the **CALL** case), the **E** relation may then do the same thing and pass the hot potato back to the local knowledge. The important thing is that, each time around this seemingly circular proof path, the respective stutter budgets (n_1 and/or n_2) must be decreased, so we know the hot potato game cannot go on forever.

The way we formulate this is that one *is* actually required to perform a reduction step on both sides, as before, but only a *logical* one. This logical reduction relation, operating on an expression and its budget, is defined as follows.

Definition 17 (Logical Reduction).

$$\frac{e \hookrightarrow e'}{n, e \hookrightarrow n', e'} \quad \frac{n' < n}{n, e \hookrightarrow n', e}$$

That is, a logical step is either a physical step, in which case one may pick an arbitrary budget n' to continue with, or a stutter step, in which case the budget must be decreased.

To get an intuition for why the proposed change to the model is sound, first observe that, since the stutter budget is finite, progress (in the form of a physical step) will eventually be made. Second, note that logical (non-)termination coincides with physical (non-)termination. Thus, logical reduction gives us more flexibility in terms of local reasoning about v_1 and v_2 , and this added flexibility is perfectly sound in that it will not enable us to equate terminating and divergent programs.

3.10.4 Eta Revisited

Using the refined model, we can now indeed prove the eta law along the lines of the earlier attempt by picking the following local knowledge:

$$L_\eta(R)(\tau) := \{(n'_1, v_1, n'_2, \lambda x. v_2 x) \mid \exists \tau', \tau''. \tau = \tau' \rightarrow \tau'' \wedge \exists n_1, n_2. n_1 < n'_1 \wedge (n_1, v_1, n_2, v_2) \in R(\tau)\}$$

Like before, we relate v_1 and $\lambda x. v_2 x$ whenever v_1 and v_2 are related by R . But note the budgets: we choose n'_1 , the budget at which we “export” v_1 to be *larger* than n_1 , the budget at which we import it. (Intuitively, $n_1 := n_1 + 1$ would suffice, but local knowledges in SPBs must be upwards-closed in their budgets.) The budget for $\lambda x. v_2 x$, on the other hand, can be arbitrary. It doesn't matter because an application of that function can always take a physical reduction step (beta reduction).

3.10.5 First-Class Continuations

It turns out that SPBs also enable reasoning about first-class continuations, *e.g.*, in the presence of `callcc` (recall that λ^μ and $F^{\mu l}$ do not feature control effects).

Basically, in such a setting, one wants the local and global knowledge to also relate continuations (evaluation contexts). Continuations can be thought of as functions, so it makes sense to treat them similarly—in particular to have an analogous consistency condition for the continuations related by the local knowledge. The issue then is that one frequently needs to show consistency of continuations that are constructed on top of continuations related by the global knowledge. A representative example is the following:

$$L(R)(\text{nat}) := \{(K_1[\text{inl } \bullet], K_2[\text{inl } \bullet]) \mid (K_1, K_2) \in R(\text{nat} + \tau)\}$$

Here, we import continuations K_1 and K_2 , and export two continuations that are defined in terms of them. To show consistency, we are given related values v_1 and v_2 and must establish that $(\text{step}(K_1[\text{inl } v_1]), \text{step}(K_2[\text{inl } v_2]))$ are related by **E**. Since the injection of a value into a sum type does not require any computation (*i.e.*, $\text{inl } v_i$ does not take a step because it is itself a value), we are stuck.

This problem is obviously very similar to the one that we faced with `eta`. It should then no longer be a surprise that SPBs solve it as well. Using the stutter budget, we can define a local knowledge like the one above as follows:

$$L(R)(\text{nat}) := \{(n'_1, K_1[\text{inl } \bullet], n'_2, K_2[\text{inl } \bullet]) \mid \exists n_1 < n'_1, n_2 < n'_2. (n_1, K_1, n_2, K_2) \in R(\text{nat} + \tau)\}$$

The consistency of this knowledge is easy to show using logical reduction for both programs.

3.10.6 Comparison to Step-Indexing

There is a superficial sense in which SPB proofs may seem similar to step-indexed KLR (SKLR) proofs, namely that they both involve a certain element of “step-counting”. The key difference is *which* steps they are counting: in SKLR proofs one counts physical reduction steps (because the model is built by induction on such steps), whereas in SPB proofs one only counts the stuttering steps (to ensure the absence of infinite stuttering). In essence, SPBs are “non-step”-indexed relations!

Consequently, with SKLRs this step-counting is essential—one cannot ignore physical steps—whereas with SPBs it is often unnecessary. Indeed, when using SPBs to reason about the many examples that PBs already handle, one may usually ignore the extra flexibility provided by SPBs’ logical reduction, avoiding step-counting by restricting one’s proofs to only rely on physical reduction.

Finally, unlike step-indexing, SPB’s stutter budgets cause no problems with transitivity. The technical report [35] contains a detailed proof of transitivity for SPBs. It largely follows the one for PBs but contains some interesting twists.

3.11 Greatest Local Knowledge

Writing down a suitable local knowledge at the beginning of a proof can be quite tedious for complex equivalences (even more so in SPBs). While not really an issue in paper proofs, this quickly becomes very tiresome in formal proofs such as these in our Coq formalization. Fortunately, we can employ *parameterized coinduction* [34, 56] to avoid this issue completely and instead write proofs in an incremental style, where we basically start with a knowledge containing just the functions in question, and extend it (in an implicit fashion) as the proof evolves. Indeed, this is how we prove a big part of soundness in Coq.

In order to get there, we need to express the property of a local knowledge being consistent as that local knowledge being a postfix point of some monotone function. Then the greatest fixed point of that function is automatically the greatest consistent local knowledge, and so we can use the incremental reasoning principle from parameterized coinduction to do proofs about it. For simplicity, here we do this in the setting of λ^μ (Section 3.5). The generalization to $F^{\mu!}$ is straightforward.

Definition 18. We define the wanted function $f \in \text{LK} \xrightarrow{\text{mon}} \text{LK}$.

$$\begin{aligned} f(L)(R)(\tau) := \{ (v_1, v_2) \mid \forall G \in \text{GK}(L). G \supseteq R \implies \\ \forall \tau'. \forall (e_1, e_2) \in \mathbf{U}(\{\tau \mapsto (v_1, v_2)\}, G)(\tau'). \\ (step(e_1), step(e_2)) \in \mathbf{E}(G)(\tau') \} \end{aligned}$$

Here, $\{\tau \mapsto (v_1, v_2)\}$ is short for the typed value relation $\lambda\tau'. \{(v_1, v_2) \mid \tau' = \tau\}$ (containing a single element).

It is easy to verify that f is well-defined, *i.e.*, that it always returns a valid local knowledge and is itself monotone. The reason for quantifying over larger G is to ensure monotonicity of the returned local knowledge.

Lemma 40. $L \subseteq f(L) \iff \text{consistent}(L)$

Definition 19 (Parameterized greatest consistent local knowledge).

$$\begin{aligned} \mathfrak{L} \in \text{LK} &\rightarrow \text{LK} \\ \mathfrak{L}(L) &:= \nu X. f(X \cup L) \end{aligned}$$

Corollary 1 (Greatest consistent local knowledge).

1. $\text{consistent}(\mathfrak{L}(\emptyset))$
2. $\forall L \in \text{LK}. \text{consistent}(L) \implies L \subseteq \mathfrak{L}(\emptyset)$

Theorem 12. $\Gamma \vdash e_1 \sim e_2 : \tau \iff \Gamma \vdash e_1 \sim_{\mathfrak{L}(\emptyset)} e_2 : \tau$

Proof. The “if” direction is trivial due to the first part of Corollary 1.

Regarding the “only if” direction: From $\Gamma \vdash e_1 \sim e_2 : \tau$ we know $\Gamma \vdash e_1 \sim_L e_2 : \tau$ with $\text{consistent}(L)$. Hence $L \subseteq \mathfrak{L}(\emptyset)$ by the second part of Corollary 1. It is easy to check that $\Gamma \vdash e_1 \sim_L e_2 : \tau$ implies $\Gamma \vdash e_1 \sim_{L'} e_2 : \tau$ for any $L' \supseteq L$, so in particular for $\mathfrak{L}(\emptyset)$. \square

Thanks to the construction of \mathfrak{L} as a greatest parameterized fixed point, we have the expected incremental reasoning principle.

Lemma 41. $L \subseteq \mathfrak{L}(L') \iff L \subseteq \mathfrak{L}(L \cup L')$

Finally, a few words about the generalisation to the $F^{\mu!}$ setting, where we have not only local knowledges but whole worlds. There we define an operation that *completes* a world by overwriting its local knowledge component with the greatest consistent local knowledge. Consequently, one constructs a world as usual (defines the STS and so on), but can leave its local knowledge empty or include only the initial functions. Then one applies the completion operator to the world and works with the resulting world instead. There is one point to note, though: since there is no consistency condition on values related at type names \mathbf{n} , completion does not touch these parts of the local knowledge. Hence the initial local knowledge must already contain any such needed values.

3.12 Comparison To Logical Relations

Let us conclude this chapter with high-level comparison of parametric (bi-)simulations and logical relations.

As we have seen, parametric bisimulations emerge from prior work on KLRs [4, 22] in an attempt to overcome its limitations concerning transitivity. In fact, both PBs and KLRs support the same high-level reasoning principles for higher-order imperative programs; they just do so in technically different ways.

Let us consider the two key principles concerning higher-order functions:

Principle 1 (Showing Behavioral Equivalence): To show that two higher-order functions v_a and v_b **behave** equivalently, it suffices to show that their applications **behave** equivalently when the arguments f_a and f_b passed in are **assumed** equivalent.

Principle 2 (Using Assumed Equivalence): If f_a and f_b are **assumed** equivalent, then $f_a \langle \rangle$ and $f_b \langle \rangle$ **behave** equivalently.

So what does it mean to “behave equivalently”? This is really the big question, for which KLRs and PBs give different answers. Rather than try to answer it directly, we will instead describe two informal proof principles concerning behavioral equivalence that, at the level of abstraction we are working at here, are supported by both proof methods, and we will finish the proof sketch by just appealing to these proof principles. After that, we will explain how the different proof methods implement these principles.

Note that these proof principles make a distinction between when two functions *behave* equivalently and when they are *assumed* equivalent. We explain the difference between these notions below. Note also that we restricted Principle 2 here to functions with unit argument; this is merely to simplify our informal discussion.

Logical relations. Both KLRs and PBs allow one to turn the above proof sketch into a proper proof. The key difference between them is how they formalize behavioral vs. assumed equivalence.

KLRs formalize this by defining a relation, which says—once and for all—what it means for two expressions to be indistinguishable at a certain type. One then uses this *same* “logical” relation as the definition of *both* behavioral and assumed equivalence. For expressions, the logical relation says that they are equivalent if they either both run forever or they both evaluate to equivalent values. For function values, which both the v ’s and the f ’s in our example are, the logical relation says that they are equivalent if they map logically-related arguments to logically-related results. Principles 1 and 2 both fall out of this definition as immediate consequences.

The main difficulty with logical relations is that, by conflating behavioral and assumed equivalence, they introduce an inherent circularity in the construction of the logical relation. In particular, the definition of equivalence of function values refers recursively to itself in a negative position (when quantifying over equivalent arguments). Traditionally, for simpler languages (*e.g.*, without recursive types or higher-order state), this circularity is handled by defining the logical relation by induction on the type structure. For richer languages, such as our source language \mathcal{S} , induction on types is no longer sufficient, but step-indexing can be used instead to stratify the construction by the number of steps of computation in the programs being related [4]. This is the approach taken by Hur and Dreyer [33] in their earlier work on compositional compiler correctness. However, it is not known how to prove *transitivity* of logical relations for step-indexed models (at least in a way that is capable of scaling to handle inter-language reasoning, which we need for compiler verification).

Parametric bisimulations. This problem with transitivity was one of the key motivations for *parametric bisimulations* (PBs). Unlike logical relations, PBs treat behavioral and assumed equivalence as distinct concepts. In particular, rather than trying to *define* assumed equivalence, PBs take assumed equivalence as a *parameter* of the model (hence the name “parametric bisimulations”). That is, a PB proof that two expressions are behaviorally equivalent is parameterized by an arbitrary *unknown* relation G (the global knowledge) representing assumed equivalence, and G could relate *any* functions f_a and f_b !

To make use of this unknown G parameter, PBs update the definition of behavioral equivalence accordingly. For function values, one can show them behaviorally equivalent precisely as suggested by Principle 1, *i.e.*, if they map G -related (assumed equivalent) arguments to behaviorally equivalent results. For expressions, behavioral equivalence extends the definition from logical relations with a new possibility, namely that the expressions may call functions f_a and f_b related by G . This amounts to baking Principle 2 directly into the definition of behavioral equivalence. The reason this is necessary—*i.e.*, the reason Principle 2 does not just fall out of the definition otherwise—is that G is a *parameter* of behavioral equivalence. Knowing that f_a and

f_b are assumed equivalent according to G tells us absolutely nothing about them! Consequently, Principle 2 must be explicitly added to the definition of behavioral equivalence as an extra case (the **CALL** disjunct).

One can understand PBs as defining a “local” notion of behavioral equivalence: two expressions are behaviorally equivalent if they behave the same in their local computations, *ignoring* what happens during calls to (G -related) external functions passed in from the environment. Intuitively, this is perfectly sound: it just means each module in a program is responsible for its own local computations, not the local computations of other modules. Moreover, as we have observed already, it is largely a technical detail: PBs can support the same high-level protocol-based reasoning as KLRs do.

The major benefit of PBs over KLRs is that they avoid the problems with the circularity of KLRs which necessitated step-indexing. In particular, note that by taking Principle 1 as the *definition* of behavioral equivalence for function values, we avoid the negative self-reference that plagues logical relations: the arguments f_a and f_b are simply drawn from the global knowledge G . This avoidance of step-indexing was essential in making it possible to establish that PBs do in fact support transitivity, but the proof of transitivity was far from easy.

Chapter 4

Parametric Inter-Language Simulations

Contents

4.1	Overview	98
4.1.1	Transitivity	100
4.2	Languages	100
4.2.1	Language-Generic Approach	100
4.2.2	Language Specification	101
4.2.3	Source Language \mathcal{S}	104
4.2.4	Intermediate Language \mathcal{I}	108
4.2.5	Target Language \mathcal{T}	113
4.3	Worlds	121
4.3.1	Queries	123
4.3.2	Value Closure	125
4.3.3	Lifting and Separating Conjunction of Local Worlds	126
4.4	Concrete Global Worlds	128
4.4.1	Global References	128
4.4.2	Unary Parts	129
4.5	Simulations	135
4.5.1	Stuttering according to algebraic well-founded orders.	141
4.5.2	The two modes of \mathbf{E} and cfg .	143
4.5.3	A Note on the Untyped Model	145
4.5.4	Convenience Lemmas	145
4.6	Example	150
4.6.1	Modules	150
4.6.2	Proof	151
4.7	Metatheory	156

4.7.1	Basics	156
4.7.2	Adequacy	157
4.7.3	Modularity	159
4.8	Proof of Transitivity	165
4.8.1	Overview	165
4.8.2	Constructing the Local World	166
4.8.3	Discussion	175

4.1 Overview

PILS generalize PBs to the inter-language setting of compiler verification. Concretely we are interested in compiling from an ML-like source language \mathcal{S} to a target machine language \mathcal{T} . In this section, we give an overview over the results of this chapter and the subsequent chapter.

The main component of our development is a relation between target modules and source modules: $\Gamma \vdash M_{\mathcal{T}} \lesssim_{\mathcal{T}\mathcal{S}} M_{\mathcal{S}} : \Gamma'$, to be defined later, intuitively states that target module $M_{\mathcal{T}}$ refines source module $M_{\mathcal{S}}$ and that they import the functions listed in Γ and export those in Γ' . The first key result, Theorem 13, applies to whole programs, *i.e.*, well-typed modules that import nothing and export at least a main function (F_{main}) of appropriate type. It states that our relation implies the standard behavioral refinement:

Theorem 13 (Adequacy for whole programs).

$$\frac{\epsilon \vdash M_{\mathcal{T}} \lesssim_{\mathcal{T}\mathcal{S}} M_{\mathcal{S}} : \Gamma \quad (F_{\text{main}} : \text{unit} \rightarrow \tau) \in \Gamma}{\text{Behav}(M_{\mathcal{T}}) \subseteq \text{Behav}(M_{\mathcal{S}})}$$

$\text{Behav}(-)$ denotes the set of I/O and termination behaviors that a program can have. The theorem implies for instance that, if $M_{\mathcal{S}}$ always successfully terminates, then so does $M_{\mathcal{T}}$ and moreover they produce the same outputs.

If we have a compiler that respects our relation $\lesssim_{\mathcal{T}\mathcal{S}}$, then Theorem 13 gives us the same result as traditional whole-program compiler verification would. However, our relation also satisfies the following crucial property (omitting a few well-formedness side conditions—see Section 4.7 for the full statement and its proof).

Theorem 14 (Preservation under linking, a.k.a. modularity).

$$\frac{\Gamma \vdash M_{\mathcal{T}}^1 \lesssim_{\mathcal{T}\mathcal{S}} M_{\mathcal{S}}^1 : \Gamma_1 \quad \Gamma, \Gamma_1 \vdash M_{\mathcal{T}}^2 \lesssim_{\mathcal{T}\mathcal{S}} M_{\mathcal{S}}^2 : \Gamma_2}{\Gamma \vdash \text{link}(M_{\mathcal{T}}^1, M_{\mathcal{T}}^2) \lesssim_{\mathcal{T}\mathcal{S}} \text{link}(M_{\mathcal{S}}^1, M_{\mathcal{S}}^2) : \Gamma_1, \Gamma_2}$$

The theorem says that if we link two target modules, each of which is related to a source module, then the resulting target module is related to the linking of those source modules. Notice that, for the linking to make sense, the types of the first module's exported functions (in Γ_1) need to match the second module's assumptions.

Of course, if a program consists of more than two modules, this theorem can be iterated as necessary and once linking results in a whole program, we can apply Theorem 13.

Observe that these properties don't mention any particular compiler but are stated in terms of arbitrary related modules. The missing link is a theorem saying that the desired compilers adhere to our relation. We prove this for *Pilsner* and *Zwikel*, our compilers from \mathcal{S} to \mathcal{T} . Their correctness theorems (slightly edited here for presentation) apply to any well-typed source module:

Theorem 22 (Correctness of Pilsner).

$$\frac{\Gamma \vdash M_{\mathcal{S}} : \Gamma'}{\Gamma \vdash \text{Pilsner}(M_{\mathcal{S}}) \lesssim_{\mathcal{T}\mathcal{S}} M_{\mathcal{S}} : \Gamma'}$$

Theorem 23 (Correctness of Zwikel).

$$\frac{\Gamma \vdash M_{\mathcal{S}} : \Gamma'}{\Gamma \vdash \text{Zwikel}(M_{\mathcal{S}}) \lesssim_{\mathcal{T}\mathcal{S}} M_{\mathcal{S}} : \Gamma'}$$

While Zwikel carries out a straightforward direct translation from \mathcal{S} to \mathcal{T} , Pilsner is more sophisticated: as shown in Figure 1.1, it compiles via an intermediate language \mathcal{I} and performs several optimizations. We will discuss Pilsner and Zwikel in detail in Chapter 5.

These results mean that we can preserve correctness not only by linking, say, Pilsner-produced code with other Pilsner-produced code, but also by linking it with code produced by Zwikel.

Moreover, we would like to stress two important points. PILS were designed with flexibility in mind and make only few assumptions about the translation of source programs, namely details of the calling convention and in-memory representation of values (see the subsequent sections). Consequently:

1. Nothing stops us from proving a theorem analogous to the previous two for *yet another* compiler from \mathcal{S} to \mathcal{T} , perhaps even using several different intermediate languages.
2. Nothing stops us from proving the relatedness of a source and target module *by hand*, *e.g.*, when the target module is not the direct result of a compiler run but was manually optimized (see Section 5.13 for an extreme example of this, where the target module overwrites its own code at run time).

Hence, we can also preserve correctness when linking with code that was produced by other compilers or even hand-translated. We only have to ensure that these translations are also correct w.r.t. $\lesssim_{\mathcal{T}\mathcal{S}}$, such that Theorem 14 applies.

4.1.1 Transitivity

Proving a property like Theorem 22 can require a lot of effort: the more complex the compiler, the more complex its correctness proof. It is thus crucial that a correctness proof can be broken up into several pieces, *e.g.*, one sub-proof per compiler pass. PILS support such a decomposition thanks to a transitivity-like property. In our setting, where Pilsner compiles via one intermediate language \mathcal{I} , we can show the following:

Theorem 15 (Transitivity).

$$\frac{\begin{array}{c} |\Gamma| \vdash M_{\mathcal{T}} \lesssim_{\mathcal{TI}} M_{\mathcal{I}} : |\Gamma'| \quad |\Gamma| \vdash M_{\mathcal{I}} \lesssim_{\mathcal{II}}^* M'_{\mathcal{I}} : |\Gamma'| \\ \Gamma \vdash M'_{\mathcal{I}} \lesssim_{\mathcal{IS}} M_{\mathcal{S}} : \Gamma' \end{array}}{\Gamma \vdash M_{\mathcal{T}} \lesssim_{\mathcal{TS}} M_{\mathcal{S}} : \Gamma'}$$

Here, $\lesssim_{\mathcal{TI}}$ relates target modules to intermediate modules, $\lesssim_{\mathcal{II}}$ relates intermediate modules to intermediate modules, and $\lesssim_{\mathcal{IS}}$ relates intermediate modules to source modules. All are very similar to $\lesssim_{\mathcal{TS}}$ and support similar reasoning principles. We will say more about them in the following sections; for now suffice it to say that, since $\lesssim_{\mathcal{TI}}$ and $\lesssim_{\mathcal{II}}$ involve only untyped languages,¹ the relations themselves are “untyped” and we erase the typing annotations in their environments (*e.g.*, written $|\Gamma|$), leaving just a list of function labels. Notice how using the transitive closure of $\lesssim_{\mathcal{II}}$ in the second premise of the rule allows us to verify each IL transformation separately.

4.2 Languages

4.2.1 Language-Generic Approach

In order to avoid duplicate work, we define PILS in a language-generic way, *i.e.*, we define similarity \lesssim_{AB} for two abstract languages A and B (for some notion of abstract language to be described), and then instantiate it with different language pairs in order to obtain the desired relations. This has two important benefits:

1. Most of the metatheory, which is quite involved, can be established once and for all. This is particularly crucial because PILS were developed from the start in Coq, and over time the definitions—and thus proofs—had to undergo countless changes. This might simply have been infeasible if it weren’t for the language-generic setup.
2. One can easily instantiate PILS with a different intermediate language (or several of them) in order to verify a different compiler. Using modularity (Theorem 14), one can then of course safely link \mathcal{T} code produced by this compiler with Pilsner-produced code.

¹In order to demonstrate that PILS are not inherently tied to typed languages, we consider a type-erasing compiler, not a completely type-directed one.

In (1), we say “most of the metatheory” because transitivity and the parts of modularity and adequacy that deal with details of module loading are actually not proven generically. Ultimately, it would be nice to do so but it would require some effort to properly axiomatize various properties of the abstract language that these proofs rely on. Moreover, it might require a distinction between *intermediate language* and *non-intermediate language*, with slightly different sets of requirements. For now, it is much easier to simply prove the theorems for the concrete instances (of course this involves many generically proven lemmas). The downside of this is that, in order to verify a compiler using different intermediate languages, one needs to reprove the corresponding transitivity property. Adequacy and modularity, on the other hand, do not need to be reproven as they do not involve the intermediate languages.

To instantiate the generic PILS model and obtain one of the desired similarity relations requires us to provide: (i) the pair of concrete languages, and (ii) the *global world* for this pair. The latter can be seen as a predefined protocol (in the sense of Chapter 2) responsible for fixing calling conventions and data representations. We will discuss global worlds further in Section 4.4.

One point that we glossed over so far is that our generic definition is also not entirely generic—as we will see in the next section, it still essentially bakes in our source language’s type structure. Consequently, instantiating PILS as-is with a different source language most likely won’t make much sense. This is fine, because we focus on a single source language in this work. Extending this to multiple source languages, perhaps even allowing interoperability between them, is clearly important but left to future work.

Another point we glossed over is that we actually define *two* generic models: a typed one and an untyped one. The former is used when the input language is \mathcal{S} , the latter is used in all other cases (where no involved language has static types). However, we will continue to refer to them as just “the generic model”, because the untyped one is obtained simply by erasing all the type arguments from the typed one.

4.2.2 Language Specification

We now describe the language abstraction in terms of which PILS are defined. In the subsequent sections, we then present the concrete languages under consideration ($\mathcal{S}, \mathcal{I}, \mathcal{T}$). Common to all languages are a set of events and a countably infinite set of labels:

$$t \in \text{Evt} ::= \iota \mid ?n \mid !n \qquad F_1, F_2, \dots \in \text{Lbl}$$

Events are produced by an executing program; they consist of internal computation (ι) and I/O operations (reading or writing a number n , respectively). Labels are used to identify module components; in this work, we consider a simplistic notion of module as the unit of compilation.

Figure 4.1 presents the abstract language in terms of a signature that any concrete language must implement. Keep in mind that we need to account for a very high-

Domains: **Val**, **Cont**, **Conf**, **Mach**, **Mod**, **Anch**

Operators and relations:

- $\cdot \in \mathbf{Conf} \rightarrow \mathbf{Conf} \rightarrow \mathbf{Conf}$
- $\emptyset \in \mathbf{Conf}$
- $\hookrightarrow \in \mathcal{P}(\mathbf{Mach} \times \mathbf{Evt} \times \mathbf{Mach})$
- $\mathbf{real} \in \mathbf{Conf} \rightarrow \mathcal{P}(\mathbf{Mach})$
- $\mathbf{cload} \in \mathbf{Mod} \rightarrow \mathbf{Anch} \rightarrow (\mathbf{Lbl} \times \mathbf{Val})^* \rightarrow \mathcal{P}(\mathbf{Conf} \times \mathbf{Conf})$
- $\mathbf{vload} \in \mathbf{Mod} \rightarrow \mathbf{Anch} \rightarrow (\mathbf{Lbl} \times \mathbf{Val})^* \rightarrow \mathbf{Lbl} \rightarrow \mathcal{P}(\mathbf{Val})$
- $\mathbf{frame} \in \mathcal{P}(\mathbf{Conf})$
- $\mathbf{core} \in \mathcal{P}(\mathbf{Conf})$
- $\mathbf{halted} := \{m \in \mathbf{Mach} \mid \nexists t, m'. m \xrightarrow{t} m'\}$
- $\mathbf{error} := \{m \in \mathbf{Mach} \mid \forall c. m \notin \mathbf{real}(c)\}$

Properties:

- **Conf** forms a commutative monoid with \cdot and \emptyset .
 - $\forall m, t, m'. m \xrightarrow{t} m' \wedge m' \notin \mathbf{error} \implies m \notin \mathbf{error}$
 - $\forall m, t, m'. m \notin \mathbf{error} \wedge m \xrightarrow{t} m' \wedge m' \in \mathbf{error} \implies t = \iota$
 - $\forall c_1, c_2, c'_1, c'_2, m. c_1 \in \mathbf{core} \wedge c_2 \in \mathbf{core} \wedge m \in \mathbf{real}(c_1 \cdot c'_1) \cap \mathbf{real}(c_2 \cdot c'_2) \implies c_1 = c_2 \wedge c'_1 = c'_2$
 - $\emptyset \in \mathbf{frame}$
 - $\forall c, c'. c \in \mathbf{frame} \wedge c' \in \mathbf{frame} \implies c \cdot c' \in \mathbf{frame}$
-

Figure 4.1: Language specification

level language (\mathcal{S}) on the one extreme and a very low-level language (\mathcal{T}) on the other extreme.

A language must come with a set **Val** of *values*, a set **Cont** of *continuations*, a set **Conf** of *configurations*, a set **Mach** of *machines*, a set **Mod** of *modules*, and a set **Anch** of *anchors* (think: load addresses). The operational semantics is given in the form of a transition system (\hookrightarrow) of machines, whose transitions are labelled with events t . Configurations can be thought of as partial machines—they play different roles in different contexts (*e.g.*, they might represent just an \mathcal{S} expression, or just an \mathcal{I} heap, or even a full \mathcal{T} machine). If a configuration c is complete, it is *realized* by a set of machines $\mathbf{real}(c)$. (In all our instantiations, this is either empty, meaning the configuration is invalid or incomplete, or it contains exactly one machine.) Configurations must form a commutative monoid with composition \cdot and neutral element \emptyset . This monoid is implicitly partial in the sense that a composition may not be realizable. (Having \cdot itself be total is more convenient for mechanization [59]).

We single out two kinds of configurations, described by **frame** and **core**. *Cores* are what PILS’s **E** will (primarily) relate—they can be thought of as expressions. *Frames* are configurations that may be contributed by other modules; they generalize the frame heaps that we have seen in Section 3.6. Intuitively, cores and frames are disjoint, although we do not formally require this. Note that there may be configurations that cannot be decomposed into only a core and a frame.

For **core** we require a uniqueness property saying that a realizable configuration can contain only one core. For **frame** we require two very natural properties: the empty configuration must be a frame; and composition of frames must again yield a frame.

We say a machine is *halted*, $m \in \mathbf{halted}$, iff it cannot step any further. It denotes an *error*, $m \in \mathbf{error}$, iff it does not realize any configuration. We require two properties of **error** (*i.e.*, of **real**). The first is that an error “cannot be undone”; apart from this we will actually never care about steps from an erroneous machine. The second says that an error can only be produced silently. It is not needed for the PILS metatheory but, if violated, the model won’t validate some basic refinements.

4.2.2.1 Behavior

Based on these constructs, we can define the set of observation traces that machines can produce.

Definition 20 (Observation traces). The set **Obs** of observation traces is defined coinductively (as the greatest fixed point of a monotone function in a powerset lattice) such that:

$$\mathbf{Obs} = \{\$, \infty\} \uplus (\mathbf{Evt} \setminus \{\iota\} \times \mathbf{Obs})$$

Accordingly, an observation trace o is a sequence of externally visible events (*i.e.*, I/O events) that

1. either goes on infinitely, representing a computation that never stops issuing I/O operations,

2. or ends with a termination marker (\$), representing a computation that produces a finite number of I/O events and then halts,
3. or ends with a divergence marker (∞), representing an execution that produces a finite number of I/O events but never stops performing internal computations.

A machine m 's *behavior* is the set of observation traces that m can possibly generate.

Definition 21 (Machine behavior). We define $\mathbf{behav} \in \mathbf{Mach} \rightarrow \mathcal{P}(\mathbf{Obs})$ coinductively such that:

$$\begin{aligned} \mathbf{behav}(m) = \{ & o \mid \exists m'. m \xrightarrow{\iota}^* m' \wedge \\ & \text{(ERR)} \quad m' \in \mathbf{error} \\ & \vee \text{(HALT)} \quad o = \$ \wedge m' \in \mathbf{halted} \\ & \vee \text{(INF)} \quad o = \infty \wedge \exists m''. m' \xrightarrow{\iota} m'' \wedge \infty \in \mathbf{behav}(m'') \\ & \vee \text{(EVT)} \quad \exists t, o', m''. o = t, o' \wedge m' \xrightarrow{t} m'' \wedge t \neq \iota \wedge o' \in \mathbf{behav}(m'') \} \end{aligned}$$

Notice that errors cannot be observed explicitly. Instead, when an execution produces an error, then from that point on anything can happen, *i.e.*, an arbitrary trace is appended.

4.2.2.2 Modules

Modules, the units of compilation (and linking), can be thought of as collections of labelled values (the *exports*). These values may refer to external functions (the *imports*) by their unique labels.

In the language specification, the module interface consists of two operations, **load** and **vload**. The former, **load**, takes an anchor saying “where” the module is to be loaded and values for each of its imports. It then returns the set of configurations—split into a global and a local part—in which the module is considered loaded as part of a complete program (this will become clearer in the next section). Given the same inputs and additionally the label of one of the exported values, **vload** looks up the module's corresponding value.

The reason why the language specification does not talk about a linking operator is that we will not prove modularity at the language-generic level (cf. Section 4.2.1).

4.2.3 Source Language \mathcal{S}

4.2.3.1 Syntax and Static Semantics

The source language \mathcal{S} is essentially a straightforward extension of $F^{\mu!}$ (Section 3.3) with simple modules and numeric I/O. Figure 4.2 presents the additions to syntax and typing, as well as new definitions.

	$p \quad ::= \dots \mid F \mid \text{input} \mid \text{output } p$	
	$\Sigma \quad ::= \epsilon \mid F:\sigma, \Sigma$	
	$M \in \text{Mod} ::= \cdot \mid F=p, M$	

$\Delta; \Gamma; \Sigma \vdash p : \sigma$

...	$\frac{\Delta \vdash \Gamma; \Sigma \quad F:\sigma \in \Sigma}{\Delta; \Gamma; \Sigma \vdash F : \sigma}$	$\frac{\Delta \vdash \Gamma; \Sigma}{\Delta; \Gamma; \Sigma \vdash \text{input} : \text{nat}}$	$\frac{\Delta; \Gamma; \Sigma \vdash p : \text{nat}}{\Delta; \Gamma; \Sigma \vdash \text{output } p : \text{unit}}$
-----	---	--	---

$\Sigma \vdash M : \Sigma'$

	$\frac{\forall \sigma \in \Sigma. \exists \tau, \tau'. \sigma = \tau \rightarrow \tau' \vee \sigma = \forall \alpha. \tau}{\Sigma \vdash \epsilon : \epsilon}$	
$\epsilon; \epsilon; \Sigma \vdash p : \sigma \quad \Sigma, F:\sigma \vdash M : \Sigma' \quad F \notin \Sigma \quad p = v \quad \sigma = \tau \rightarrow \tau' \vee \sigma = \forall \alpha. \tau$	$\frac{}{\Sigma \vdash (F=p, M) : (F:\sigma, \Sigma')}$	
$\Sigma \vdash M_1 : \Sigma_1 \quad \Sigma, \Sigma_1 \vdash M_2 : \Sigma_2$	$\frac{}{\Sigma \vdash \text{link}(M_1, M_2) : \Sigma_1, \Sigma_2}$	

Figure 4.2: Syntax and Static Semantics of the Source language \mathcal{S}

The syntax is extended with labels (F), a construct for reading input, and one for writing output. The typing judgment now also carries an environment Σ mapping labels to types. Types are the same as in $F^{\mu!}$. All typing rules from $F^{\mu!}$ carry over (they simply ignore the new Σ component), and the three new ones are straightforward.

A module M is simply an ordered list of labelled “programs” p . The module typing judgment $\Sigma \vdash M : \Sigma'$ is also presented in Figure 4.2. For now, ignore the third rule, which we discuss in the next section. The typing ensures the following:

- The module has no free variables.
- The module’s imports are faithfully described by Σ .
- The module’s exports are faithfully described by Σ' as well as uniquely labelled and disjoint from its imports.
- The module’s components may refer to the labels of other components from the same module, but only in a strict left-to-right dependency order. (Notice how, in the premise of the second rule, F moved from the exports to the imports when checking the remaining module.)
- The module’s components all are function values (values of function or \forall types).

The last two restrictions are made to keep the module semantics simple and avoid too much distraction. Note, however, that \mathcal{S} supports universal and existential types and thus can in principle be used to code up ML-style modules [70]. Also, PILS themselves do not bake in any restriction on the dependency order and, being a coinductive method, are perfectly compatible with mutual recursion.

e	$::= \dots \mid F \mid \text{input} \mid \text{output } e$	
K	$::= \dots \mid \text{output } K$	
$h \in \text{Heap}$	$:= (\text{Loc} \xrightarrow{\text{fin}} \text{Val})_{\perp}$	
$\rho \in \text{Env}$	$:= \text{Lbl} \rightarrow \text{Val}$	

$\langle h; \rho; e \rangle \xrightarrow{t} \langle h'; \rho; e' \rangle$

$\langle h; \rho; e \rangle$	$\xrightarrow{t} \langle h'; \rho; e' \rangle$	$((\langle h; e \rangle \hookrightarrow \langle h'; e' \rangle))$
$\langle h; \rho; K[F] \rangle$	$\xrightarrow{t} \langle h; \rho; K[v] \rangle$	$(v = \rho(F))$
$\langle h; \rho; K[\text{input}] \rangle$	$\xrightarrow{?n} \langle h; \rho; n \rangle$	
$\langle h; \rho; K[\text{output } n] \rangle$	$\xrightarrow{!n} \langle h; \rho; \langle \rangle \rangle$	
$\langle h; \rho; e \rangle$	$\xrightarrow{t} \langle \perp; \rho; e \rangle$	$(e \neq v \text{ and none of the previous rules apply})$

Linking and loading:

$$\begin{aligned} \text{link} &\in \text{Mod} \times \text{Mod} \rightarrow \text{Mod} \\ \text{link}(M_1, M_2) &:= M_1, M_2 \\ \\ \text{load} &\in \text{Mod} \rightarrow \text{Heap} \times \text{Env} \times \text{Exp} \\ \text{load}(M) &:= (\emptyset, M, |p| \langle \rangle) \quad \text{if } M(F_{\text{main}}) = p \end{aligned}$$

Figure 4.3: Dynamic Semantics of Source language \mathcal{S}

4.2.3.2 Dynamic Semantics

As shown in Figure 4.3, expressions e are extended in the obvious way, as are evaluation contexts K . (Since we are not interested in \mathcal{S} 's contextual equivalence, we no longer care to define arbitrary contexts.) The notion of value is unchanged from that of $F^{\mu!}$, *i.e.*, labels are not values. Instead, they are looked up in a read-only environment ρ when in evaluation position. The reduction relation is therefore defined between triples of heap, expression, and label environment (rather than just heap and expression like before).

With the goal of a convenient implementation of configuration composition (\cdot) , we change our notion of heap to include an undefined heap, *i.e.*, we work with $\text{Heap} = (\text{Loc} \rightarrow \text{Val})_{\perp}$ rather than $\text{Loc} \rightarrow \text{Val}$. The undefined heap \perp is not to be confused with the empty heap $(\lambda l. \perp)$, for which we continue to write \emptyset . We overload the \sqcup notation to refer to the obvious merging of two lifted heaps that returns \perp if one of the heaps is \perp (and, as before, when the heaps overlap). Note that the empty heap \emptyset is its neutral element and \perp its absorbing element.

The reduction relation is now also labelled with events $t \in \text{Evt}$. Let us look at the four rules in Figure 4.3. The first rule incorporates any step from $F^{\mu!}$, labelling it as

internal computation (this includes $F^{\mu!}$'s non-deterministic allocation). The second rule describes label lookups via ρ . The third and fourth describe numeric input and output, respectively, which are labelled accordingly. Note that the rule for input is (externally) non-deterministic. Finally, whenever none of these rules are applicable and the expression is not a value, the last rule fires. It transitions to a designated error state, where the heap component is invalidated (we will see in a moment how this naturally fits the language specification). For instance, the last rule fires whenever a label is looked up that is not covered by ρ ; it also fires whenever (i) the expression is a read or write memory operation but the heap component doesn't contain the location being read or written, respectively, (ii) the expression is an allocation but the heap already contains any possible location, (iii) the expression is any of these but the heap is invalid (\perp).

Module linking is simply concatenation (and thus asymmetric). Indeed, the third module typing rule in the previous Figure 4.2, which talks about linking, can be derived from the first two rules. A complete program is a module without imports (be it the result of linking or not) that exports a designated main function (F_{main}) of type $\text{unit} \rightarrow \tau$ (for an arbitrary τ). In order to execute it, it has to be *load*-ed: the initial heap is empty, the (constant) label environment is the module itself, and the initial expression is a call to the main function. Note that since this call is done in the empty evaluation context, the machine will be halted when the function eventually returns (if it returns) thanks to the side condition $e \neq v$ in the last rule. This matters for adequacy, as will be explained in Section 4.7.2.

4.2.3.3 Implementation of the Language Specification

Figure 4.4 presents the way in which we decide \mathcal{S} to implement the abstract language from Section 4.2.1. Values, continuations, and modules are self-explanatory. Since there is no need for anchors in the source language, we set them to be the singleton set 1 . Machines are the triples for which we defined the reduction \hookrightarrow , and \hookrightarrow is precisely \hookrightarrow . We write $m.\text{hp}$ to project the heap component out of the tuple m .

In order to define configurations, we make use of the following lifting operation.

Definition 22. Given a set X , we write $X_{\perp, \emptyset}$ for the commutative monoid with carrier $X \uplus \{\emptyset, \perp\}$ (i.e., X extended with two new elements) and a composition (\oplus) that is defined as follows:

$$\begin{aligned} \forall y \in X_{\perp, \emptyset}. \emptyset \oplus y &= y = y \oplus \emptyset \\ \forall y \in X_{\perp, \emptyset}. \perp \oplus y &= \perp = y \oplus \perp \\ \forall x_1, x_2 \in X. x_1 \oplus x_2 &= \perp \end{aligned}$$

Accordingly, \emptyset is neutral, \perp is absorbing, and composition is *exclusive* in that it maps any two original elements to \perp .

Configurations are machines where one or more components may be missing or invalid. For a machine m to realize a configuration c , it must match the configuration

Val	$:= \text{Val}$
Cont	$::= \text{Cont}$
Mod	$::= \text{Mod}$
Anch	$:= 1$
Mach	$:= \text{Heap} \times \text{Env} \times \text{Exp}$
Conf	$:= \text{Heap} \times \text{Env}_{\perp, \emptyset} \times \text{Exp}_{\perp, \emptyset}$
Conf monoid:	
\emptyset	$:= (\emptyset, \emptyset, \emptyset) \quad (h, \rho, e) \cdot (h', \rho', e') := (h \sqcup h', \rho \cdot \rho', e \cdot e')$
\hookrightarrow	$:= \hookrightarrow$
real (c)	$:= \{m \mid m = c \wedge m.\text{hp} \neq \perp \wedge m.\text{hp} \text{ finite}\}$
clload (M)($_$)(ρ)	$:= \{(c, \emptyset) \mid \exists \rho'. c = (\emptyset, \rho \uplus M \uplus \rho', \emptyset)\}$
vload (M)($_$)($_$)(F)	$:= \{v \mid \exists p. v = p \wedge (F=p) \in M\}$

Figure 4.4: Implementation of Language Specification for \mathcal{S}

($m = c$ with the obvious embedding of **Mach** in **Conf**). Moreover, it must carry a valid and finite heap (finiteness guarantees that allocation will succeed). Notice that the definition of **Conf** and its pointwise composition operation (\cdot) imply that heaps can successfully be split across several configurations, but label environments and expressions cannot—they must be defined in exactly one component in order for the composition to be realizable.

The implementation of **vload** is straightforward: it simply looks up the given label in the module. Since the anchor argument carries no information, it is ignored (we write $_$ to indicate this). The imported values are ignored as well since labels are looked up at runtime (this is also why linking in \mathcal{S} is just concatenation). The implementation of **clload** permits any configuration whose heap is empty, whose expression component is missing, and whose environment contains the imported values and the exported values (*i.e.*, the module itself). Note that the environment may contain additional values (ρ'), thus accounting for other modules that M may be linked with.

4.2.4 Intermediate Language \mathcal{I}

Pilsner’s intermediate language \mathcal{I} is an untyped, or rather, dynamically typed CPS-variant of \mathcal{S} , inspired by Kennedy’s intermediate language [39]. Its syntax is shown in Figure 4.5, together with a notion of well-formedness carving out the subset of the language that we are interested in (to be discussed in a moment).

$$a \in \text{BExp} ::= \langle \rangle \mid n \mid x.1 \mid x.2 \mid \text{inl } x \mid \text{inr } x \mid \langle x_1, x_2 \rangle \mid x_1 == x_2 \mid x_1 \odot x_2 \mid \text{fix } f(y, k).e \mid \Lambda k.e$$

$$e \in \text{Exp} ::= \text{let } y = a \text{ in } e \mid \text{let } k = \text{cont } y.e \text{ in } e \mid y \leftarrow \text{input}; e \mid \text{output } x; e \mid y \leftarrow \text{ref } x; e \mid x_1 := x_2; e \mid y \leftarrow !x; e \mid \text{ifnz } x \text{ then } e_1 \text{ else } e_2 \mid \text{case } x (y. e_1) (y. e_2) \mid x_1 x_2 k \mid x [] k \mid k x \mid \text{halt } n$$

$$M \in \text{Mod} ::= \cdot \mid F=a, M$$

$$Z ::= \epsilon \mid x, Z \mid k, Z$$

$$\boxed{Z \vdash a}$$

$$\frac{}{Z \vdash \langle \rangle} \quad \frac{}{Z \vdash n} \quad \frac{x \in Z}{Z \vdash x.1} \quad \frac{x \in Z}{Z \vdash x.2} \quad \frac{x \in Z}{Z \vdash \text{inl } x} \quad \frac{x \in Z}{Z \vdash \text{inr } x}$$

$$\frac{x_1, x_2 \in Z}{Z \vdash \langle x_1, x_2 \rangle} \quad \frac{x_1, x_2 \in Z}{Z \vdash x_1 \odot x_2} \quad \frac{x_1, x_2 \in Z}{Z \vdash x_1 == x_2}$$

$$\frac{Z^\dagger, f, y, k \vdash e}{Z \vdash \text{fix } f(y, k).e} \quad \frac{Z^\dagger, k \vdash e}{Z \vdash \Lambda k.e}$$

$$\boxed{Z \vdash e}$$

$$\frac{Z \vdash a \quad Z, y \vdash e}{Z \vdash \text{let } y = a \text{ in } e} \quad \frac{Z, y \vdash e \quad Z, k \vdash e'}{Z \vdash \text{let } k = \text{cont } y.e \text{ in } e'}$$

$$\frac{Z, y \vdash e}{Z \vdash y \leftarrow \text{input}; e} \quad \frac{x \in Z \quad Z \vdash e}{Z \vdash \text{output } x; e}$$

$$\frac{x \in Z \quad Z \vdash e_1 \quad Z \vdash e_2}{Z \vdash \text{ifnz } x \text{ then } e_1 \text{ else } e_2} \quad \frac{x \in Z \quad Z, y \vdash e_1 \quad Z, y \vdash e_2}{Z \vdash \text{case } x (y. e_1) (y. e_2)}$$

$$\frac{x_1, x_2, k \in Z}{Z \vdash x_1 x_2 k} \quad \frac{x, k \in Z}{Z \vdash x [] k} \quad \frac{k, x \in Z}{Z \vdash k x}$$

$$\frac{x \in Z \quad Z, y \vdash e}{Z \vdash y \leftarrow \text{ref } x; e} \quad \frac{x \in Z \quad Z, y \vdash e}{Z \vdash y \leftarrow !x; e} \quad \frac{x_1, x_2 \in Z \quad Z \vdash e}{Z \vdash x_1 := x_2; e}$$

$$\boxed{L \vdash M : L'}$$

$$\frac{}{L \vdash \epsilon : \epsilon} \quad \frac{L \vdash a \quad L, F \vdash M : L' \quad F \notin L}{L \vdash (F=a, M) : (F, L')}$$

Figure 4.5: Intermediate language \mathcal{I} : Syntax and Well-formedness

4.2.4.1 Syntax

Being in continuation-passing style, every subexpression is explicitly named and functions never “return”. Concretely, we distinguish between (i) *pure expressions* $a \in \text{BExp}$, which are evaluated in let-bindings and always result in a value but never in any side effects, and (ii) *control expressions* $e \in \text{Exp}$. Ignoring conditionals, every control expression is essentially a sequence of let-bindings ending in a function or continuation invocation. For instance, let $k = \text{cont } y. e_1$ in e_2 defines a new continuation k with argument y and body e_1 , and then executes e_2 (which may use k). Here $y \in \text{TVar}$ is term variable, while $k \in \text{KVar}$ is a continuation variable. Any x in the language syntax is a meta-variable standing for either a term variable or a label, *i.e.*, ranges over $\text{TVar} \uplus \text{Lbl}$ but not KVar . Any y is in binding position.

Functions in \mathcal{I} explicitly take a continuation argument k , and that the syntax for function calls accounts for that as well. Jumping to a continuation k , which can either be the one initially given, or a self-defined one, is written $k \ x$, where x is the continuation’s argument (in case k is a function’s initial continuation, x should be thought of as the function’s return value). Note that one cannot evade the initial continuation: defining a new continuation requires a pre-existing continuation for the final control expression. Also note that continuations are second-class citizens in the sense that the syntax does not allow them to be components of pairs, regular arguments to functions, etc.

The other syntactic forms are self-explanatory (but see below for `halt n`).

4.2.4.2 Well-formedness

The well-formedness predicates in Figure 4.5 ensure that programs are well-scoped. For instance, $Z \vdash a$ guarantees that the only labels and free variables that occur in a are (a subset of) those in Z .

More interestingly, however, well-formedness also enforces that continuations are used in an affine fashion [39], *i.e.*, called at most once. Affinity is exploited by Pilsner’s code generation pass, as discussed in Section 5.10. Responsible for enforcing affinity are the rules for functions (the one for $\text{fix } f(y, k). e$ and the one for $\Lambda k. e$): when checking the body e in the premise, we not only extend the original environment Z with the function’s bindings, but we also *remove all continuations variables* from Z , written Z^\dagger . Hence the only continuation initially available to a function is the one it is passed, namely k . This enforces affinity because, roughly, the only way a continuation could be used multiple times is by embedding it into a function closure. Consider the following example, which is ruled out by well-formedness:

$$\begin{array}{l} \text{fix } f(y, k). \\ \quad \text{let } g = (\text{fix } g(-, k'). k \ y) \text{ in} \\ \quad k \ g \end{array}$$

The function f defines a function g that refers to f ’s continuation k (which well-formedness forbids). It then passes g to k . If k were now to call its argument g , this

would cause a second invocation of k .

Well-formedness for modules is what one would expect.

4.2.4.3 Dynamic Semantics

Figure 4.6 presents \mathcal{I} 's dynamic semantics (which does not care about well-formedness). In contrast to the source language, \mathcal{I} employs an environment-based semantics, where continuations and functions evaluate to closures of code and environment. This avoids the need to reason about substitutions when verifying optimizations, which is often a hassle.

Run-time values consist of the unit value, numbers, pairs, left/right tagged values, and closures $\langle \sigma; a \rangle$. A closure's environment component σ is a mapping from variables (including labels) to values².

In order to give the semantics of let-bindings, we first define a straightforward evaluation function that evaluates a pure expression a in an environment σ , written $\llbracket a \rrbracket_\sigma$. It is total on inputs whose free variables are covered by σ . The small-step reduction relation is fairly straightforward as well. As one can easily see, the pure expression component a of a closure is always either a function or continuation expression. For instance, we execute a continuation definition `let $k = \text{cont } y. e_1$ in e_2` by creating a new closure consisting of the current environment σ as well as the continuation's code `cont $y. e_1$` , and add this value to the environment under name k . We would then proceed to execute e_2 in this extended environment.

The `halt n` expression form is merely a technical device that we use in the transitivity proof (Section 4.8). The compiler is oblivious to it. Since \mathcal{I} is an intermediate language, used internally by Pilsner, adding such a dummy expression is unproblematic. Note that reduction is defined to be stuck (**halted**) for such an expression. It is even ruled out by the well-formedness relation above.

4.2.4.4 Implementation of the Language Specification

Figure 4.7 shows how we implement the language specification with \mathcal{I} . Modules, anchors, configurations, etc. are fairly similar to those in the source language. We define **Cont** simply as **Val** because continuations are already values in the language.

Since \mathcal{I} 's semantics is closure-based, expressions and their run-time environment (σ) are tightly coupled, and we do not need to allow configurations containing an environment but no expression, or vice versa. Hence, in the implementation of **Conf**, we lift **Env** and **Exp** not individually but together as $(\text{Env} \times \text{Exp})_{\perp, \emptyset}$.

For the same reason, the implementation of **load** is trivial: because function values already contain their environment, there is nothing else to load, so both the global and local configuration are empty. The implementation of **vload**, on the other hand, is consequently more complicated than before. Defined by recursion on the structure of the given module, it looks up the function expression a associated

²The operational semantics of \mathcal{I} treats term variables, continuation variables, and labels all in a uniform way.

$l \in \text{Loc}$
 $v \in \text{Val} ::= \langle \rangle \mid n \mid l \mid \langle v_1, v_2 \rangle \mid \text{inl } v \mid \text{inr } v \mid \langle \sigma; a \rangle \mid \langle \sigma; \text{cont } y. e \rangle$
 $h \in \text{Heap} := (\text{Loc} \xrightarrow{\text{fin}} \text{Val})_{\perp}$
 $\sigma \in \text{Env} := \text{Lbl} \uplus \text{TVar} \uplus \text{KVar} \rightarrow \text{Val}$

$\llbracket a \rrbracket_{\sigma} \in \text{Val}_{\perp}$

$\llbracket \langle \rangle \rrbracket_{\sigma} := \langle \rangle$
 $\llbracket n \rrbracket_{\sigma} := n$
 $\llbracket \langle x_1, x_2 \rangle \rrbracket_{\sigma} := \langle v_1, v_2 \rangle \quad (\sigma(x_1) = v_1 \wedge \sigma(x_2) = v_2)$
 $\llbracket x.1 \rrbracket_{\sigma} := v_1 \quad (\sigma(x) = \langle v_1, v_2 \rangle)$
 $\llbracket x.2 \rrbracket_{\sigma} := v_2 \quad (\sigma(x) = \langle v_1, v_2 \rangle)$
 $\llbracket \text{inl } x \rrbracket_{\sigma} := \text{inl } v \quad (\sigma(x) = v)$
 $\llbracket \text{inr } x \rrbracket_{\sigma} := \text{inr } v \quad (\sigma(x) = v)$
 $\llbracket \text{fix } f(y, k). e \rrbracket_{\sigma} := \langle \sigma; \text{fix } f(y, k). e \rangle$
 $\llbracket \Lambda k. e \rrbracket_{\sigma} := \langle \sigma; \Lambda k. e \rangle$
 $\llbracket x_1 \odot x_2 \rrbracket_{\sigma} := \llbracket \odot \rrbracket (n_1, n_2) \quad (\sigma(x_1) = n_1 \wedge \sigma(x_2) = n_2)$
 $\llbracket x_1 == x_2 \rrbracket_{\sigma} := 1 \quad (\sigma(x_1) = l = \sigma(x_2))$
 $\llbracket x_1 == x_2 \rrbracket_{\sigma} := 0 \quad (\sigma(x_1) = l_1 \neq l_2 = \sigma(x_2))$

$\langle h; \sigma, e \rangle \xrightarrow{t} \langle h'; \sigma', e' \rangle$

$\langle h; \sigma, \text{let } y = a \text{ in } e \rangle \xrightarrow{t} \langle h; \sigma[y \mapsto v], e \rangle \quad (\llbracket a \rrbracket_{\sigma} = v)$
 $\langle h; \sigma, \text{let } k = \text{cont } y. e_1 \text{ in } e_2 \rangle \xrightarrow{t} \langle h; \sigma[k \mapsto \langle \sigma; \text{cont } y. e_1 \rangle], e_2 \rangle$
 $\langle h; \sigma, (y \leftarrow \text{input}; e) \rangle \xrightarrow{?n} \langle h; \sigma[y \mapsto n], e \rangle$
 $\langle h; \sigma, (\text{output } x; e) \rangle \xrightarrow{!n} \langle h; \sigma, e \rangle \quad (\sigma(x) = n)$
 $\langle h; \sigma, \text{ifnz } x \text{ then } e_1 \text{ else } e_2 \rangle \xrightarrow{t} \langle h; \sigma, e_2 \rangle \quad (\sigma(x) = 0)$
 $\langle h; \sigma, \text{ifnz } x \text{ then } e_1 \text{ else } e_2 \rangle \xrightarrow{t} \langle h; \sigma, e_1 \rangle \quad (\sigma(x) = n \neq 0)$
 $\langle h; \sigma, \text{case } x (y. e_1) (y. e_2) \rangle \xrightarrow{t} \langle h; \sigma[y \mapsto v], e_1 \rangle \quad (\sigma(x) = \text{inl } v)$
 $\langle h; \sigma, \text{case } x (y. e_1) (y. e_2) \rangle \xrightarrow{t} \langle h; \sigma[y \mapsto v], e_2 \rangle \quad (\sigma(x) = \text{inr } v)$
 $\langle h; \sigma, k \ x \rangle \xrightarrow{t} \langle h; \sigma'[y \mapsto v], e \rangle \quad (\sigma(k) = \langle \sigma'; \text{cont } y. e \rangle)$
 $\quad \quad \quad (\sigma(x) = v)$
 $\langle h; \sigma, x_1 \ x_2 \ k \rangle \xrightarrow{t} \langle h; \sigma'[f, y, k' \mapsto v_1, v_2, v], e \rangle$
 $\quad \quad \quad (\sigma(x_1), \sigma(x_2), \sigma(k) = v_1, v_2, v)$
 $\quad \quad \quad (v_1 = \langle \sigma'; \text{fix } f(y, k'). e \rangle)$
 $\langle h; \sigma, x \ \square \ k \rangle \xrightarrow{t} \langle h; \sigma'[k' \mapsto v], e \rangle \quad (\sigma(x) = \langle \sigma'; \Lambda k'. e \rangle)$
 $\quad \quad \quad (\sigma(k) = v)$
 $\langle h; \sigma, (y \leftarrow \text{ref } x; e) \rangle \xrightarrow{t} \langle h \sqcup [l \mapsto v]; \sigma[y \mapsto l], e \rangle \quad (\sigma(x) = v)$
 $\quad \quad \quad (h \sqcup [l \mapsto v] \neq \perp)$
 $\langle h; \sigma, (y \leftarrow !x; e) \rangle \xrightarrow{t} \langle h; \sigma[y \mapsto v], e \rangle \quad (\sigma(x) = l)$
 $\quad \quad \quad (h(l) = v)$
 $\langle h \sqcup [l \mapsto v]; \sigma, (x_1 := x_2; e) \rangle \xrightarrow{t} \langle h \sqcup [l \mapsto v']; \sigma, e \rangle \quad (\sigma(x_1) = l)$
 $\quad \quad \quad (\sigma(x_2) = v')$
 $\langle h; \sigma, e \rangle \xrightarrow{t} \langle \perp; \sigma, e \rangle \quad (e \neq \text{halt } n)$
 $\quad \quad \quad (\text{no other rule applicable})$

Figure 4.6: Intermediate language \mathcal{I} : Dynamics

Val $:= \text{Val}$
Cont $::= \text{Val}$
Mod $::= \text{Mod}$
Anch $::= 1$
Mach $::= \text{Heap} \times (\text{Env} \times \text{Exp})$
Conf $::= \text{Heap} \times (\text{Env} \times \text{Exp})_{\perp, \emptyset}$

Conf monoid:
 $\emptyset := (\emptyset, \emptyset) \quad (h, x) \cdot (h', x') := (h \sqcup h', x \oplus x')$

$\hookrightarrow := \hookrightarrow$
real(c) $:= \{m \mid m = c \wedge m.\text{hp} \neq \perp \wedge m.\text{hp} \text{ finite}\}$

vload(ϵ)(ρ)(F) $:= \emptyset$
vload($F'=a, M$)(ρ)(F) $:= \{\langle \rho; a \rangle\} \quad (F' = F)$
vload($F'=a, M$)(ρ)(F) $:= \mathbf{vload}(M)(\rho, (F', \langle \rho; a \rangle))(F) \quad (F' \neq F)$

clload(M)(ρ) $:= \{(\emptyset, \emptyset)\}$

Figure 4.7: Intermediate language \mathcal{I} : implementation of the language specification

with the given label F and constructs a suitable environment for it. This environment is the given one (σ) extended with all the values previously defined in the module, which are computed recursively.

4.2.5 Target Language \mathcal{T}

As shown in Figure 4.8, our target language \mathcal{T} is an idealized machine language featuring instructions for arithmetic, control flow (*e.g.*, conditional jump), memory access, and numeric I/O. Some of them support multiple addressing modes. For instance, if operand o is $\langle r_1 \pm n \rangle$, then the instruction **sto** o r_2 stores the contents of register r_2 on the stack at the address contained in register r_1 , offset by $\pm n$. If $o = [r_1 \pm n]$, then it stores it on the heap instead. The **lpc** instruction loads the current program counter into the given register.

Code is encoded as data (using an injective function \mathbb{E}) and can thus be modified using the standard store instruction (**sto**).

The **alloc** instruction allocates a chunk of heap memory of the given size. We assume the choice of memory to be deterministic (but otherwise arbitrary).

\mathcal{T} is idealized also in the sense that machine words are unbounded natural numbers and stack and heap are unbounded as well. The set of registers, though, is fixed (their names, by the way, are merely suggestive and the language semantics does not differentiate them).

$r \in \text{Reg}$	$::=$	$\text{sp} \mid \text{clo} \mid \text{arg} \mid \text{env} \mid \text{ret} \mid \text{aux}_1 \mid \text{aux}_2 \mid \text{aux}_3$
$o \in \text{Opr}$	$::=$	$n \mid r \mid \langle r \pm n \rangle \mid [r \pm n]$
$z \in \text{Instr}$	$::=$	$\text{jmp } o \mid \text{jnz } r \ o \mid \text{ld } r \ o \mid \text{sto } o \ r \mid \text{lpc } r \mid \text{bop } \odot \ r \ o_1 \ o_2 \mid$ $\text{input } r \mid \text{output } r \mid \text{alloc } r_1 \ r_2$
\mathbb{E}	\in	$\text{Instr} \rightarrow \text{Word} \quad (\text{injective})$
Word	$::=$	\mathbb{N}
$R \in \text{RegFile}$	$::=$	$\text{Reg} \rightarrow \text{Word}$
$q \in \text{Stack}$	$::=$	$(\text{Word} \rightarrow \text{Word})_{\perp}$
$h \in \text{Heap}$	$::=$	$(\text{Word} \rightarrow \text{Word})_{\perp}$
$\llbracket - \rrbracket (-, -, -)$	\in	$\text{Opr} \rightarrow \text{RegFile} \times \text{Stack} \times \text{Heap} \rightarrow \text{Word}$
$\llbracket n \rrbracket (R, q, h)$	$::=$	n
$\llbracket r \rrbracket (R, q, h)$	$::=$	$R(r)$
$\llbracket \langle r \pm n \rangle \rrbracket (R, q, h)$	$::=$	$w \quad (q(R(r) \pm n) = w)$
$\llbracket [r \pm n] \rrbracket (R, q, h)$	$::=$	$w \quad (h(R(r) \pm n) = w)$
$\text{store} \in \text{Word} \times \text{RegFile} \times \text{Stack} \times \text{Heap} \rightarrow \text{Word} \rightarrow \text{Opr} \rightarrow \text{Word} \times \text{RegFile} \times \text{Stack} \times \text{Heap}$		
$\text{store}(n, R, q, h)(w)(r)$	$::=$	$(n + 1, R[r \mapsto w], q, h)$
$\text{store}(n, R, q, h)(w)(\langle r \pm n' \rangle)$	$::=$	$(n + 1, R, q \sqcup [R(r) \pm n' \mapsto w], h)$
$\text{store}(n, R, q, h)(w)([r \pm n'])$	$::=$	$(n + 1, R, q, h \sqcup [R(r) \pm n' \mapsto w])$
$(n, R, q, h) \xrightarrow{t} (n', R', q', h')$	$(n \neq 0 \text{ and matching entry in table})$	
$(n, R, q, h) \xrightarrow{t} (n, R, q, \perp)$	$(n \neq 0 \text{ and no matching entry in table})$	

Instruction z	Condition	Reduction
$\text{jmp } o$	$\llbracket o \rrbracket (R, q, h) = n'$	$\xrightarrow{t} (n', R, q, h)$
$\text{jnz } r \ o$	$R(r) = 0$	$\xrightarrow{t} (n + 1, R, q, h)$
$\text{jnz } r \ o$	$R(r) \neq 0 \wedge \llbracket o \rrbracket (R, q, h) = n'$	$\xrightarrow{t} (n', R, q, h)$
$\text{ld } r \ o$	$\llbracket o \rrbracket (R, q, h) = w$	$\xrightarrow{t} (n + 1, R[r \mapsto w], q, h)$
$\text{sto } o \ r$	$\text{store}(n, R, q, h)(R(r))(o) = \Omega$	$\xrightarrow{t} \Omega$
$\text{lpc } r$		$\xrightarrow{t} (n + 1, R[r \mapsto n], q, h)$
$\text{bop } \odot \ r \ o_1 \ o_2$	$\llbracket o_1 \rrbracket (R, q, h) = n_1 \wedge \llbracket o_2 \rrbracket (R, q, h) = n_2$ $\wedge \text{store}(n, R, q, h)(\llbracket \odot \rrbracket (n_1, n_2))(r) = \Omega$	$\xrightarrow{t} \Omega$
$\text{alloc } r_1 \ r_2$	$\text{alloc}(h, R(r_2)) = w$	$\xrightarrow{t} (n + 1, R[r_1 \mapsto w], q,$ $h \sqcup [w \mapsto 0^{R(r_2)}])$
$\text{input } r$		$\xrightarrow{?i} (n + 1, R[r \mapsto i], q, h)$
$\text{output } r$		$\xrightarrow{!R(r)} (n + 1, R, q, h)$

Figure 4.8: Target language \mathcal{T}

The reduction relation is defined between quadruples consisting of a program counter n , a register file R , a stack q , and a heap h . The program counter always refers to the heap.

The table in Figure 4.8 lists the regular reduction rules and is to be interpreted as follows: (n, R, q, h) steps according to the “Reduction” column if $h(n) = \mathbb{E}(z)$ (*i.e.*, the current instruction is z) and the condition in the second column holds. For instance, if $h(n)$ encodes the load instruction `ld r o` and $\llbracket o \rrbracket(R, q, h) = w$, then $(n, R, q, h) \xrightarrow{c} (n+1, R[r \mapsto w], q, h)$. This rule makes use of the auxiliary function $\llbracket - \rrbracket$ for evaluating the instruction’s operand o , resulting in a word w . Its definition, also given in the figure, is straightforward. A second auxiliary function, *store*, is used to compute the result of a memory write (it is undefined for a numeric operand). Finally, if none of the rules from the table apply, then an error rule can fire, analogous to \mathcal{S} and \mathcal{I} . This is the case, for instance, if the program counter does not point to a valid instruction, or if a store instruction tries to access an invalid memory address. The exception is when the program counter is 0, which we take to be the system continuation used in loading a whole program (see below). For adequacy (Section 4.7.2) we want such machines to be stuck, *i.e.*, not take any step at all, not even an error step.

4.2.5.1 Stack

Stack and heap are separate simply because both are unbounded in size. There are actually no special instructions for operating the stack, just the standard load and store instructions (`ld` and `sto`) in combination with the stack addressing operand $\langle r \pm n \rangle$. We do, however, adopt the convention that the `sp` register functions as the stack pointer (high-water mark).

We define the set of stacks analogous to that of heaps such that it includes \perp . Regarding the operational semantics, this is merely for the sake of uniformity: in fact, we only ever care about stacks that are *fully* defined, *i.e.*, store a value for every address and thus are infinite. This makes sense because the way to “allocate” stack cells is by bumping up the stack pointer and not via a designated instruction whose reduction rule could extend the stack accordingly.

4.2.5.2 Calling Convention

In order to support linking of machine language modules (to be defined in a moment)—no matter their origin—all modules have to agree upon a calling convention. We therefore dictate the following contract for calls to imported functions:

1. Write the function value into register `clo`.
2. Write the argument value into register `arg`.
3. Write the return address into register `ret`.

4. Jump to $[\text{clo} + 0]$, i.e., to wherever the value in the heap at address `clo` points to.

If and when control eventually reaches the return address, the following must hold³:

5. The function’s result resides in register `arg`.
6. Registers `env` and `sp` have been preserved, i.e., these are callee-save registers while the rest are caller-save.
7. The contents of the stack (up to the stack pointer) have been preserved as well.

(For simplicity, we will also follow this convention for intra-module calls, but this is an implementation detail of our compilers.)

4.2.5.3 Modules, Loading, and Linking

We now come to Figure 4.9. An \mathcal{S} or \mathcal{I} module will be translated to a *group* in \mathcal{T} . A group g consists of two components: a *code pointer table* and a *data block*. The table associates each exported function (label) with the index in the data block where its code starts. The data block is just a sequence of words.

If we define \mathcal{T} modules to be such groups, then linking and relocation would become tricky. Instead, we define \mathcal{T} modules to be *collections* of groups. The idea is that the code in each group can assume that its imported functions will at run time be located in memory right before the group’s data block (this will be ensured by the loader), so the code can address them relative to the current program counter. This way, linking can be defined simply as concatenation (think: taking the union) and there is no need for maintaining information about—or performing operations related to—relocation.

Loading requires some work, of course. We first define a recursive auxiliary function *load'*. Here and elsewhere in the figure, some constructions use meta-level standard operations from functional programming, such as *enumerate*, which turns a list a, b, c into a list of pairs $(1, a), (2, b), (3, c)$. *load'* constructs the initial heap (split into a global and local part, so that it can be reused for the implementation of **clload**) and also returns a mapping from the exported labels to their function values. The argument n denotes the address at which to load the module. In the recursive use of *load'*, this address is increased by $n_u + n_c + n_c + n_d$, the amount of heap cells used for the current group g . The last n_d of these cells simply contain g ’s data block (see the second part in the definition of h_2). The preceding cells form the *group header*.

The first n_u cells of the group header (n to $n + n_u - 1$) contain the imported “function values” for g (in the terminology of the calling convention above), i.e., they enable code in g to use its imports. Note that these function values are not the code pointers but *point* to the code pointers. This is to enable a uniform compilation of

³A module may assume these for its imports and must guarantee them for its exports.

module-level and internal functions, the latter of which cannot simply be represented by code pointers because they generally also need to carry an environment.

It remains to explain the $n_c + n_c$ cells in between. Recall that in the source language's module typing rules, when checking a function, there is no formal distinction between those labels that are imports of the module and those that are defined by the module itself ("left" of the function in question) and thus may be referenced by the function. It is also convenient not to make such a distinction when implementing a compiler from \mathcal{S} to \mathcal{T} . For this reason, the loader prepares for each group g a homogeneous environment not only consisting of the group's imports but also of its exports, to be referenced by the group's code. We call this the group's *label environment*. We already discussed the first part of it, namely the first n_u cells of h_2 . Cells $n + n_u$ to $n + n_u + n_c - 1$ make up the second part: they contain the exported function values.

What are these values? They must be the addresses of heap cells containing the corresponding code pointers. These cells need to be placed somewhere, so the loader simply places them right after the label environment and right before the code. This means that cells $n + n_u$ to $n + n_u + n_c - 1$ contain the addresses $n + n_u + n_c$ to $n + n_u + n_c + n_c - 1$ (see h_1), and cells $n + n_u + n_c$ to $n + n_u + n_c + n_c - 1$ contain the code pointers listed in the group's table c (after shifting them by the appropriate offset). The reason for putting the exported function values into the global heap (h_1) and all the rest into the local heap (h_2) will become clear in Section 4.4. For now, this separation should be ignored.

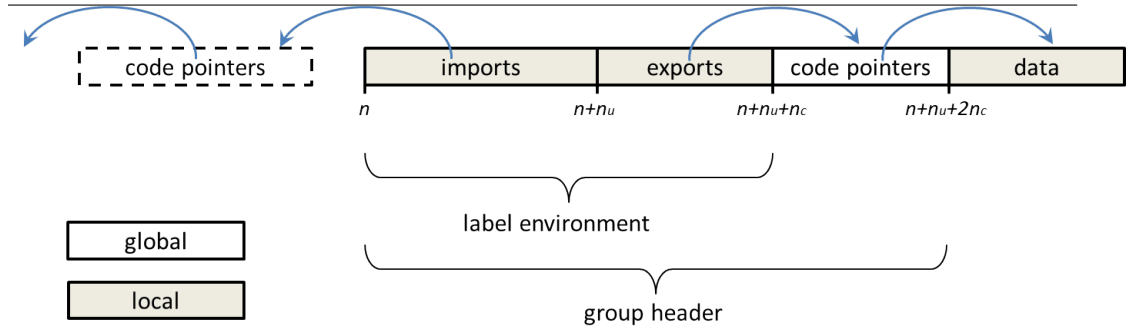
Figure 4.10 sums up the heap layout created by $load'$. The dashed chunk is not included in the global heap created by $load'$ but will be added to it in the definition of **load** in the next section.

We define a basic well-formedness predicate on target modules, $\vdash M : L$. Any module produced by our compilers will satisfy this predicate. It states that L lists the labels that M exports (in the order in which they appear in M 's groups). Moreover, the predicate ensures that the code pointer tables are valid: if (a, b) is a group in M and its table a says that n is the code pointer for export F , then n is a valid index into the data b .

Defining the actual $load$ function is now easy. $load(M)$ uses $load'$ with address 0 (any would work) and empty import list (since M is assumed to be complete) to compute heaps h_1 , h_2 , and export mapping x . It then returns a machine whose heap is the union of h_1 and h_2 (note that these are always valid disjoint heaps); whose stack is filled with zeroes; whose register file is zeroed (R_0) except that clo stores w , the value of the program's main function; and whose program counter is n , the value stored at w (*i.e.*, the main function's code pointer). Notice that this simulates a function call to the main function. Also note that since sp is 0, the stack is *logically* empty. Requiring the stack to *physically* contain only zeroes is an arbitrary choice.

Note that the definition shown in the figure assumes that M provides a main function, *i.e.*, $x(F_{\text{main}}) \neq \perp$, and that its value is a valid heap address, *i.e.*, $h_1(w) \neq \perp$. When that is not the case, then $load(M)$ is not defined either, *i.e.*, $load(M) = \perp$.

$$\begin{aligned}
g \in \text{Group} &:= (\text{Lbl} \times \text{Word})^* \times \text{Word}^* \\
M \in \text{Mod} &:= \text{Group}^* \\
\\
\text{link} &\in \text{Mod} \times \text{Mod} \rightarrow \text{Mod} \\
\text{link}(M_1, M_2) &:= M_1, M_2 \\
\\
\text{load}' &\in \text{Mod} \rightarrow \text{Word} \rightarrow \text{Word}^* \rightarrow \text{Heap} \times \text{Heap} \times (\text{Lbl} \times \text{Word})^* \\
\text{load}'(\epsilon)(n)(u) &:= (\emptyset, \emptyset, \epsilon) \\
\text{load}'(g, M)(n)(u) &:= (h_1 \sqcup h'_1, h_2 \sqcup h'_2, (x, x')) \\
&\text{where } (c, d) := g \\
&\quad n_\rho := \text{length}(\rho) \\
&\quad n_c := \text{length}(c) \\
&\quad n_d := \text{length}(d) \\
&\quad x := \text{enumerate } (\text{map } \text{fst } c) \ (n + n_\rho + n_c) \\
&\quad h_1 := [n + n_u + n_c \mapsto \text{map } (\lambda(F, w).n + n_u + n_c + n_c + w) \ c] \\
&\quad h_2 := [n \mapsto u, \text{map } \text{snd } x] \sqcup [n + n_u + n_c + n_c \mapsto d] \\
&\quad (h'_1, h'_2, x') := \text{load}'(M)(n + n_u + n_c + n_c + n_d)(u, \text{map } \text{snd } x) \\
\\
\text{load} &\in \text{Mod} \rightarrow \text{Word} \times \text{RegFile} \times \text{Stack} \times \text{Heap} \\
\text{load}(M) &:= (n, R_0[\text{clo} \mapsto w], (\lambda_.0), h_1 \sqcup h_2) \\
&\text{where } (h_1, h_2, x) := \text{load}'(M)(0)(\epsilon) \\
&\quad w := x(F_{\text{main}}) \\
&\quad n := h_1(w) \\
\\
\boxed{\vdash M : L} &:= \\
&\quad L = \text{flatten } (\text{map } (\text{map } \text{fst} \circ \text{fst}) M) \wedge \\
&\quad \forall (a, b) \in M. \forall (F, n) \in a. n < \text{length}(b)
\end{aligned}$$

Figure 4.9: Target language \mathcal{T} : ModulesFigure 4.10: Heap layout as set up by load' .

Example To illustrate the heap construction by *load*, let us look at an example. Suppose we have two modules M_a and M_b , each consisting of a single group (*e.g.*, as procuded by one of our compilers):

$$\begin{aligned} M_a &= g_a \\ g_a &= ((F_f, 0), (F_g, 20), d_a) \\ \\ M_b &= g_b \\ g_b &= ((F_{\text{main}}, 0), d_b) \end{aligned}$$

Module M_a exports two functions F_f and F_g , whose code starts at offsets 0 and 20 in d_a , respectively. Module M_b exports one function (the “main” function) F_{main} , whose code starts at offset 0 in d_b . We are interested in the complete program

$$M = \text{link}(M_a, M_b) = g_a, g_b$$

obtained by linking M_a and M_b , so the assumption is that M_a has no imports, while M_b ’s imports are M_a ’s exports.

The result of $\text{load}'(M)(0)(\epsilon)$ is (h_1, h_2, x) , which looks as follows:

$$\begin{aligned} h_1 &= h_1^a \sqcup h_1^b \\ h_2 &= h_2^a \sqcup h_2^b \\ x &= (F_f, 2), (F_g, 3), (F_{\text{main}}, n_a + 7) \\ \\ h_1^a &= [2 \mapsto 4, 24] \\ h_1^b &= [n_a + 7 \mapsto n_a + 8] \\ \\ h_2^a &= [0 \mapsto 2, 3] \sqcup [4 \mapsto d_a] \\ h_2^b &= [n_a + 4 \mapsto 2, 3, n_a + 7] \sqcup [n_a + 8 \mapsto d_b] \\ \\ n_a &= \text{length}(d_a) \end{aligned}$$

For instance, M_b ’s only label environment, used by the code in d_b to look up functions, starts at address $n_a + 4$ and takes up three cells. (Since M_b is the “last” module, this environment’s contents correspond to the export map x .) So if F_{main} wants to call F_g , then F_{main} ’s code in d_b will read the heap at address $n_a + 5$ and find the number 3. Indeed, 3 is F_g ’s function value: the heap at address 3 stores 24, which is F_g ’s code pointer because d_a starts at address 4.

Note that F_g ’s function value, 3, is not known at M_b ’s compile time. On the other hand, d_b knows where to look for F_g because it knows the size of its own environment and thus can access $n_a + 5$ relative to its starting address ($n_a + 8$).

4.2.5.4 Implementation of the Language Specification

Let us now show how we implement the language specification with \mathcal{T} . Values are words (numbers, heap addresses, stack addresses), anchors are words (heap ad-

Val	$:=$	Word
Anch	$:=$	Word
Mod	$::=$	Mod
Cont	$:=$	Word
Mach	$:=$	Word \times RegFile \times Stack \times Heap
Conf	$:=$	Word $_{\perp, \emptyset} \times$ RegFile $_{\perp, \emptyset} \times$ Stack \times Heap
Conf monoid:		
\emptyset	$:=$	$(\emptyset, \emptyset, \emptyset, \emptyset)$
$(n, R, q, h) \cdot (n', R', q', h')$	$:=$	$(n \oplus n', R \oplus R', q \sqcup q', h \sqcup h')$
$\hookrightarrow := \hookrightarrow$		
real (c)	$:=$	$\{m \mid m = c \wedge m.\text{hp} \neq \perp \wedge m.\text{hp} \text{ finite} \wedge m.\text{st} \neq \perp\}$
vload (M)(n)(ρ)(F)	$:=$	$\{v \mid \exists h_1, h_2, x.$ $\text{load}'(M)(n)(\text{map } \text{snd } \rho) = (h_1, h_2, x) \wedge x(F) = v\}$
clload (M)(n)(ρ)	$:=$	$\{(c_1, c_2) \mid \exists h_1, h_2, x, R, h.$ $\text{load}'(M)(n)(\text{map } \text{snd } \rho) = (h_1, h_2, x) \wedge \text{dom}(h) = \{w \mid \exists F. (F, w) \in \rho\} \wedge$ $h \sqcup h_1 \sqcup h_2 \neq \perp \wedge R(\text{sp}) = 0 \wedge c_1 = (\emptyset, R, (\lambda_.0), h \sqcup h_1) \wedge c_2 = (\emptyset, \emptyset, \emptyset, h_2)\}$

Figure 4.11: Target language \mathcal{T} : Implementation of the Language Specification

dresses), and continuations are words (code pointers, *i.e.*, heap addresses). The configuration monoid and **real** are defined mostly analogously to the previous languages (heap and stack can be split, the rest cannot).

Both the implementation of **vload** and that of **clload** make use of load' ⁴. **vload** simply looks up the given label F in the generated export map x . **clload** relies on the other two components instead, the global heap h_1 and the local heap h_2 . It allows for any global configuration c_1 whose program counter is missing; whose register file is defined and stores 0 in the stack pointer register **sp**; whose stack is constant zero (matching load); and whose heap is h_1 plus some disjoint h . Intuitively, h contains the imported function values (this will become clearer once we see how **clload** is used). Accordingly, **clload** requires that h 's domain corresponds to the import map ρ ; it must also be disjoint from h_2 , but apart from that, h can be arbitrary. The local configuration, c_2 , is restricted to be empty except for its heap component, which must be exactly h_2 .

⁴Note that load' does not care if the module is complete or not.

4.3 Worlds

Like in PBs (Chapter 3), worlds play a central role in PILS. Before, we distinguished between local worlds and full worlds. A local world was constructed on a per-proof basis, and turned into a full world by linking it with a predefined shared world. This basic story is still true for PILS. However, we have to generalize the notion of shared world somewhat. Instead of only governing global references, the shared world’s purpose is now also to formalize the calling convention and data representation that all modules have to agree upon. Moreover, since we are dealing with multiple languages, instead of having a single shared world, we will have one for each language pair of interest: we will define the global worlds $\Omega_{\mathcal{TL}}$, $\Omega_{\mathcal{IL}}$, $\Omega_{\mathcal{IS}}$, and $\Omega_{\mathcal{TS}}$.

As previously mentioned, we are going to define PILS generically, *i.e.*, parameterized by two languages A and B and a global world for them, and we are going to have two versions of this generic definition: one where the relations are type-indexed, and one where they are not. We present both at the same time (in several figures to come) by showing the typed version but highlighting (in **color**) all the type-related components that simply need to be erased in order to obtain the untyped version.

The structure of worlds is shown in Figure 4.12, following some preliminary definitions. First, the set of types used by the typed versions of the model is the same as for PBs, because \mathcal{S} ’s type language is the same as $F^{\mu!}$ ’s. Recall that we extend this language with type names \mathbf{n} . The set CTyF of (closed) flexible types is slightly different from before: only function types and type names are considered flexible, reference types are not. In PBs, the treatment of reference types as flexible was convenient, but did not make so much sense conceptually since local worlds could not influence their interpretation. The generalization of PBs to PILS naturally leads to a cleaner treatment of reference types.

Global worlds $\Omega \in \text{GWorld}_{A,B}$ generalize PBs’s shared world W_{ref} (from Section 3.6.3.1). They consist of a transition system \mathbf{T} , a *configuration relation* \mathbf{C} , and a *query component* \mathbf{Q} . The notion of transition systems is the same as before (TrSys was defined in Section 3.6.1). Configuration relations generalize PBs’s heap relations. The query component is completely new; it provides several methods for “querying” the global state (see below). We write $A.\mathbf{Val}$ to denote language A ’s implementation of the abstract notion **Val**, and so on.

Full worlds $W \in \text{World}_{A,B}$ look like extensions of global worlds. They additionally contain an *algebraic well-founded ordered set* \mathbf{O} , which will be used to facilitate *stuttering* in proofs (Section 4.5.1). For now, it can be safely ignored.

In the typed setting only, full worlds also contain a *type name interpretation* \mathbf{N} . This component provides a set \mathbf{NS} of owned type names and a relational interpretation \mathbf{NR} of these names (in PBs, this was part of the local knowledge). Local worlds $w \in \text{LWorld}_{A,B}^T$ are like full worlds except that they don’t contain the global query handlers and that their configuration relation must only relate frames. As before, local worlds can depend on the state of the global world. We will see the lifting of local worlds to full worlds in a moment.

$\mathbf{n} \in \text{TyNam}$	
$\sigma \in \text{Ty}$	$::= \alpha \mid \mathbf{n} \mid \text{unit} \mid \text{nat} \mid \sigma_1 \times \sigma_2 \mid \sigma_1 + \sigma_2 \mid \sigma_1 \rightarrow \sigma_2 \mid \forall \alpha. \sigma \mid \exists \alpha. \sigma \mid \mu \alpha. \sigma \mid \text{ref } \sigma$
CTy	$::= \{ \sigma \in \text{Ty} \mid \text{fv}(\sigma) = \emptyset \}$
CTyF	$::= \{ (\tau_1 \rightarrow \tau_2) \in \text{CTy} \} \cup \{ (\forall \alpha. \tau) \in \text{CTy} \} \cup \{ \mathbf{n} \in \text{CTy} \}$
$\text{VRelF}_{A,B}$	$::= \text{CTyF} \rightarrow \mathcal{P}(A.\text{Val} \times B.\text{Val})$
$\text{VRel}_{A,B}$	$::= \text{CTy} \rightarrow \mathcal{P}(A.\text{Val} \times B.\text{Val})$
$\text{CRel}_{A,B}$	$::= \mathcal{P}(A.\text{Conf} \times B.\text{Conf})$
$\text{MRel}_{A,B}$	$::= \mathcal{P}(A.\text{Mach} \times B.\text{Mach})$
$\text{ERel}_{A,B}$	$::= \text{CTy} \rightarrow \text{CRel}_{A,B}$
$\text{KRel}_{A,B}$	$::= \text{CTy} \rightarrow \text{CTy} \rightarrow \mathcal{P}(A.\text{Cont} \times B.\text{Cont})$
Queries ($\mathbf{v}, \mathbf{v}' \in L.\text{Val}$ and $\mathbf{k} \in L.\text{Cont}$):	
$vq \in \text{VQry}_L$	$::= \text{unit} \mid \text{nat } n \mid \text{pair } \mathbf{v} \mathbf{v}' \mid \text{inl } \mathbf{v} \mid \text{inr } \mathbf{v} \mid \text{roll } \mathbf{v} \mid \text{pack } \mathbf{v} \mid \text{fix} \mid \text{gen} \mid \text{goodfix} \mid \text{goodgen}$
$cq \in \text{CQry}_L$	$::= \text{app } \mathbf{v} \mathbf{v}' \mathbf{k} \mid \text{inst } \mathbf{v} \mathbf{k} \mid \text{ret } \mathbf{v} \mathbf{k}$
$\text{QH}_{A,B}^T$	$::= \{ (\text{vqh}_a \in T.S \xrightarrow{\text{mon}} \text{VQry}_A \rightarrow \mathcal{P}(A.\text{Val})$ $\quad, \text{vqh}_b \in T.S \xrightarrow{\text{mon}} \text{VQry}_B \rightarrow \mathcal{P}(B.\text{Val})$ $\quad, \text{cqh}_a \in T.S \rightarrow \text{CQry}_A \rightarrow \mathcal{P}(A.\text{Conf})$ $\quad, \text{cqh}_b \in T.S \rightarrow \text{CQry}_B \rightarrow \mathcal{P}(B.\text{Conf})$ $\quad, \text{rqh} \in T.S \xrightarrow{\text{mon}} \text{VRel}_{A,B}) \}$
$\text{CR}_{A,B}^T$	$::= (T.S \rightarrow \text{VRelF}_{A,B}) \xrightarrow{\text{mon}} T.S \rightarrow \text{CRel}_{A,B}$
$\text{LCR}_{A,B}^T$	$::= (T.S \rightarrow \text{VRelF}_{A,B}) \xrightarrow{\text{mon}} T.S \rightarrow \mathcal{P}(A.\text{frame} \times B.\text{frame})$
$\text{NI}_{A,B}^T$	$::= \{ (\text{NS} \in \mathcal{P}(\text{TyNam})$ $\quad, \text{NR} \in (T.S \rightarrow \text{VRelF}_{A,B}) \xrightarrow{\text{mon}} T.S \xrightarrow{\text{mon}} \text{TyNam} \rightarrow \mathcal{P}(A.\text{Val} \times B.\text{Val})) \mid$ $\quad \forall G, s. \forall \mathbf{n} \notin \text{NS}. \text{NR}(G)(s)(\mathbf{n}) = \emptyset \}$
$\text{GWorld}_{A,B}$	$::= \{ (T \in \text{TrSys}, C \in \text{CR}_{A,B}^T, Q \in \text{QH}_{A,B}^T) \}$
$\text{LWorld}_{A,B}^T$	$::= \{ (T \in \text{TrSys}, C \in \text{LCR}_{A,B}^{T \times T}, \mathbf{N} \in \text{NI}_{A,B}^{T \times T}, O \in \text{AWFO}) \}$
$\text{World}_{A,B}$	$::= \{ (T \in \text{TrSys}, C \in \text{CR}_{A,B}^T, Q \in \text{QH}_{A,B}^T, \mathbf{N} \in \text{NI}_{A,B}^T, O \in \text{AWFO}) \}$

Figure 4.12: Worlds

At this point, the reader is probably wondering where the local knowledge component has gone (apart from the interpretation of type names), which played such a crucial role in PBs. We have seen in Section 3.11 how one can in principle always pick the local knowledge to be the greatest consistent relation. Here we go one step further by getting rid of the local knowledge component in worlds altogether. We will explain this in detail in Section 4.5.

4.3.1 Queries

Let us talk about a (global) world’s query component. It provides five query methods, as shown in the definition of $\text{QH}_{A,B}^T$ in Figure 4.12. Part of their purpose is to act as the interface to access information stored in the global world’s state. We will see later how they are used exactly but we want to give some intuition here.

Configuration Queries Recall that PBs’s **RET** case asserts that the two computations have finished and returned similar values (*e.g.*, Section 3.5). In our source language \mathcal{S} , termination is a syntactic property and is trivial to check. But what does it mean for a low-level \mathcal{T} computation to be “finished”?

In order to define PILS in a language-generic fashion, we require the global world to provide a way of answering such a question. This is done in the form of two query handlers cqh_a (for the “output” language A) and cqh_b (for the input language B). There are three queries $cq \in \text{CQry}_L$ concerning the configurations of a language L . The query $\text{app } f \mathbf{v} \mathbf{k}$ asks for the configurations that currently represent a call of function f with argument \mathbf{v} and continuation \mathbf{k} ; $\text{inst } g \mathbf{k}$ asks for the configurations that represent an instantiation of generalization g with continuation \mathbf{k} ; and $\text{ret } \mathbf{v} \mathbf{k}$ asks for those that represent a return of value \mathbf{v} to continuation \mathbf{k} . The reason why we care about continuations here as well is that in low-level languages such as \mathcal{T} , we cannot tell whether a computation has returned without knowing its initial continuation (think: return address).

In effect, this means that the global world’s implementation of cqh_a and cqh_b determines a major aspect of the calling conventions.

Value Queries Recall that PBs’s **RET** and **STEP/CALL** cases require that certain values (*e.g.*, the returned results) are related by the value closure of the global knowledge. The value closure is a straightforward lifting of a relation from functions to other value forms, *e.g.*, by saying that two pairs are related iff their first projections are related (recursively) and their second projections as well.

In order to define the value closure operation generically in the PILS setting, we first need to have a uniform way of determining how \mathcal{S} ’s value forms are represented by the two languages A and B . This is easy for our source and intermediate language: given a value, it is trivial to tell whether it represents, say, a certain pair. On the machine language side, though, values are just machine words, and whether or not

$$\begin{aligned}
\langle\langle - \rangle\rangle^{(-)} &\in \mathbf{VRelF}_{A,B} \rightarrow T.S \rightarrow \mathbf{VRel}_{A,B} \quad (\text{implicit: } A, B; T \in \mathbf{TrSys}; Q \in \mathbf{QH}_{A,B}^T) \\
\langle\langle R \rangle\rangle^s &:= \{ (\tau \rightarrow \tau', \mathbf{v}_a, \mathbf{v}_b) \in R \mid \mathbf{v}_a \in Q.\mathbf{vqh}_a(s)(\text{fix}) \wedge \mathbf{v}_b \in Q.\mathbf{vqh}_b(s)(\text{fix}) \} \\
&\cup \{ (\forall \alpha. \tau, \mathbf{v}_a, \mathbf{v}_b) \in R \mid \mathbf{v}_a \in Q.\mathbf{vqh}_a(s)(\text{gen}) \wedge \mathbf{v}_b \in Q.\mathbf{vqh}_b(s)(\text{gen}) \} \\
&\cup \{ (\mathbf{n}, \mathbf{v}_a, \mathbf{v}_b) \in R \mid \mathbf{v}_b \in Q.\mathbf{vqh}_b(s)(\text{name}) \} \\
&\cup \{ (\text{unit}, \mathbf{v}_a, \mathbf{v}_b) \mid \mathbf{v}_a \in Q.\mathbf{vqh}_a(s)(\text{unit}) \wedge \mathbf{v}_b \in Q.\mathbf{vqh}_b(s)(\text{unit}) \} \\
&\cup \{ (\text{nat}, \mathbf{v}_a, \mathbf{v}_b) \mid \exists n. \mathbf{v}_a \in Q.\mathbf{vqh}_a(s)(\text{nat } n) \wedge \mathbf{v}_b \in Q.\mathbf{vqh}_b(s)(\text{nat } n) \} \\
&\cup \{ (\tau_1 \times \tau_2, \mathbf{v}_a, \mathbf{v}_b) \mid \exists \mathbf{v}_a^1, \mathbf{v}_a^2, \mathbf{v}_b^1, \mathbf{v}_b^2. (\mathbf{v}_a^1, \mathbf{v}_b^1) \in \langle\langle R \rangle\rangle^s(\tau_1) \wedge (\mathbf{v}_a^2, \mathbf{v}_b^2) \in \langle\langle R \rangle\rangle^s(\tau_2) \wedge \\
&\quad \mathbf{v}_a \in Q.\mathbf{vqh}_a(s)(\text{pair } \mathbf{v}_a^1 \mathbf{v}_a^2) \wedge \mathbf{v}_b \in Q.\mathbf{vqh}_b(s)(\text{pair } \mathbf{v}_b^1 \mathbf{v}_b^2) \} \\
&\cup \{ (\tau_1 + \tau_2, \mathbf{v}_a, \mathbf{v}_b) \mid \exists \mathbf{v}_a^1, \mathbf{v}_b^1. (\mathbf{v}_a^1, \mathbf{v}_b^1) \in \langle\langle R \rangle\rangle^s(\tau_1) \wedge \\
&\quad \mathbf{v}_a \in Q.\mathbf{vqh}_a(s)(\text{inl } \mathbf{v}_a^1) \wedge \mathbf{v}_b \in Q.\mathbf{vqh}_b(s)(\text{inl } \mathbf{v}_b^1) \} \\
&\cup \{ (\tau_1 + \tau_2, \mathbf{v}_a, \mathbf{v}_b) \mid \exists \mathbf{v}_a^2, \mathbf{v}_b^2. (\mathbf{v}_a^2, \mathbf{v}_b^2) \in \langle\langle R \rangle\rangle^s(\tau_2) \wedge \\
&\quad \mathbf{v}_a \in Q.\mathbf{vqh}_a(s)(\text{inr } \mathbf{v}_a^2) \wedge \mathbf{v}_b \in Q.\mathbf{vqh}_b(s)(\text{inr } \mathbf{v}_b^2) \} \\
&\cup \{ (\mu \alpha. \tau, \mathbf{v}_a, \mathbf{v}_b) \mid \exists \mathbf{v}'_a, \mathbf{v}'_b. (\mathbf{v}'_a, \mathbf{v}'_b) \in \langle\langle R \rangle\rangle^s(\tau[\mu \alpha. \tau / \alpha]) \wedge \\
&\quad \mathbf{v}_a \in Q.\mathbf{vqh}_a(s)(\text{roll } \mathbf{v}'_a) \wedge \mathbf{v}_b \in Q.\mathbf{vqh}_b(s)(\text{roll } \mathbf{v}'_b) \} \\
&\cup \{ (\text{ref } \tau, \mathbf{v}_a, \mathbf{v}_b) \mid (\mathbf{v}_a, \mathbf{v}_b) \in Q.\mathbf{rqh}(s)(\tau) \}
\end{aligned}$$

$$\overline{G}(s) := \langle\langle G(s) \rangle\rangle^s$$

Figure 4.13: Value Closure.

a given word represents a pair typically very much depends on the memory in which it is considered (and similarly for other forms).

We therefore require that the global world provide a mechanism to answer questions such as: does value \mathbf{v} currently represent a pair of values \mathbf{v}_1 and \mathbf{v}_2 ? This mechanism takes the form of two value query handlers \mathbf{vqh}_a and \mathbf{vqh}_b , very much like the configuration query handlers that we have just seen. They take a state and a value query $vq \in \mathbf{VQry}_L$ (for $L \in \{A, B\}$), and return a set of values. The above example translates into $\mathbf{v} \in \mathbf{vqh}_a(s)(\text{pair } \mathbf{v}_1 \mathbf{v}_2)$, for instance. Crucially, these handlers must be monotone in their state argument w.r.t. \sqsubseteq , reflecting the immutability of values (e.g., “once a pair, always a pair”).

For global reference values, the global world also needs to keep track of which A locations correspond to which B locations (in the two modules one reasons about). Recall the global world W_{ref} from PBs, where a state s_{ref} was a partial bijection between heap locations. In PILS, a global world must provide \mathbf{rqh} for querying this correspondence. (Essentially, each of the global worlds that we will define, will still carry such a partial bijection in its state, and \mathbf{rqh} is simply the interface to it.) Actually, since we will never be interested in asking whether some value is a reference without also asking about its counterpart, the language of value queries (\mathbf{VQry}_L) mentioned above does not include a construct for references.

4.3.2 Value Closure

With the help of these three query handlers, the generalization of PBs’s value closure operator is fairly straightforward. It is shown in Figure 4.13⁵ as $\langle\langle - \rangle\rangle^{(-)}$. Like before, it is defined as a least fixed point. Unlike before, however, it now takes the world’s transition system T and query method collection Q (both implicitly) as well as the current state s (explicitly) as arguments, so that it can invoke Q ’s query methods on s to ensure that the values have the right “shape”. Note that references are dealt with using $Q.\text{rgh}$, and all the rest using $Q.\text{vqh}_a$ and $Q.\text{vqh}_b$. There are a few points to note.

- Admittedly, the name “value closure” is no longer justified since $\langle\langle R \rangle\rangle^s$ is not always a superset of R . We nevertheless keep the name for the sake of consistency.

Related to this, we will see in the later parts of this section that PILS are set up such that whenever we take related values from the global knowledge, we do so via its value closure—never directly from the unfiltered global knowledge (in contrast to PBs).

- The reason why the operator does not simply take the whole global world as implicit argument is that in the definition of our particular global worlds (Section 4.4), we already want to make use of the value closure operator—hence there would be a cycle.
- As a convenient shorthand, we write $\overline{G}(s)$ for $\langle\langle G(s) \rangle\rangle^s$, *i.e.*, the value closure over $G(s)$ relative to state s .
- The asymmetry in the case of type names seems odd: we require \mathbf{v}_b to satisfy the name query as expected, but we require nothing of \mathbf{v}_a . What is **name** anyway?

Like in PBs, type names allow us to attach relational interpretations to abstract types, and it is important that these relations can be essentially arbitrary. In our PILS transitivity proof, it turns out that it is also important that the *untyped* value closure has a “type name” case, namely when showing something along the lines of $\overline{G}(\mathbf{n}) \subseteq \overline{G}_1 \circ \overline{G}_2(\mathbf{n})$ (see Lemma 64 in Section 4.8). There we must be able to somehow decompose the relational interpretation of \mathbf{n} . However, we cannot simply add an unrestricted case of the form

$$\cup \{ (\mathbf{n}, \mathbf{v}_a, \mathbf{v}_b) \in R \}$$

to the value closure, because then inverting \overline{G} , in order to reason about programs that use related values, would force us to deal with the ever-present possibility of essentially not knowing anything at all about these values.

⁵Here and in other places, we sometimes write *W.S* short for *W.T.S.*, and similar for other components when there is no ambiguity.

To avoid this hopeless situation, we introduced the **name** query and require in the value closure's type name case that the left-side values satisfy it. Moreover, we defined \mathcal{T} 's value query handler such that it only admits a very particular class of values (cf. Figure 4.18). These values suffice for the purposes of our transitivity proof, and, crucially, they cannot ever be successfully used: any attempt to deconstruct them results in an error (this implies that they are distinct from all other value forms occurring in the value closure). This knowledge makes it easy to deal with the artificial type name case in proofs.

- The reader may also wonder about the two value queries that are not used by the value closure operator: **goodfix** and **goodgen**. These are the PILS versions of PBs's **FixVal** and **GenVal**. They will be imposed on what the PILS substitute for PBs's local knowledge is. Like before, their purpose is to ensure that certain relational compositions that come up in the proof of transitivity are empty. This will become clear in Section 4.8.2.

4.3.3 Lifting and Separating Conjunction of Local Worlds

Figure 4.14 shows the PILS versions of lifting and separating conjunction of local worlds (the PB versions were defined in Section 3.6.4). All definitions in this figure take as implicit arguments two languages A and B ; some also implicitly take the global world $\Omega \in \text{GWorld}_{A,B}$ or just a transition system T . We will make sure that it is clear from the context how these implicit arguments are instantiated when using the operators later on.

The first construct in the figure is the separating conjunction \otimes , generalizing \otimes (Figure 3.14) from heap relations to configuration relations.

Next is the lifting operator. It composes its argument w with the global world Ω much like in PBs. Note how the resulting full world $w\uparrow$ forwards any queries to Ω . Also note that, when accessing Ω 's configuration relation $\Omega.C$, we need to pass it a global knowledge indexed by $\Omega.T$ states, but the one we are given (G) is indexed by $w\uparrow.T$ states; thus, instead of passing G , we pass $G(-, s)$. (This does not apply to $w.C$, because it uses the product transition system anyways.)

The product \otimes of local worlds is defined iff their name sets are disjoint, like for PBs. The definition is straightforward except for one point: both in the construction of the product's configuration relation and its name interpretation, we are given a G indexed by $T \times (w_1 \otimes w_2).T$ states, but need to pass one indexed by $T \times w_1.T$ states to w_1 , and one indexed by $T \times w_2.T$ states to w_2 . We would like to pass $\lambda(s'_g, s'_1). G(s'_g, (s'_1, s_2))$ to w_1 (and similarly for w_2), but this is not guaranteed to be monotone in s_2 . Consequently, we wouldn't be able to show the required monotonicity of the composed NR. For this reason, we first monotonize the global knowledge constructions, written $G_{\langle 1 \rangle}^{s_2}$ and $G_{\langle 2 \rangle}^{s_1}$. This monotization is harmless because ultimately we are only interested in monotone G anyways.

$$\begin{aligned}
(-) \otimes (-) &\in \mathbf{CRel}_{A,B} \rightarrow \mathbf{CRel}_{A,B} \rightarrow \mathbf{CRel}_{A,B} \\
R \otimes R' &:= \{ (c_1 \cdot c'_1, c_2 \cdot c'_2) \mid (c_1, c_2) \in R \wedge (c'_1, c'_2) \in R' \} \\
\\
(-) \uparrow &\in \mathbf{LWorld}_{A,B}^{\Omega, \mathbf{T}} \rightarrow \mathbf{World}_{A,B} \\
w \uparrow. \mathbf{T} &:= \Omega. \mathbf{T} \times w. \mathbf{T} \\
w \uparrow. \mathbf{Q.vqh}_a(s_g, s) &:= \Omega. \mathbf{Q.vqh}_a(s_g) \\
w \uparrow. \mathbf{Q.vqh}_b(s_g, s) &:= \Omega. \mathbf{Q.vqh}_b(s_g) \\
w \uparrow. \mathbf{Q.cqh}_a(s_g, s) &:= \Omega. \mathbf{Q.cqh}_a(s_g) \\
w \uparrow. \mathbf{Q.cqh}_b(s_g, s) &:= \Omega. \mathbf{Q.cqh}_b(s_g) \\
w \uparrow. \mathbf{Q.rqh}(s_g, s) &:= \Omega. \mathbf{Q.rqh}(s_g) \\
w \uparrow. \mathbf{C}(G)(s_g, s) &:= \Omega. \mathbf{C}(G(-, s))(s_g) \otimes w. \mathbf{C}(G)(s_g, s) \\
w \uparrow. \mathbf{N} &:= w. \mathbf{N} \\
w \uparrow. \mathbf{O} &:= w. \mathbf{O} \\
\\
(-)_{\langle 1 \rangle}^{(-)} &\in ((\Omega. \mathbf{T} \times (T_1 \times T_2)). \mathbf{S} \rightarrow \mathbf{VRelF}_{A,B}) \rightarrow T_2. \mathbf{S} \rightarrow (\Omega. \mathbf{T} \times T_1). \mathbf{S} \rightarrow \mathbf{VRelF}_{A,B} \\
G_{\langle 1 \rangle}^{s_2}(s_g, s_1) &:= \bigcup \{ G(s_g, (s_1, s'_2)) \mid s'_2 \sqsubseteq s_2 \} \\
\\
(-)_{\langle 2 \rangle}^{(-)} &\in ((\Omega. \mathbf{T} \times (T_1 \times T_2)). \mathbf{S} \rightarrow \mathbf{VRelF}_{A,B}) \rightarrow T_1. \mathbf{S} \rightarrow (\Omega. \mathbf{T} \times T_2). \mathbf{S} \rightarrow \mathbf{VRelF}_{A,B} \\
G_{\langle 2 \rangle}^{s_1}(s_g, s_2) &:= \bigcup \{ G(s_g, (s'_1, s_2)) \mid s'_1 \sqsubseteq s_1 \} \\
\\
(-) \otimes (-) &\in \mathbf{LWorld}_{A,B}^T \rightarrow \mathbf{LWorld}_{A,B}^T \rightarrow \mathbf{LWorld}_{A,B}^T \\
(w_1 \otimes w_2). \mathbf{T} &:= w_1. \mathbf{T} \times w_2. \mathbf{T} \\
(w_1 \otimes w_2). \mathbf{C}(G)(s_g, (s_1, s_2)) &:= \Omega. \mathbf{C}(G_{\langle 1 \rangle}^{s_2})(s_g, s_1) \otimes w. \mathbf{C}(G_{\langle 2 \rangle}^{s_1})(s_g, s_2) \\
(w_1 \otimes w_2). \mathbf{N.NS} &:= w_1. \mathbf{N.NS} \uplus w_2. \mathbf{N.NS} \\
(w_1 \otimes w_2). \mathbf{N.NR}(G)(s_g, (s_1, s_2)) &:= w_1. \mathbf{N.NR}(G_{\langle 1 \rangle}^{s_2})(s_g, s_1) \cup w_2. \mathbf{N.NR}(G_{\langle 2 \rangle}^{s_1})(s_g, s_2) \\
(w_1 \otimes w_2). \mathbf{O} &:= w_1. \mathbf{O} \times w_2. \mathbf{O}
\end{aligned}$$

Figure 4.14: Lifting and separating conjunction of local worlds.

$$\begin{aligned}
\Omega^{AB}.\top &:= \top^A \times \top_{\text{ref}}^{AB} \times \top^B \\
\Omega^{AB}.\text{Q.vqh}_a(s_a, s_{\text{ref}}, s_b) &:= \text{vqh}^A(s_a) \\
\Omega^{AB}.\text{Q.vqh}_b(s_a, s_{\text{ref}}, s_b) &:= \text{vqh}^B(s_b) \\
\Omega^{AB}.\text{Q.cqh}_a(s_a, s_{\text{ref}}, s_b) &:= \text{cqh}^A(s_a) \\
\Omega^{AB}.\text{Q.cqh}_b(s_a, s_{\text{ref}}, s_b) &:= \text{cqh}^B(s_b) \\
\Omega^{AB}.\text{Q.rqh}(s_a, s_{\text{ref}}, s_b) &:= \text{rqh}^{AB}(s_{\text{ref}}) \\
\Omega^{AB}.\text{C}(G)(s_a, s_{\text{ref}}, s_b) &:= (\text{P}^A(s_a) \times \{B.\emptyset\}) \otimes \text{C}_{\text{ref}}^{AB}(G)(s_a, s_{\text{ref}}, s_b) \otimes (\{A.\emptyset\} \times \text{P}^B(s_b))
\end{aligned}$$

Figure 4.15: Schematic Definition of Concrete Global Worlds.

4.4 Concrete Global Worlds

As we are going to see in Section 4.5, PILS define at the top-level a generic *module similarity*, written \lesssim_Ω . This basically generalizes PBs open program equivalence (here we consider modules because they are the objects that we compile and link). This relation is relative to a global world $\Omega \in \text{GWorld}_{A,B}$ for the language pair A, B , with A intuitively being the language *to* which is being translated and B the one *from* which is being translated. Eventually we will instantiate this relation with four global worlds in order to obtain the four models of interest. In this section, we define these four global worlds:

- $\Omega_{\mathcal{TS}} \in \text{GWorld}_{\mathcal{T},\mathcal{S}}$ (typed)
- $\Omega_{\mathcal{TI}} \in \text{GWorld}_{\mathcal{T},\mathcal{I}}$ (untyped)
- $\Omega_{\mathcal{II}} \in \text{GWorld}_{\mathcal{I},\mathcal{I}}$ (untyped)
- $\Omega_{\mathcal{IS}} \in \text{GWorld}_{\mathcal{I},\mathcal{S}}$ (typed)

Some details in this section may be hard to understand without familiarity with the actual PILS relations that will be discussed in Section 4.5. It can therefore be helpful to revisit this section at a later point. (We did consider a different order of presentation but prefer the current one.)

To avoid duplicate work, we define all four global worlds following the schema shown in Figure 4.15. Accordingly, each global world consists of three parts: one solely regarding the output language A , one solely regarding the input language B , and one regarding both.

4.4.1 Global References

The last of these deals with global references. It consists of a transition system \top_{ref}^{AB} , a configuration relation $\text{C}_{\text{ref}}^{AB}$ and a reference query handler rqh^{AB} , the implementation of $\Omega^{AB}.\text{Q.rqh}$. These are defined in Figure 4.16 and closely follow the construction of W_{ref} in PBs: states are partial bijections on locations that can grow over time,

and the configuration relation $\mathbf{C}_{\text{ref}}^{AB}$ ensures that these locations are allocated in the heaps and store related values. Reference queries are resolved by rqh^{AB} via simple membership tests on the state.

$\mathbf{T}_{\text{ref}}^{AB}$ comes in a typed and an untyped version (recall that the untyped one is obtained by erasing the type-related components and conditions, which are highlighted in *color*). For $\Omega_{\mathcal{TS}}$ and $\Omega_{\mathcal{IS}}$, we pick the typed $\mathbf{T}_{\text{ref}}^{\mathcal{TS}}$ and $\mathbf{T}_{\text{ref}}^{\mathcal{IS}}$, respectively, since they involve the typed source language. For $\Omega_{\mathcal{TI}}$ and $\Omega_{\mathcal{IL}}$, we pick the untyped $\mathbf{T}_{\text{ref}}^{\mathcal{TI}}$ and $\mathbf{T}_{\text{ref}}^{\mathcal{IL}}$, respectively. (This is reflected in the concrete definitions of the different \mathbf{C}_{ref} implementations, namely in the quantification over elements of s_{ref}). We pick the versions of rqh^{AB} analogously.

4.4.2 Unary Parts

The unary part of Ω^{AB} concerning language A consists of a transition system \mathbf{T}^A , query handlers vqh^A and cqh^A (implementing $\Omega^{AB}.\text{vqh}_a$ and $\Omega^{AB}.\text{cqh}_a$, respectively), and a configuration predicate \mathbf{P}^A . We will discuss these in a moment. The unary part concerning language B is analogous. The overall configuration relation $\Omega^{AB}.\mathbf{C}$ then is the separating conjunction of three parts: $\mathbf{C}_{\text{ref}}^{AB}$; the lifting of \mathbf{P}^A by pairing each element with an empty B configuration; and the analogous lifting of \mathbf{P}^B . It remains to define \mathbf{T}^L , vqh^L and cqh^L , for $L \in \{\mathcal{S}, \mathcal{I}, \mathcal{T}\}$.

4.4.2.1 Global World Components Regarding \mathcal{S}

Figure 4.17 shows these constructions for the source language \mathcal{S} . States are the run-time environments used by the operational semantics (mapping labels to values, see Section 4.2.3). It makes sense that the global world “owns” the environment because it is shared by all modules constituting the program. In reflectance of the semantics, the transition relation for these states is trivial: environments are not allowed to change in any way, not even to grow.

The value and configuration query handlers are purely syntactic (they do not care about the state) and mostly straightforward. For instance, value v represents a pair of v_1 and v_2 iff it is *the* pair value $\langle v_1, v_2 \rangle$. Note the following points:

- There are no restrictions on inhabitants of type names (we allow the full set of values).
- Similarly for `goodfix` and `goodgen`, which only matter for the intermediate language.
- We find it convenient to restrict the configurations accepted by $\text{cqh}_{\mathcal{S}}$ to cores (Section 4.2.2). We will do so also for the other languages.
- Our characterization of application (and instantiation) is actually “one step ahead”: what we characterize is not the beta redex but rather the result of performing the beta reduction. This is done for consistency with how we set things up for the machine language in Section 4.4.2.3 below.

$$\begin{aligned}
T_{\text{ref}}^{AB}.S &:= \{ s_{\text{ref}} \in \mathcal{P}_{\text{fin}}(\text{CTy} \times \text{Loc}_A \times \text{Loc}_B) \mid \forall (\tau, l_1, l_2) \in s_{\text{ref}}. \forall (\tau', l'_1, l'_2) \in s_{\text{ref}}. \\
&\quad (l_1 = l'_1 \implies \tau = \tau' \wedge l_2 = l'_2) \wedge (l_2 = l'_2 \implies \tau = \tau' \wedge l_1 = l'_1) \} \\
T_{\text{ref}}^{AB}.\sqsubseteq &:= \sqsubseteq \\
T_{\text{ref}}^{AB}.\sqsubseteq_{\text{pub}} &:= \sqsubseteq \\
\text{rqh}^{AB}(s_{\text{ref}})(\tau) &:= \{ (l_a, l_b) \mid (\tau, l_a, l_b) \in s_{\text{ref}} \} \\
C_{\text{ref}}^{AB}(G)(s_a, s_{\text{ref}}, s_b) &:= C_{\text{ref}}^{AB}(s_{\text{ref}})(\overline{G}(s_a, s_{\text{ref}}, s_b)) \\
C_{\text{ref}}^{TS}(s_{\text{ref}})(R) &:= \{ ((\emptyset, \emptyset, \emptyset, h_a), (h_b, \emptyset, \emptyset)) \mid h_a \neq \perp \wedge h_b \neq \perp \wedge \\
&\quad \text{dom}(h_a) = \{ l_a \mid \exists \tau, l_b. (\tau, l_a, l_b) \in s_{\text{ref}} \} \wedge \\
&\quad \text{dom}(h_b) = \{ l_b \mid \exists \tau, l_a. (\tau, l_a, l_b) \in s_{\text{ref}} \} \wedge \\
&\quad \forall (\tau, l_a, l_b) \in s_{\text{ref}}. (h_a(l_a), h_b(l_b)) \in R(s)(\tau) \} \\
C_{\text{ref}}^{TI}(s_{\text{ref}})(R) &:= \{ ((\emptyset, \emptyset, \emptyset, h_a), (h_b, \emptyset)) \mid h_a \neq \perp \wedge h_b \neq \perp \wedge \\
&\quad \text{dom}(h_a) = \{ l_a \mid \exists l_b. (l_a, l_b) \in s_{\text{ref}} \} \wedge \\
&\quad \text{dom}(h_b) = \{ l_b \mid \exists l_a. (l_a, l_b) \in s_{\text{ref}} \} \wedge \\
&\quad \forall (l_a, l_b) \in s_{\text{ref}}. (h_a(l_a), h_b(l_b)) \in R(s) \} \\
C_{\text{ref}}^{II}(s_{\text{ref}})(R) &:= \{ ((h_a, \emptyset), (h_b, \emptyset)) \mid h_a \neq \perp \wedge h_b \neq \perp \wedge \\
&\quad \text{dom}(h_a) = \{ l_a \mid \exists l_b. (l_a, l_b) \in s_{\text{ref}} \} \wedge \\
&\quad \text{dom}(h_b) = \{ l_b \mid \exists l_a. (l_a, l_b) \in s_{\text{ref}} \} \wedge \\
&\quad \forall (l_a, l_b) \in s_{\text{ref}}. (h_a(l_a), h_b(l_b)) \in R(s) \} \\
C_{\text{ref}}^{IS}(s_{\text{ref}})(R) &:= \{ ((h_a, \emptyset), (h_b, \emptyset, \emptyset)) \mid h_a \neq \perp \wedge h_b \neq \perp \wedge \\
&\quad \text{dom}(h_a) = \{ l_a \mid \exists \tau, l_b. (\tau, l_a, l_b) \in s_{\text{ref}} \} \wedge \\
&\quad \text{dom}(h_b) = \{ l_b \mid \exists \tau, l_a. (\tau, l_a, l_b) \in s_{\text{ref}} \} \wedge \\
&\quad \forall (\tau, l_a, l_b) \in s_{\text{ref}}. (h_a(l_a), h_b(l_b)) \in R(s)(\tau) \}
\end{aligned}$$

Figure 4.16: Global World Components Related to Global References.

$T^{\mathcal{S}}.S$	$:= \text{Env}$
$T^{\mathcal{S}}.\sqsubseteq$	$:= (=)$
$T^{\mathcal{S}}.\sqsubseteq_{\text{pub}}$	$:= (=)$
$\text{vqh}_{\mathcal{S}}(-)(\text{unit})$	$:= \{ \langle \rangle \}$
$\text{vqh}_{\mathcal{S}}(-)(\text{nat } \mathbf{v})$	$:= \{ n \mid n = \mathbf{v} \}$
$\text{vqh}_{\mathcal{S}}(-)(\text{pair } \mathbf{v} \mathbf{v}')$	$:= \{ \langle \mathbf{v}, \mathbf{v}' \rangle \}$
$\text{vqh}_{\mathcal{S}}(-)(\text{inl } \mathbf{v})$	$:= \{ \text{inl } \mathbf{v} \}$
$\text{vqh}_{\mathcal{S}}(-)(\text{inr } \mathbf{v})$	$:= \{ \text{inr } \mathbf{v} \}$
$\text{vqh}_{\mathcal{S}}(-)(\text{pack } \mathbf{v})$	$:= \{ \text{pack } \mathbf{v} \}$
$\text{vqh}_{\mathcal{S}}(-)(\text{roll } \mathbf{v})$	$:= \{ \text{roll } \mathbf{v} \}$
$\text{vqh}_{\mathcal{S}}(-)(\text{fix})$	$:= \{ \text{fix } f(x).e \}$
$\text{vqh}_{\mathcal{S}}(-)(\text{gen})$	$:= \{ \Lambda.e \}$
$\text{vqh}_{\mathcal{S}}(-)(\text{name})$	$:= \mathbf{Val}$
$\text{vqh}_{\mathcal{S}}(-)(\text{goodfix})$	$:= \mathbf{Val}$
$\text{vqh}_{\mathcal{S}}(-)(\text{goodgen})$	$:= \mathbf{Val}$
$\text{cqh}_{\mathcal{S}}(-)(\text{app } \mathbf{v} \mathbf{v}' \mathbf{k})$	$:= \{ (\emptyset, \emptyset, \mathbf{k}[e[\mathbf{v}/f][\mathbf{v}'/x]]) \mid \mathbf{v} = \text{fix } f(x).e \}$
$\text{cqh}_{\mathcal{S}}(-)(\text{inst } \mathbf{v} \mathbf{k})$	$:= \{ (\emptyset, \emptyset, \mathbf{k}[e]) \mid \mathbf{v} = \Lambda.e \}$
$\text{cqh}_{\mathcal{S}}(-)(\text{ret } \mathbf{v} \mathbf{k})$	$:= \{ (\emptyset, \emptyset, \mathbf{k}[\mathbf{v}]) \}$
$P^{\mathcal{S}}(s) \in \mathcal{P}(\mathbf{Conf})$	
$P^{\mathcal{S}}(s) := \{ (\emptyset, s, \emptyset) \}$	

Figure 4.17: Global World Components Concerning Language \mathcal{S} .

$T^{\mathcal{I}}.S$	$:= 1$
$T^{\mathcal{I}}.\sqsubseteq$	$:= 1 \times 1$
$T^{\mathcal{I}}.\sqsubseteq_{\text{pub}}$	$:= 1 \times 1$
$\text{vqh}_{\mathcal{I}}(-)(\text{unit})$	$:= \{ \langle \rangle \}$
$\text{vqh}_{\mathcal{I}}(-)(\text{nat } \mathbf{v})$	$:= \{ n \mid n = \mathbf{v} \}$
$\text{vqh}_{\mathcal{I}}(-)(\text{pair } \mathbf{v} \ \mathbf{v}')$	$:= \{ \langle \mathbf{v}, \mathbf{v}' \rangle \}$
$\text{vqh}_{\mathcal{I}}(-)(\text{inl } \mathbf{v})$	$:= \{ \text{inl } \mathbf{v} \}$
$\text{vqh}_{\mathcal{I}}(-)(\text{inr } \mathbf{v})$	$:= \{ \text{inr } \mathbf{v} \}$
$\text{vqh}_{\mathcal{I}}(-)(\text{pack } \mathbf{v})$	$:= \{ \mathbf{v} \}$
$\text{vqh}_{\mathcal{I}}(-)(\text{roll } \mathbf{v})$	$:= \{ \mathbf{v} \}$
$\text{vqh}_{\mathcal{I}}(-)(\text{fix})$	$:= \{ \langle \sigma; \text{fix } f(y, k). e \rangle \}$
$\text{vqh}_{\mathcal{I}}(-)(\text{gen})$	$:= \{ \langle \sigma; \Lambda. e \rangle \}$
$\text{vqh}_{\mathcal{I}}(-)(\text{name})$	$:= \{ \langle \sigma; n \rangle \}$
$\text{vqh}_{\mathcal{I}}(-)(\text{goodfix})$	$:= \{ \langle \sigma; \text{fix } f(y, k). e \rangle \mid \forall n. e \neq \text{halt } n \}$
$\text{vqh}_{\mathcal{I}}(-)(\text{goodgen})$	$:= \{ \langle \sigma; \Lambda k. e \rangle \mid \forall n. e \neq \text{halt } n \}$
$\text{cqh}_{\mathcal{I}}(-)(\text{app } \mathbf{v} \ \mathbf{v}' \ \mathbf{k})$	$:= \{ (\emptyset, (\sigma[f \mapsto \mathbf{v}, y \mapsto \mathbf{v}', k \mapsto \mathbf{k}], e)) \mid \mathbf{v} = \langle \sigma; \text{fix } f(y, k). e \rangle \}$
$\text{cqh}_{\mathcal{I}}(-)(\text{inst } \mathbf{v} \ \mathbf{k})$	$:= \{ (\emptyset, (\sigma[k \mapsto \mathbf{k}], e)) \mid \mathbf{v} = \langle \sigma; \Lambda k. e \rangle \}$
$\text{cqh}_{\mathcal{I}}(-)(\text{ret } \mathbf{v} \ \mathbf{k})$	$:= \{ (\emptyset, (\sigma, k \ x)) \mid \sigma(k) = \mathbf{k} \wedge \sigma(x) = \mathbf{v} \}$
$P^{\mathcal{I}}(s) \in \mathcal{P}(\mathbf{Conf})$	
$P^{\mathcal{I}}(s) := \{ \emptyset \}$	

Figure 4.18: Global World Components Concerning Language \mathcal{I} .

The associated configuration predicate $\text{cpred}^{\mathcal{S}}$ simply requires that the configuration consists solely of the state's environment.

4.4.2.2 Global World Components Regarding \mathcal{I}

Figure 4.18 shows the constructions for the intermediate language \mathcal{I} . In contrast to \mathcal{S} , the run-time environment in \mathcal{I} is not shared by all modules in a program but is local to a computation. In fact, each function value, being a closure, comes with its own environment. Therefore, there is no need to keep track of anything purely \mathcal{I} -specific in the global world and we can pick the trivial transition system (with a singleton state space).

Like for the source language, the value and configuration query handlers are purely syntactic and mostly straightforward. Note the following points:

- Being untyped, \mathcal{I} does not provide a “pack” or “roll” construct and thus there are no “packed” or “rolled” values. We decide to represent such \mathcal{S} values the

same as the underlying values. This decision is arbitrary but natural.

- We restrict inhabitants of type names to be numbers.
- We require that the body of “good” functions is not a `halt n` expression. We will use such “bad” functions when crafting intermediate values in the proof of transitivity.

4.4.2.3 Global World Components Regarding \mathcal{T}

The constructions for the machine language \mathcal{T} , shown in Figure 4.19, are more interesting. Recall that \mathcal{T} configurations consist of (possibly) a program counter, (possibly) a register file, a stack and a heap. We set things up such that the global world always contributes the register file, the free part of the stack, and the part of the heap used for allocating values. Formally, the state space $\mathsf{T}^{\mathcal{T}}.\mathsf{S}$ ranges over pairs of register file and *value registry*. Let us first focus on the register file.

Registers. The predicate $\mathsf{P}^{\mathcal{T}}$ only allows configurations whose register file is precisely the one in the current state. The full transition relation allows the contents of the register file to change arbitrarily, because any code may write to a register. The public transition relation, however, requires that the callee-saved registers (*i.e.*, `sp` and `env`) retain their value. This corresponds to what the calling convention prescribes for the end-to-end behavior of functions.

Stack. Recall that `sp` is intended to function as the stack pointer. Hence the stack portion above and including address $R(\text{sp})$ is considered free. Accordingly, $\mathsf{P}^{\mathcal{T}}$ only allows configurations whose stack is exactly this infinite, cofinite, and continuous chunk (its contents do not matter). This formally ties the stack to register `sp`. Moreover, it rules out that any particular module puts constraints on that part. Note that this, together with the public transition relation, implies that PILS do not expect functions to preserve the unused part of the stack.

Value registry. The value registry’s purpose is to keep track of allocated values, primarily such that the value query handler can do its job. As the figure shows, it is a mapping from *value descriptions* to sets of values, and `vqh $_{\mathcal{T}}$` consults this mapping to find out which pairs, sums, and functions are currently known. Other value forms don’t require an entry in the registry (for instance, numbers are unboxed and simply represent themselves). We collapse fix-points and generalizations as they are represented the same way in \mathcal{T} .

This representation, and that of the other value forms, manifests itself in the configuration predicate $\mathsf{cpred}^{\mathcal{T}}$, which connects the registry to the memory (and thus to reality). It dictates that in order for a value v to represent a pair of v_1 and v_2 ($v \in s.\rho(\text{pair}_{\mathcal{T}} v_1 v_2)$), v must be a heap pointer to v_1 , with v_2 being stored in the adjacent cell. Similarly, sum values are represented by a heap pointer to a tag (zero for

$$\begin{aligned}
d \in \text{ValDescr} &::= \text{pair}_{\mathcal{T}} \mathbf{v} \mathbf{v}' \mid \text{inl}_{\mathcal{T}} \mathbf{v} \mid \text{inr}_{\mathcal{T}} \mathbf{v} \mid \text{fun}_{\mathcal{T}} \mathbf{v} & (\mathbf{v} \in \mathbf{Val}) \\
\rho \in \text{Registry} &::= \text{ValDescr} \rightarrow \mathcal{P}(\mathbf{Val}) \\
\\
\mathsf{T}^{\mathcal{T}}.S &::= \text{RegFile} \times \text{Registry} \\
\mathsf{T}^{\mathcal{T}}.\sqsubseteq &::= \{ ((R, \kappa), (R', \kappa')) \mid \kappa \subseteq \kappa' \} \\
\mathsf{T}^{\mathcal{T}}.\sqsubseteq_{\text{pub}} &::= \{ ((R, \kappa), (R', \kappa')) \mid \kappa \subseteq \kappa' \wedge R(\text{sp}) = R'(\text{sp}) \wedge R(\text{env}) = R'(\text{env}) \} \\
\\
\text{vqh}_{\mathcal{T}}(-)(\text{unit}) &::= \mathbf{Val} \\
\text{vqh}_{\mathcal{T}}(-)(\text{nat } \mathbf{v}) &::= \{ \mathbf{v} \} \\
\text{vqh}_{\mathcal{T}}(s)(\text{pair } \mathbf{v} \mathbf{v}') &::= s.\rho(\text{pair}_{\mathcal{T}} \mathbf{v} \mathbf{v}') \\
\text{vqh}_{\mathcal{T}}(s)(\text{inl } \mathbf{v}) &::= s.\rho(\text{inl}_{\mathcal{T}} \mathbf{v}) \\
\text{vqh}_{\mathcal{T}}(s)(\text{inr } \mathbf{v}) &::= s.\rho(\text{inr}_{\mathcal{T}} \mathbf{v}) \\
\text{vqh}_{\mathcal{T}}(-)(\text{roll } \mathbf{v}) &::= \{ \mathbf{v} \} \\
\text{vqh}_{\mathcal{T}}(-)(\text{pack } \mathbf{v}) &::= \{ \mathbf{v} \} \\
\text{vqh}_{\mathcal{T}}(s)(\text{fix}) &::= \{ \mathbf{v} \mid \exists \mathbf{v}'. \mathbf{v} \in s.\rho(\text{fun}_{\mathcal{T}} \mathbf{v}') \} \\
\text{vqh}_{\mathcal{T}}(s)(\text{gen}) &::= \{ \mathbf{v} \mid \exists \mathbf{v}'. \mathbf{v} \in s.\rho(\text{fun}_{\mathcal{T}} \mathbf{v}') \} \\
\text{vqh}_{\mathcal{T}}(-)(\text{name}) &::= \mathbf{Val} \\
\text{vqh}_{\mathcal{T}}(-)(\text{goodfix}) &::= \mathbf{Val} \\
\text{vqh}_{\mathcal{T}}(-)(\text{goodgen}) &::= \mathbf{Val} \\
\\
\text{cqh}_{\mathcal{T}}(s)(\text{app } \mathbf{v} \mathbf{v}' \mathbf{k}) &::= \{ (n, \emptyset, \emptyset, \emptyset) \mid \mathbf{v} = s.R(\text{clo}) \wedge \mathbf{v}' = s.R(\text{arg}) \wedge \\
&\quad \mathbf{k} = s.R(\text{ret}) \wedge \mathbf{v} \in s.\rho(\text{fun}_{\mathcal{T}} n) \} \\
\text{cqh}_{\mathcal{T}}(s)(\text{inst } \mathbf{v} \mathbf{k}) &::= \{ (n, \emptyset, \emptyset, \emptyset) \mid \mathbf{v} = s.R(\text{clo}) \wedge \\
&\quad \mathbf{k} = s.R(\text{ret}) \wedge n \in s.\rho(\text{fun } \mathbf{v}) \} \\
\text{cqh}_{\mathcal{T}}(s)(\text{ret } \mathbf{v} \mathbf{k}) &::= \{ (\mathbf{k}, \emptyset, \emptyset, \emptyset) \mid \mathbf{v} = s.R(\text{arg}) \} \\
\\
\mathsf{P}^{\mathcal{T}} \in \mathsf{T}^{\mathcal{T}}.S \rightarrow \mathcal{P}(\mathbf{Conf}) \\
\mathsf{P}^{\mathcal{T}}(R, \kappa) &::= \{ (\emptyset, R, q, h) \mid q \neq \perp \wedge \text{dom}(q) = \{n \mid n \geq R(\text{sp})\} \wedge \forall \mathbf{v}, \mathbf{v}', \mathbf{v}''. \\
&\quad (\mathbf{v} \in \kappa(\text{pair}_{\mathcal{T}} \mathbf{v}' \mathbf{v}'') \implies h(\mathbf{v}) = \mathbf{v}' \wedge h(\mathbf{v} + 1) = \mathbf{v}'') \wedge \\
&\quad (\mathbf{v} \in \kappa(\text{inl}_{\mathcal{T}} \mathbf{v}') \implies h(\mathbf{v}) = 0 \wedge h(\mathbf{v} + 1) = \mathbf{v}') \wedge \\
&\quad (\mathbf{v} \in \kappa(\text{inr}_{\mathcal{T}} \mathbf{v}') \implies \exists n \neq 0. h(\mathbf{v}) = n \wedge h(\mathbf{v} + 1) = \mathbf{v}') \wedge \\
&\quad (\mathbf{v} \in \kappa(\text{fun}_{\mathcal{T}} \mathbf{v}') \implies h(\mathbf{v}) = \mathbf{v}') \}
\end{aligned}$$

Figure 4.19: Global World Components Concerning Language \mathcal{T} .

“left”, non-zero for “right”), with the data being in the adjacent cell. Function values are just heap addresses (intuitively: pointers to closures). The contents stored at such an address is the code pointer, which is part of the value description: $s.\rho(\text{fun}_{\mathcal{T}} \mathbf{v}')$ is the set of functions with code pointer \mathbf{v}' . From the point of view of the global world, only this heap cell matters. Code and environment, possibly stored in the consecutive cells, are subject to a local world.

Both public and full transition relation allow the registry to grow over time, but forbid any other modification. This ensures that any registered value stays valid forever, and never changes intrinsically⁶. Of course, the (required) monotonicity of $\text{vqh}_{\mathcal{T}}$ relies on this.

Configuration queries. The registry also plays a role in $\text{cqh}_{\mathcal{T}}$. How should the `app` and `inst` queries be implemented for \mathcal{T} , *i.e.*, what shall count as function calls? Following Hur and Dreyer [32], we choose to be very liberal in what we accept here. For instance, we do not want to require that the program counter points to a particular “call” instruction (in fact, we don’t even assume such an instruction in our language). To be as flexible as possible, we define an application as the (relevant part of the) machine state *after* the function call happened. That is, we require that the program counter points to the first instruction of the function’s code and that the appropriate registers are filled with return address and argument value (in the case of a fix-point application), as dictated by the calling convention. We also require that the function is registered as such. This is important because otherwise, when reasoning about a function call, one may not know whether reading the code pointer (before jumping) succeeds.

The `ret \mathbf{v} \mathbf{k}` query is easier to answer: the configuration’s program counter must be exactly the continuation \mathbf{k} and register `arg` must contain value \mathbf{v} .

4.5 Simulations

The definitions of the various PILS simulation relations and auxiliary constructions are shown in Figures 4.20 to 4.22. (As before, the parts unique to the typed model are **highlighted**.) Let us first give a brief overview.

Most notions and concepts from PBs still exist in PILS. For some of the constructions, the main difference in generalizing to the inter-language setting is simply that they are now abstracted over the details of the particular languages being related, with the help of language specification and worlds. That said, there are significant changes as well. Recall that (i) we want PILS to be asymmetric (because so is refinement), (ii) we are getting rid of an explicit local knowledge component, (iii) we want to avoid having two sources of coinduction (as were present in PBs), and (iv) we both construct a typed and an untyped version.

⁶Benign changes to the concrete representation are allowed. For instance, the tag of an “inr” value may change from 1 to 2, because both are non-zero.

$$\begin{aligned}
& \text{ES}_{A,B}^W := W.O \rightarrow A.\mathbf{Cont} \times B.\mathbf{Cont} \rightarrow (W.S \rightarrow \mathbf{VRelF}_{A,B}) \rightarrow W.S \rightarrow W.S \rightarrow \\
& \quad (A.\mathbf{Conf} \times B.\mathbf{Conf})_{\perp} \rightarrow \mathbf{ERel}_{A,B} \\
& \text{KS}_{A,B}^W := W.O \rightarrow A.\mathbf{Cont} \times B.\mathbf{Cont} \rightarrow (W.S \rightarrow \mathbf{VRelF}_{A,B}) \rightarrow W.S \rightarrow W.S \rightarrow \mathbf{KRel}_{A,B} \\
& \mathbf{E}_W \in \text{ES}_{A,B}^W \\
& \mathbf{E}_W(i)(\mathbf{k}_a^0, \mathbf{k}_b^0)(G)(s_0)(s)(\phi)(\tau) = \{ (e_a, e_b) \mid G \in \mathbf{GK}_W \implies \\
& \quad \forall c_a, c_b. \forall \eta_a \in \mathbf{frame}. \forall \eta_b \in \mathbf{frame}. \\
& \quad \forall (m_a, m_b) \in \mathbf{cfg}(W.C)(G)(s)(\phi)(e_a, e_b)(c_a, c_b)(\eta_a, \eta_b). \\
& \quad (\mathbf{ERR}) \exists m'_b. m_b \xrightarrow{\iota}^* m'_b \wedge m'_b \in \mathbf{error} \\
& \vee (\mathbf{RET}) \exists m'_b, s', \mathbf{v}_a, \mathbf{v}_b, e'_a, e'_b, c'_a, c'_b. m_b \xrightarrow{\iota}^* m'_b \wedge s' \sqsupseteq s \wedge s' \sqsupseteq_{\text{pub}} s_0 \wedge \\
& \quad (m_a, m'_b) \in \mathbf{cfg}(W.C)(G)(s')(\perp)(e'_a, e'_b)(c'_a, c'_b)(\eta_a, \eta_b) \wedge (\mathbf{v}_a, \mathbf{v}_b) \in \overline{G}(s')(\tau) \wedge \\
& \quad (e'_a, e'_b) \in W.\mathbf{cqh}_a(s')(\text{ret } \mathbf{v}_a \mathbf{k}_a^0) \times W.\mathbf{cqh}_b(s')(\text{ret } \mathbf{v}_b \mathbf{k}_b^0) \\
& \vee (\mathbf{STEP}) m_a \notin \mathbf{halted} \wedge \forall t, m'_a. m_a \xrightarrow{t} m'_a \implies \exists i', e'_a, e'_b, c'_a, c'_b, m'_b, m''_b, s', \phi'. \\
& \quad m_b \xrightarrow{\iota}^* m'_b \wedge s' \sqsupseteq s \wedge (m'_a, m''_b) \in \mathbf{cfg}(W.C)(G)(s')(\phi')(e'_a, e'_b)(c'_a, c'_b)(\eta_a, \eta_b) \wedge \\
& \quad (\mathbf{REC}) (m'_b \xrightarrow{t} m''_b \vee m'_b = m''_b \wedge t = \iota \wedge i' < i) \wedge \phi' \neq \perp \wedge \\
& \quad (e'_a, e'_b) \in \mathbf{E}_W(i')(\mathbf{k}_a^0, \mathbf{k}_b^0)(G)(s_0)(s')(\phi')(\tau) \\
& \vee (\mathbf{CALL}) m'_b \xrightarrow{t} m''_b \wedge \phi' = \perp \wedge \\
& \quad (e'_a, e'_b) \in \mathbf{U}(s')(G(s'))(G(s'))(\mathbf{K}_W(i')(\mathbf{k}_a^0, \mathbf{k}_b^0)(G)(s_0)(s'))(\tau) \} \\
& \mathbf{K}_W \in \text{KS}_{A,B}^W \\
& \mathbf{K}_W(i)(\mathbf{k}_a^0, \mathbf{k}_b^0)(G)(s_0)(s)(\tau')(\tau) = \{ (\mathbf{k}_a, \mathbf{k}_b) \mid \forall G' \supseteq G. \forall s' \sqsupseteq_{\text{pub}} s. \forall (v_a, v_b) \in \overline{G'}(s')(\tau'). \\
& \quad W.\mathbf{cqh}_a(s')(\text{ret } v_a \mathbf{k}_a) \times W.\mathbf{cqh}_b(s')(\text{ret } v_b \mathbf{k}_b) \subseteq \mathbf{E}_W(i)(\mathbf{k}_a^0, \mathbf{k}_b^0)(G')(s_0)(s')(\perp)(\tau) \} \\
& \mathbf{F}_W \in (W.S \rightarrow \mathbf{VRelF}_{A,B}) \rightarrow W.S \rightarrow \mathbf{VRelF}_{A,B} \\
& \mathbf{F}_W(G)(s)(\tau) := \{ (\mathbf{v}_a, \mathbf{v}_b) \in \mathbf{GoodFuns}(s)(\tau) \mid \exists i. \forall \mathbf{k}_a, \mathbf{k}_b. \forall G' \supseteq G. \forall s' \sqsupseteq s. \\
& \quad \mathbf{U}(s')(\mathbf{sgn}(\tau)(\mathbf{v}_a, \mathbf{v}_b))(G'(s'))(\mathbf{sgn}(\mathbf{k}_a, \mathbf{k}_b)) \subseteq \mathbf{E}_W(i)(\mathbf{k}_a, \mathbf{k}_b)(G')(s')(s')(\perp) \} \\
& \mathbf{GK}_W \in \mathcal{P}(W.S \rightarrow \mathbf{VRelF}_{A,B}) \\
& \mathbf{GK}_W := \{ G \mid \mathbf{F}_W(G) \subseteq G \wedge (\forall s, s'. s' \sqsupseteq s \implies G(s') \supseteq G(s)) \wedge \\
& \quad \forall s. \forall \mathbf{n} \in W.\mathbf{NS}. G(s)(\mathbf{n}) = W.\mathbf{NR}(G)(s)(\mathbf{n}) \} \\
& \Gamma \vdash M_a \lesssim_{\Omega} M_b : \Gamma' := \\
& \quad \forall \mathcal{N} \in \mathcal{P}(\mathbf{TyNam}). \mathcal{N} \text{ countably infinite} \implies \exists w \in \mathbf{LWorld}_{A,B}^{\Omega, \mathbf{T}}. \\
& \quad \forall \gamma_a, \gamma_b. \text{dom}(\gamma_a) = \mathbf{dom}(\Gamma) = \text{dom}(\gamma_b) \implies \forall \Psi_a, \Psi_b. \\
& \quad \forall (h_a^g, h_a^l) \in \mathbf{cload}(M_a)(\Psi_a)(\gamma_a). \forall (h_b^g, h_b^l) \in \mathbf{cload}(M_b)(\Psi_b)(\gamma_b). \exists s_0. \\
& \quad w.\mathbf{NS} \subseteq \mathcal{N} \wedge \\
& \quad w \in \mathbf{stable}(\Omega) \wedge \\
& \quad (\forall G \in \mathbf{GK}_{w\uparrow}. (h_a^g, h_b^g) \in \Omega.\mathbf{C}(G(-, \text{snd } s_0))(\text{fst } s_0)) \wedge \\
& \quad (\forall G \in \mathbf{GK}_{w\uparrow}. (h_a^l, h_b^l) \in w.\mathbf{C}(G)(s_0)) \wedge \\
& \quad \Omega.\mathbf{rqh}(\text{fst } s_0) = \emptyset \wedge \\
& \quad \forall f' : \tau' \in \Gamma'. \exists (v_a, v_b) \in \mathbf{vload}(M_a)(\Psi_a)(\gamma_a)(f') \times \mathbf{vload}(M_b)(\Psi_b)(\gamma_b)(f'). \\
& \quad \forall s \sqsupseteq s_0. \forall G \in \mathbf{GK}_{w\uparrow}. (s \in \mathbf{sat}(G) \wedge \forall f : \tau \in \Gamma. (\gamma_a f, \gamma_b f) \in \overline{G}(s)(\tau)) \implies \\
& \quad (v_a, v_b) \in \overline{G}(s)(\tau')
\end{aligned}$$

Figure 4.20: Key components of PILS

At the top-level, PILS define *module similarity*, written $\Gamma \vdash M_a \lesssim_{\Omega} M_b : \Gamma'$. This basically generalizes PBs open program equivalence (here we consider modules because they are the objects that we compile and link). This relation is relative to a global world $\Omega \in \text{GWorld}_{A,B}$ for the language pair A, B , with A intuitively being the language *to* which is being translated and B the one *from* which is being translated. Γ and Γ' list the imported and exported labels, respectively (each label is paired with its type in the typed version). At the end, we instantiate this relation with four global worlds in order to obtain the four models of interest: we write $\lesssim_{\mathcal{TS}}$ for the typed model $\lesssim_{\Omega_{\mathcal{TS}}}$, we write $\lesssim_{\mathcal{TI}}$ for the untyped model $\lesssim_{\Omega_{\mathcal{TI}}}$, and so on.

The \mathbf{E} relation. The main ingredient of the top-level relation is \mathbf{E} , the generalization of PBs’s expression equivalence, which retains the familiar distinction of three cases.

In our generic setting there is no notion of expressions—only configurations—but we find it helpful to still refer to the configurations related by \mathbf{E} as such. Indeed, as we will see, for the source language \mathcal{S} these configurations will usually just be expressions (empty heap and no environment). For \mathcal{T} , on the other hand, they will usually just be program counters (empty heap, empty stack, no register file). These configurations will be completed with the help of configurations provided by the world, as discussed in a moment.

Let us come to the formal definition of \mathbf{E} , shown in Figure 4.20. At a high level, it relates two “expressions” e_a (the “target” of a transformation) and e_b (the “source” of a transformation) iff one of three cases holds:

- (ERR) e_b can silently (*i.e.*, without I/O) produce an error.
- (RET) e_a is finished, and e_b can silently finish returning a related value.
- (STEP) e_a can take a step and e_b can match it (perhaps after some internal computation). “Match” means that both steps produce the same event and that the remaining computations either (REC) are again related by \mathbf{E} , or (CALL) are about to call related “external” functions, *i.e.*, functions related by the global knowledge G . This “external call” case is the characteristic feature of PBs, as discussed at length in Chapter 3. In the concrete PILS instances we care about, the event t is in the CALL case guaranteed to always be ι .

We now go into details.

Asymmetric small-step formulation. In contrast to the symmetric big-step formulation of \mathbf{E} in PBs, PILS employ an asymmetric small-step formulation (ASF). Notice how \mathbf{E} ’s STEP case asks us to consider each possible step of the “target” program e_a in turn (when using REC repeatedly), each time asking us to match it with steps of the “source” program e_b . Besides being seemingly necessary to properly deal with events (here: I/O), such an ASF is also important in the context of compiler

verification because it gives the compiler the flexibility to remove erroneous behaviors of the source program and resolve some of its nondeterminism.

In a naive ASF, each step of the target must be matched with one or more steps of the source, thus forcing the source to take at least as many steps as the target. Of course, this is overly restrictive and our actual definition allows *stuttering*, *i.e.*, it allows the source to not take any steps at all. In order for this to be sound, however, some other form of progress must occur; otherwise one could trivially show any divergent target program related to any source program. We follow the standard approach to stuttering by indexing the **E** relation with an element i of a well-founded partially ordered set ($W.O$). We postpone a discussion of the details to Section 4.5.1.

Notice that the check that $m_a \notin \mathbf{halted}$, *i.e.*, that m_a can take a step, is necessary to keep the **STEP** case from being trivially satisfied whenever the program cannot continue.

From configurations to machines and back. In order to talk about the execution of e_a and e_b , we first need to “complete” these configurations and convert them into physical machines. These completions should not be completely arbitrary; they should adhere to the world’s constraints at the current state s . Hence we first quantify over c_a and c_b , representing the portion of the machine state constrained by the world W (an implicit argument), and require that they are indeed related by $W.C(G)(s)$. This is done with the help of the auxiliary construct \mathbf{cfg} from Figure 4.21. Its ϕ argument, which is also an argument to **E**, plays an important role, but we postpone its explanation. For now let us assume that it is always \perp (even in **E**’s **REC** case), so that we don’t need to worry about the **MORE** disjunct in \mathbf{cfg} .

Then, given such c_a and c_b , we attach these to e_a and e_b , together with arbitrary frame configurations η_a and η_b representing the rest of the running program state. Finally, we only consider machines m_a and m_b that realize these composed configurations⁷.

The two other occurrences of \mathbf{cfg} , namely in **STEP** and **RET**, are proof obligations. For instance, **STEP** requires us to show that, after the step(s), each resulting machine can again be decomposed into a (possibly new) expression e'_a , a (possibly new) configuration c'_a , and the *original* frame configuration η_a (similarly for the B -side), since we should not have touched the frame’s private memory. Moreover, c'_a and c'_b must again satisfy W ’s constraints, but we may advance s to a future state s' in order to achieve that.

Configuration queries. In the **RET** case of **E**, we generalize PB’s “terminating with a value” to “returning a value to the initial continuation”, where the initial continuations k_a^0, k_b^0 are given as additional arguments to **E**. As expected, we do so by appeal to the world’s configuration query handlers $W.cqh_a$ and $W.cqh_b$. Similarly,

⁷Note that, in our three concrete languages, a machine realizing a configuration is identical to the configuration (modulo the obvious embedding).

$$\begin{aligned}
& \text{cfg} \in \text{CR}_{A,B}^T \rightarrow (T.S \rightarrow \text{VRelF}_{A,B}) \rightarrow T.S \rightarrow (A.\mathbf{Cont} \times B.\mathbf{Cont})_{\perp} \rightarrow \\
& \quad A.\mathbf{Conf} \times B.\mathbf{Conf} \rightarrow A.\mathbf{Conf} \times B.\mathbf{Conf} \rightarrow A.\mathbf{Conf} \times B.\mathbf{Conf} \rightarrow \text{MRel}_{A,B} \\
& \text{cfg}(C)(G)(s)(\phi)(e_a, e_b)(c_a, c_b)(\eta_a, \eta_b) = \{ (m_a, m_b) \mid \\
& \quad (\text{CORE}) \ \phi = \perp \wedge m_a \in \mathbf{real}(e_a \cdot c_a \cdot \eta_a) \wedge m_b \in \mathbf{real}(e_b \cdot c_b \cdot \eta_b) \wedge \\
& \quad \quad e_a \in \mathbf{core} \wedge e_b \in \mathbf{core} \wedge (c_a, c_b) \in C(G)(s) \\
& \quad \vee (\text{MORE}) \ \phi = (\eta_a, \eta_b) \wedge m_a \in \mathbf{real}(e_a) \wedge m_b \in \mathbf{real}(e_b) \wedge \\
& \quad \quad c_a = \emptyset \wedge c_b = \emptyset \} \\
\\
& \mathbf{U} \in W.S \rightarrow \text{VRelF}_{A,B} \rightarrow \text{VRelF}_{A,B} \rightarrow \text{KRel}_{A,B} \rightarrow \mathbf{CTy} \rightarrow \text{CRel}_{A,B} \\
& \mathbf{U}(s)(R_f)(R_v)(R_k)(\tau) := \{ (e_a, e_b) \mid \\
& \quad (\text{APP}) \ \exists \tau_v, \tau_r. \exists (\mathbf{k}_a, \mathbf{k}_b) \in R_k(\tau_r)(\tau). \\
& \quad \quad \exists (f_a, f_b) \in \langle\langle R_f \rangle\rangle^s(\tau_v \rightarrow \tau_r). \exists (v_a, v_b) \in \langle\langle R_v \rangle\rangle^s(\tau_v). \\
& \quad \quad (e_a, e_b) \in W.\text{cqh}_a(s)(\text{app } f_a v_a \mathbf{k}_a) \times W.\text{cqh}_b(s)(\text{app } f_b v_b \mathbf{k}_b) \\
& \quad \vee (\text{INST}) \ \exists \alpha, \tau_r, \tau_i. \exists (\mathbf{k}_a, \mathbf{k}_b) \in R_k(\tau_r[\tau_i/\alpha])(\tau). \\
& \quad \quad \exists (f_a, f_b) \in \langle\langle R_f \rangle\rangle^s(\forall \alpha. \tau_r). \\
& \quad \quad (e_a, e_b) \in W.\text{cqh}_a(s)(\text{inst } f_a \mathbf{k}_a) \times W.\text{cqh}_b(s)(\text{inst } f_b \mathbf{k}_b) \} \\
\\
& \text{sat} \in (w \uparrow .S \rightarrow \text{VRelF}_{A,B}) \rightarrow \mathcal{P}(w \uparrow .S) \\
& \text{sat}(G) := \{ s \mid \exists (c_a, c_b) \in w.\mathbf{C}(G)(s) \} \\
\\
& \text{stable}(\Omega) \in \mathcal{P}(\text{LWorld}_{A,B}^{\Omega}) \\
& \text{stable}(\Omega) := \{ w \mid \forall G, s_g, s, s'_g, c_a, c_b, e_a, e_b, \eta_a, \eta_b, c'_a, c'_b, m_a, m_b. \\
& \quad G \in \mathbf{GK}_{w \uparrow} \wedge (c_a, c_b) \in w.\mathbf{C}(G)(s_g, s) \wedge s'_g \sqsupseteq s_g \wedge \\
& \quad (m_a, m_b) \in \text{cfg}(\Omega.\mathbf{C})(G(-, s))(s_g)(e_a, e_b)(c'_a, c'_b)(c_a \cdot \eta_a, c_b \cdot \eta_b) \implies \\
& \quad \exists s' \sqsupseteq_{\text{pub}} s. (c_a, c_b) \in w.\mathbf{C}(G)(s'_g, s') \} \\
\\
& \text{GoodFuns} \in W.S \rightarrow \text{VRelF}_{A,B} \\
& \text{GoodFuns}(s)(\tau) = \{ (\mathbf{v}_a, \mathbf{v}_b) \mid \\
& \quad (\text{FIX}) \ \exists \tau_v, \tau_r. \tau = \tau_v \rightarrow \tau_r \wedge (\mathbf{v}_a, \mathbf{v}_b) \in W.\text{vqh}_a(s)(\text{goodfix}) \times W.\text{vqh}_b(s)(\text{goodfix}) \\
& \quad \vee (\text{GEN}) \ \exists \alpha, \tau_r. \tau = \forall \alpha. \tau_r \wedge (\mathbf{v}_a, \mathbf{v}_b) \in W.\text{vqh}_a(s)(\text{goodgen}) \times W.\text{vqh}_b(s)(\text{goodgen}) \}
\end{aligned}$$

Figure 4.21: Key components of PILS (continued)

in the **CALL** case, we use queries to check if the configurations in question represent function calls (see the definition of **U** in Figure 4.21).

Notice that **E**’s argument type τ (in the typed version) is the type of the values being returned to the initial continuations.

The K relation. Like for PBs, the continuation relation **K**—referenced in **E**’s **RET** case—is itself defined straightforwardly in terms of **E**, but is now phrased with the help two return queries to express the use of the two continuations language-independently.

The F relation. Recall from Section 4.3 that, in PILS, we want to always work with what we knew in PBs as the greatest consistent local knowledge, but we do not want to have it as an explicit component of the world, because that would be pointless. We achieve this as follows. We define once and for all the *function similarity* **F**, which can be thought of as the greatest consistent local knowledge but is not part of the world. This relation is easy to define on top of **E**, as shown in Figure 4.20. It roughly says that future uses of the functions with arguments related by the global knowledge and with arbitrary initial continuations must be related by **E**. The second argument to **U**, $\text{sng}(\tau)(\mathbf{v}_a, \mathbf{v}_b)$, is a shortcut for the singleton value relation $\lambda\tau'. \{(\mathbf{v}_a, \mathbf{v}_b) \mid \tau' = \tau\}$. The last argument, $\text{sng}(\mathbf{k}_a, \mathbf{k}_b)$, stands for the continuation relation $\lambda\tau, \tau'. \{(\mathbf{k}_a, \mathbf{k}_b) \mid \tau = \tau'\}$, whose untyped version is a singleton and whose typed version relates the two continuations whenever input type and output type coincide. Also note how we “remember” the initial continuations by passing them on to **E**.

How does **F** tie in with the other relations? Note that **E** is defined with the help of **GK**, the set of valid global knowledges. In PBs, this set in turn was defined to contain a global knowledge G iff G included the world’s local knowledge ($G \supseteq W.L(G)$ in the simple version from Section 3.5). In PILS, we are forced to define it in terms of **F** instead, as is shown in the figure. Hence, instead of having a local knowledge component in worlds, in PILS we have an explicit function relation **F**. Moreover, **F**, **GK**, and **E** are all mutually recursive. This leads to some significant differences in the metatheory and the coinductive reasoning in PILS, as we will discover in Section 4.7.

Guardedness. Guardedness ensures the soundness of coinductive reasoning. PBs enforced guardedness via their consistency condition, by requiring that we beta-reduce the function applications before showing them related by **E** (Section 3.5). In PILS, this might naturally be done in an analogous way via the definition of **F**. However, PILS actually work a little differently, for the following reason: for our machine language, we implement the configuration query such that function calls are actually configurations whose program counter already points to the first instruction of the function’s code. For uniformity and to avoid mismatches in proofs, we do the same for the other languages as well. Now, requiring such function calls to take a

step before showing them related by **E** would be somewhat restrictive as it would rule out functions that don't actually do anything (*e.g.*, $\text{fix } f(x).42$ in \mathcal{S}).

Therefore, things are slightly different in PILS. We do not require any step-taking in **F**, as evident from the formal definition. So, when showing functions like the above related by **F**, one can jump right to the **RETURN** case in **E**. However, when attempting to appeal to **E**'s **CALL** case, we are forced to take a step on both sides. This formulation actually fits very naturally with the asymmetric small-step formulation of **E**. Contrast this with PB's symmetric big-step formulation (Figure 3.15 in Section 3.6.5), where in the **CALL** case both reduction sequences can be empty.

Module similarity. At the bottom of Figure 4.20 we finally define module similarity. Simply put, this says that there exists a local world such that if the imported values are related, then so are the exported values. The details are more complicated because the exports may make sense only in the proper context.

More formally, we are first given a countably infinite set of type names \mathcal{N} , just like in PBs. We then require that there exists a stable local world w (potentially dependent on the global world Ω), whose type names are covered by \mathcal{N} . Next, given global and local configurations obtained from loading the modules, there must exist an initial state s_0 of the full world such that these configurations are related at s_0 by $\Omega.C$ and $w.C$, respectively. (We also require that the initial state does not relate any global references; this is a technical condition that simplifies the proof of modularity.) Finally, for any exported label, the two modules must provide values that are related at any satisfiable future state and global knowledge at which the imported values (in γ_a and γ_b) are also related. Satisfiable here means that there are some configurations related by w at this state; it is a technical condition later used in the proof of transitivity.

4.5.1 Stuttering according to algebraic well-founded orders.

The **E** relation allows stuttering: in the **STEP** case, when we are given a silent step of the target program e_a , we are *not* forced to take a step in the source program e_b . Instead, following a standard technique [75], we can decrease i , which is an element of a well-founded set⁸ and was given to **E** as an argument. The well-foundedness guarantees that we can procrastinate only finitely many times, thus avoiding bogus “proofs” where we stutter ad infinitum and could thereby trivially relate any divergent target program to any source program.

Let us look at the details. Whenever we need to show two functions related by **F**, we get to choose an i for that particular (sub-)proof (cf. Figure 4.20). This i must be an element of the well-founded set $W.O$. In **E**'s **STEP/REC** case, we can avoid taking a step in e_b by providing a smaller i' , meaning $i' < i$ according to the $W.O$'s well-founded order. Then we can continue our proof with respect to i' (notice how it is passed to the recursive occurrence of **E**). It thus makes sense to think of this

⁸A *well-founded set* X consists of a carrier set $X.C$ and a well-founded order $(X.<) \subset X \times X$.

$$\begin{aligned}
\text{AWFO} := \{ & (\mathbf{C} \in \text{Set}, (<) \in \mathcal{P}(\mathbf{C} \times \mathbf{C}), 0 \in \mathbf{C}, 1 \in \mathbf{C}, (+) \in \mathbf{C} \rightarrow \mathbf{C} \rightarrow \mathbf{C}) \mid \\
& (<) \text{ transitive and well-founded} \wedge \\
& (\forall i \in \mathbf{C}. i + 0 = i = 0 + i) \wedge \\
& (\forall i, j, k \in \mathbf{C}. i + (j + k) = (i + j) + k) \wedge \\
& (\forall i, i', j \in \mathbf{C}. i < i' \implies i + j < i' + j) \wedge \\
& (\forall i, j, j' \in \mathbf{C}. j < j' \implies i + j < i + j') \wedge \\
& (\forall i. 0 \leq i) \wedge \\
& 0 \neq 1 \}
\end{aligned}$$

$$\begin{aligned}
\text{Symmetric product: } (& \times) \in \text{AWFO} \rightarrow \text{AWFO} \rightarrow \text{AWFO} \\
(X \times Y).C &:= X.C \times Y.C \\
(X \times Y).0 &:= (0, 0) \\
(X \times Y).1 &:= (1, 1) \\
(i, j) < (i', j') &:= (i < i' \wedge j \leq j') \vee (i \leq i' \wedge j < j') \\
(i, j) + (i', j') &:= (i + i', j + j')
\end{aligned}$$

$$\begin{aligned}
\text{Lexicographic product: } (& \#) \in \text{AWFO} \rightarrow \text{AWFO} \rightarrow \text{AWFO} \\
(X \# Y).C &:= X.C \times Y.C \\
(X \# Y).0 &:= \langle 0, 0 \rangle \\
(X \# Y).1 &:= \langle 1, 1 \rangle \\
\langle i, j \rangle < \langle i', j' \rangle &:= i < i' \vee (i = i' \wedge j < j') \\
\langle i, j \rangle + \langle i', j' \rangle &:= \langle i + i', j + j' \rangle
\end{aligned}$$

Figure 4.22: Algebraically well-founded sets.

argument as a credit or budget. Actually, we get to pick a new i' even if we don't stutter, but then i' does not need to be smaller than i (in fact, there is no restriction on i' at all). Hence as soon as we take a proper step on the source side after stuttering for a while, we are once again eligible for stuttering a finite number of times.

Since i 's universe is a component of the world, we can choose it as we like when we construct the world. Proofs of, say, different program transformations can therefore employ different metrics.

As Figure 4.22 shows, we actually require a little extra monoidal structure of $W.O$, in addition to it being well-founded. It must also provide an order-respecting “addition” with unit 0 and its elements must be non-negative, *e.g.*, such that $i \leq i + j$ and $j \leq i + j$ for any i, j . This is important for composing programs and their proofs—in particular the PILS version of the external call lemma (see Section 4.7.1). We also require that the carrier be non-trivial, *i.e.*, contain at least another element (here written 1) in addition to 0 . This allows us to embed the natural numbers, which is sometimes convenient. We call such monoids *algebraically well-founded sets*.

The figure also shows two product constructions (of two algebraically well-founded sets). The symmetric product is used in the definition of the product of two local

worlds that we have seen earlier in Figure 4.14. The lexicographic product will be used in the proof of transitivity in Section 4.8.

In order to make life easier for the user, we provide an embedding of any ordinary well-founded order $X \in \mathbf{WFO}$ into a particular algebraically well-founded order, namely into $\mathbf{gwf} \in \mathbf{AWFO}$:

Definition 23.

$$\begin{aligned}
 \mathbf{gwf} &\in \mathbf{AWFO} \\
 \mathbf{gwf}.C &:= \text{List}(\{(X, x) \in \mathbf{WFO} \times X.C\}) \\
 \mathbf{gwf}.0 &:= \epsilon \\
 \mathbf{gwf}.1 &:= (\mathbb{N}_{<}, 42) \\
 i + j &:= i, j \\
 \frac{i < j}{i < (X, x), j} &\quad \frac{i < j}{(X, x), i < (X, x), j} \quad \frac{x < x'}{(X, x), i < (X, x'), i}
 \end{aligned}$$

The elements of \mathbf{gwf} are heterogeneous lists of elements from arbitrary well-founded orders (formally, each list component is a pair of a well-founded order X and an element $x \in X.C$ thereof). Addition is defined as concatenation and the empty list is the neutral element. The 1 element is arbitrarily chosen to be a natural number. The order $\mathbf{gwf}.<$ is defined as the least fixed point of the stated rules.

An arbitrary well-founded set X can be embedded in \mathbf{gwf} in the sense of the lemma below. Therefore, instead of worrying about how to make her well-founded set of choice have the required extra structure, a user may simply pick \mathbf{gwf} as her local world's $w.O$.

Lemma 42. Given a well-founded set $X \in \mathbf{WFO}$ and elements $x, x' \in X.C$, we have:

$$x < x' \iff (X, x) < (X, x')$$

In the Coq formalization, we do not require $<$ to be transitive. Instead, we simply use its transitive closure in \mathbf{E} , as well as in the conclusions (but not premises) of the three order-specific axioms. This is equivalent but slightly more convenient for the user who has to define the order and prove these conditions.

4.5.2 The two modes of \mathbf{E} and \mathbf{cfg} .

Recall the ϕ argument of \mathbf{E} and \mathbf{cfg} , which we haven't explained yet. In the discussion of \mathbf{cfg} above, we assumed for simplicity that ϕ 's value is always \perp . Indeed, it is \perp initially, *i.e.*, when showing two function calls related (cf. the definition of \mathbf{F}). However, it may actually change from \perp to a pair (η_a, η_b) of frame configurations—and back and forth—depending on how the proof proceeds. Whenever it changes, the meanings of \mathbf{E} and/or \mathbf{cfg} change with it, so we speak of two different modes in which these construct can function, controlled by the mode argument ϕ .

The purpose of the ϕ argument is to indicate *internal* steps of computation. If it were always \perp , then essentially \mathbf{E} would treat each step as if it might result in

control being passed to the environment. Concretely, after *each* step of the target program and matching steps of the source program, we would be obliged to show that the memory constraints currently imposed by the world are again met. And then, in reasoning about the next step of the target program, we would be forced to quantify over completely new configurations yet again.

Although sound, this is of course unnecessarily strict. Intuitively, we don't need to show that the world's conditions are satisfied again until the point where we pass control to the environment; similarly, we don't need to quantify over new configurations except at points where control is passed to *us*, because there is no way that the memory could have changed in between the internal steps of our local computation.

PILS facilitate this with the help of ϕ as follows. Initially, we have cores e_a, e_b and ϕ is \perp (e.g., as set up by **F**). We call this the *default mode*. Inside **E**, ϕ is used to compose machines via **cfg**. That is, we complete the cores with world-related configurations c_a, c_b and arbitrary frame configurations η_a, η_b to machines m_a, m_b , as dictated by the **CORE** clause in **cfg**. Now, suppose we use **E**'s **REC** case to take a target step to m'_a and possibly some source steps to m''_b . **STEP** asks us to find ϕ' and s' and e'_a, e'_b and c'_a, c'_b such that

$$(m'_a, m''_b) \in \text{cfg}(W.C)(G)(s')(\phi')(e'_a, e'_b)(c'_a, c'_b)(\eta_a, \eta_b)$$

and **REC** requires that e'_a, e'_b are related by **E** for ϕ' . Moreover, **REC** insists that $\phi' \neq \perp$, i.e., it initiates a mode change. Note that by definition of **cfg** (the **MORE** clause), ϕ' must thus necessarily be “chosen” to be (η_a, η_b) . What implications does this have?

First, in order to show that m'_a, m''_b are related as stated above, we are *not* required to show that the world is currently satisfied. Second, **cfg** now enforces that we pick e'_a, e'_b such that they actually account for the whole machines m'_a, m''_b , i.e., the configurations that we need to show related recursively for ϕ' are no longer cores but essentially are the whole machines, including the frame configurations. We thus call this the *fixed mode*.

Third, when we continue the reasoning by setting out to show e'_a, e'_b related in fixed mode (for ϕ'), the new machines that we are given by **cfg**'s **MORE** do not comprise any new parts—they simply realize e'_a and e'_b , respectively.

Assume that we repeat this process for a few times—essentially just taking a few of steps with the target machine and a few steps with the source machine—and then wish to appeal to one of the other **E** cases. In **RET**, we are forced to switch back to the default mode (**cfg** is used with the \perp argument) and thus finally have to get back in synch with the world. That is, we must find a proper future state and decompositions of the machines that satisfy the world at that state and also contain the original pristine frame configurations (this makes sense because control moves on to code that may expect the world's constraints to be satisfied and its local memory to be uncompromised). We know what the original frame configurations are because we stored them in the mode argument (when entering fixed mode earlier) and repeatedly passed on this argument since then.

The same happens in the `CALL` case. Here, of course, we must continue reasoning about the continuations in which the calls are made. Like **F**, **K** is defined in terms of **E**’s default mode. This crucially forces the proof to consider fresh completing configurations in a future state (in which the functions potentially return).

4.5.3 A Note on the Untyped Model

Since our source programs are typesafe and therefore “don’t go wrong”, neither will correctly produced \mathcal{I} or \mathcal{T} programs. One may thus wonder why our model takes faulty programs into account (the `ERR` case in **E**). This feature is actually crucial for verifying transformations in the untyped version of the model. (Recall that we obtain this version by erasing all the type arguments, which are highlighted, from the definitions in the relevant figures.)

To see this, first consider the following optimization at the *source* level (where x is a variable of type `nat`):

$$\text{fix } f(x). \text{ ifnz } x \text{ then } e \text{ else } e \quad \rightsquigarrow \quad \text{fix } f(x). e$$

In the process of showing that $\text{fix } f(x). e$ is similar to the original function, we will be given arguments related at some unknown relation G and state s , $(\text{nat}, v_a, v_b) \in \overline{G}(s)$. Now, by inverting the definition of the value closure, we learn that $v_a = v_b = n$, for some number n .

Let us ignore the remaining proof steps and instead consider this transformation at the IL level, where we would be working in the untyped version of the model. There, we will still be given related arguments, $(v_a, v_b) \in \overline{G}(s)$, but this time the type information is missing. Consequently, when inverting the value closure, we don’t end up with the single case above (where $v_a = v_b = n$), but must also consider all the other cases such as v_a and v_b being pairs. Now, note two important points: First, the global world $\Omega_{\mathcal{II}}$ ensures that whenever v_b is a number, then so is v_a . In that case we can proceed as we would above in the typed model. Second, if v_b is *not* a number, then the original program produces an error and, thanks to `ERR`, there is nothing more to show.

4.5.4 Convenience Lemmas

Reasoning directly using the definition of **E** can be tedious. Here we present a few lemmas that aim at making proofs more convenient and cleaner by providing higher-level reasoning principles. These lemmas (and variants thereof) are used extensively throughout our Coq development. In particular when reasoning about \mathcal{T} code, they have become indispensable. To better grasp the meaning of these lemmas, it often helps to read them bottom-up, a la “in order to show ..., it suffices to show ...”.

The first two express the relationship between **E**’s default mode and its fixed mode. Most remaining lemmas are stated in terms of the fixed mode, so we typically start out our reasoning by applying Lemma 43, which acquires world-related configurations and frames.

Lemma 43 (Acquire). If

$$\begin{aligned} \forall (c_a, c_b) \in W.C(G)(s). \forall (\eta_a, \eta_b) \in \mathbf{frame} \times \mathbf{frame}. \\ (e_a \cdot c_a \cdot \eta_a, e_b \cdot c_b \cdot \eta_b) \in \mathbf{E}_W(i)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)(\eta_a, \eta_b)(\tau) \end{aligned}$$

then $(e_a, e_b) \in \mathbf{E}_W(i)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)(\perp)(\tau)$.

Lemma 44 (Release). If

- $(e_a, e_b) \in \mathbf{E}_W(i)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)(\perp)(\tau)$
- $(c_a, c_b) \in W.C(G)(s)$
- $e_a \in \mathbf{core} \wedge e_b \in \mathbf{core}$

then $(e_a \cdot c_a \cdot \eta_a, e_b \cdot c_b \cdot \eta_b) \in \mathbf{E}_W(i)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)(\eta_a, \eta_b)(\tau)$.

Lemma 45 states that in fixed mode it is always safe to assume the configurations to be realizable.

Lemma 45. If

$$\begin{aligned} (\exists (m_a, m_b) \in \mathbf{real}(c_a) \times \mathbf{real}(c_b)) \implies \\ (c_a, c_b) \in \mathbf{E}_W(i)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)(\eta_a, \eta_b)(\tau) \end{aligned}$$

then $(c_a, c_b) \in \mathbf{E}_W(i)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)(\eta_a, \eta_b)(\tau)$.

Similarly, Lemma 46 allows us to assume that the global knowledge is valid (this holds even in default mode).

Lemma 46. If

$$G \in \mathbf{GK}_W \implies (c_a, c_b) \in \mathbf{E}_W(i)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)(\eta_a, \eta_b)(\tau)$$

then $(c_a, c_b) \in \mathbf{E}_W(i)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)(\eta_a, \eta_b)(\tau)$.

Lemma 47 lets us take several (silent) steps in the source program. If this results in an error, we are done immediately. Otherwise it suffices to show the original target program related to the new source program for budget j . Since we take at least one step, we may choose j arbitrarily (there is no point in allowing zero steps in this lemma).

Lemma 47 (Stepping Source). If

$$\begin{aligned} \forall m_b \in \mathbf{real}(c_b). \exists m'_b. m_b \xrightarrow{\tau}^+ m'_b \wedge \\ (m'_b \in \mathbf{error} \vee \\ \exists c'_b. m'_b \in \mathbf{real}(c'_b) \wedge (c_a, c'_b) \in \mathbf{E}_W(j)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)(\eta_a, \eta_b)(\tau)) \end{aligned}$$

then $(c_a, c_b) \in \mathbf{E}_W(i)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)(\eta_a, \eta_b)(\tau)$.

How about taking steps in the target program? Recall that **E** asks us to consider a single target step at a time and that such a formulation is generally necessary because of possible non-determinism in the target language. However, reasoning in such a way can be extremely tedious since low-level programs are typically very long. Moreover, usually a single source step translates into many target steps, so for most of the target steps one would simply stutter on the source side. Often one actually knows exactly how the given target program will execute. In these cases, instead of being given one target step after the other, each time arguing that it must be a particular step, one would like to just provide a step sequence of the target program and then continue reasoning with its result.

Fortunately, we can prove a lemma that allows exactly this. Of course, it applies only when such reasoning is sound, *i.e.*, when the target execution in question is guaranteed to be deterministic. The lemma thus rests on the idea of lowering determinism from being a property of a language to being a property of a machine and execution.

Definition 24 (Locally deterministic machines). A machine m is *locally deterministic*, written $m \in \mathbf{LDet}$, iff

$$\begin{aligned} & \forall c, t, m', c'. \\ & m \in \mathbf{real}(c) \wedge m \xrightarrow{t} m' \wedge m' \in \mathbf{real}(c') \implies \forall m_1 \in \mathbf{real}(c). \\ & m_1 \notin \mathbf{halted} \wedge \forall t_1, m'_1. m_1 \xrightarrow{t_1} m'_1 \implies m'_1 \in \mathbf{real}(c') \wedge t = t_1 \end{aligned}$$

In words: if m realizes a configuration c and takes a step to a machine realizing c' then any other machine realizing c can also take a step. Moreover, any such step necessarily produces the same event and also results in a machine realizing c' .

Definition 25 (Deterministic step). A *deterministic step* is a step of a locally deterministic machine:

$$m \xrightarrow{t} m' \quad := \quad m \in \mathbf{LDet} \wedge m \xrightarrow{t} m'$$

Obviously, in our \mathcal{T} language, any step that doesn't read input (*i.e.*, $t \neq ?n$) is locally deterministic. In \mathcal{S} and \mathcal{I} , allocations are also excluded.

Lemma 48 makes use of this notion of determinism. It allows us to take multiple (n) silent steps on the target side *if* these steps are deterministic. The “difference” between the original budget i and the new budget j (for which the new target program must be shown related to the original source program) has to be at least n . This is because we need to stutter n times in the proof. We also require n to be non-zero as the property would otherwise not hold generically—but allowing $n = 0$ would be pointless anyways.

Lemma 48 (Stepping Target). If

$$\begin{aligned} & m_a \in \mathbf{real}(c_a) \wedge m_a \xrightarrow{\iota}^n m'_a \wedge m'_a \in \mathbf{real}(c'_a) \wedge \\ & n > 0 \wedge j <^n i \wedge \\ & (c'_a, c_b) \in \mathbf{E}_W(j)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)(\eta_a, \eta_b)(\tau) \end{aligned}$$

then $(c_a, c_b) \in \mathbf{E}_W(i)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)(\eta_a, \eta_b)(\tau)$.

Proof. Since $n > 0$, we have $m_a \in \text{LDet}$ and $m_a \hookrightarrow m_a''$ with $m_a'' \xrightarrow{\iota}^{n-1} m_a'$ for some m_a'' .

We prove the goal by induction on n . After unfolding \mathbf{E} , we are given (via `cfg`) machines $m \in \mathbf{real}(c_a)$ and $m_b \in \mathbf{real}(c_b)$. Since m realizes the same configuration as m_a , we can exploit m_a 's determinacy: the fact that m_a is not halted implies that m is not halted either. This allows us to appeal to \mathbf{E} 's `STEP` case. So suppose $m \xrightarrow{t} m''$.

Case $n = 1$: Further exploiting the determinacy of m_a tells us that t must be ι and that $m'' \in \mathbf{real}(c_a')$. By using the `REC` case and stuttering, it suffices to show that there is $i' < i$ such that

$$(c_a', c_b) \in \mathbf{E}_W(i')(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)(\eta_a, \eta_b)(\tau),$$

which is a premise of the lemma (for $i' := j$).

Case $n > 1$: Since m_a' is not an error (it realizes a configuration) and errors cannot be undone (as required by the language specification), m_a'' cannot be an error either. Hence there is c_a'' such that $m_a'' \in \mathbf{real}(c_a'')$. Further exploiting the determinacy of m_a thus tells us that t must be ι and that $m'' \in \mathbf{real}(c_a'')$ as well. By using the `REC` case and stuttering, it suffices to show that there is $i' < i$ such that

$$(c_a'', c_b) \in \mathbf{E}_W(i')(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)(\eta_a, \eta_b)(\tau).$$

From $j <^n i$ we have i' with $j <^{n-1} i' < i$, so we are done by using the inductive hypothesis on $m_a'' \xrightarrow{\iota}^{n-1} m_a'$.

□

Of course, if one wants to take steps on both sides, one can simply compose Lemma 48 with Lemma 47. In that case, Lemma 47 should be applied first, because it allows one to pick the intermediate budget such that Lemma 48's condition on it trivially holds.

Lemma 48 has a minor shortcoming, though. Suppose that after taking some steps, we want to finish the proof of relatedness by appeal to \mathbf{E} 's `CALL` case. Because the premise is stated in terms of \mathbf{E} , however, the lemma does not allow us to make the last step—the one leading to a machine representing a call—part of the deterministic execution that we provide. Instead, we are forced to provide a deterministic execution that ends one step earlier, then unfold \mathbf{E} and manually take the last step. Since this can sometimes be a little tedious, we also provide the following specialized variant of the lemma, where the deterministic execution ends in the actual call. Since the source program also must take at least one step in order for the `CALL` case to be applicable, we bake in part of Lemma 47 here.

Lemma 49 (Stepping to Call). If

- $m_a \in \mathbf{real}(c_a)$
- $m_a \xrightarrow{\iota}^n m'_a$
- $\exists j' <^{n-1} i$
- $n > 0$
- $\forall m_b \in \mathbf{real}(c_b). \exists m'_b, e_a, e_b, c'_a, c'_b.$
 $m_b \xrightarrow{\iota}^+ m'_b \wedge$
 $(m'_a, m'_b) \in \mathbf{cfg}(W.C)(G)(s)(\perp)(e_a, e_b)(c'_a, c'_b)(\eta_a, \eta_b) \wedge$
 $(e_a, e_b) \in \mathbf{U}(s)(G(s))(G(s))(\mathbf{K}_W(j)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s))(\tau)$

then $(c_a, c_b) \in \mathbf{E}_W(i)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)(\eta_a, \eta_b)(\tau)$.

(The condition $\exists j' <^{n-1} i$ is actually only needed for the case where $m_b \xrightarrow{\iota}^+ m'_b$ takes exactly one step. For simplicity, we require it in any case.)

What if we can't show that the resulting programs constitute calls in state s ? After all, that state may be completely out of synch with the machines by now. Fortunately, we can first apply Lemma 50, which lets us advance the current state arbitrarily when in fixed mode.

Lemma 50 (Advancing the State). If

$$(e_a, e_b) \in \mathbf{E}_W(i)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s')(\eta_a, \eta_b)(\tau) \wedge s' \sqsupseteq s$$

then $(e_a, e_b) \in \mathbf{E}_W(i)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)(\eta_a, \eta_b)(\tau)$.

For the sake of completeness, we also state a variant of Lemma 48 for ending the proof using \mathbf{E} 's **RET** case.

Lemma 51 (Stepping Target to Return). If

- $m_a \in \mathbf{real}(c_a)$
- $m_a \xrightarrow{\iota}^n m'_a$
- $s \sqsupseteq_{\text{pub}} s_0$
- $\exists j <^n i$
- $n > 0$
- $\forall m_b \in \mathbf{real}(c_b). \exists e_a, e_b, c'_a, c'_b, \mathbf{v}_a, \mathbf{v}_b.$
 $(m'_a, m_b) \in \mathbf{cfg}(W.C)(G)(s)(\perp)(e_a, e_b)(c'_a, c'_b)(\eta_a, \eta_b) \wedge$
 $(e_a, e_b) \in W.\mathbf{cqh}_a(s)(\mathbf{ret} \mathbf{v}_a \mathbf{k}_a) \times W.\mathbf{cqh}_b(s)(\mathbf{ret} \mathbf{v}_b \mathbf{k}_b) \wedge$
 $(\mathbf{v}_a, \mathbf{v}_b) \in \overline{G}(s)(\tau)$

then $(c_a, c_b) \in \mathbf{E}_W(i)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)(\eta_a, \eta_b)(\tau)$.

In combination, the lemmas shown in this section provide a powerful abstraction and in most cases keep us from having to work directly with the definition of \mathbf{E} . Moreover, although PILS are defined in an asymmetric small-step style, these lemmas essentially let us reason most of the time in the symmetric big-step style of PBs (cf. Section 3.6.5). We present a simple example that employs these lemmas in Section 4.6.

Remark. In all the stepping lemmas above, we only allow silent steps. In principle, we could also allow output steps in Lemma 49, such that both the target and the source execution produces the same sequence of outputs. This would require a slightly different proof (each target step that produces output must be matched immediately by the corresponding source step) but should be straightforward. Similarly, we could also have a straightforward combination of Lemmas 47 and 48 that allows output steps. For our compiler proofs, these generalizations are not important since there is basically only two places where we need to reason about an output step, namely the translation of an \mathcal{S} or \mathcal{I} output expression.

4.6 Example

As an example of using the PILS model, we now sketch the proof of a very simple example simulation between a source module $M_{\mathcal{S}}$ and a target module $M_{\mathcal{T}}$.

4.6.1 Modules

The source module is defined as follows.

$$\begin{aligned} \Gamma &:= F : \text{unit} + \text{nat} \rightarrow \text{nat} \\ \Gamma' &:= E : \text{nat} \times \text{nat} \rightarrow \text{nat} \\ M_{\mathcal{S}} &:= E = \lambda x. F (\text{inr } x.2) \\ \Gamma &\vdash M_{\mathcal{S}} : \Gamma' \end{aligned}$$

It consists of (and therefore exports) a single function E . This function takes as argument a pair from which it extracts the second component. It then right-injects this value into a sum, on which it calls the imported function F , returning whatever F returns (if it does).

The target module is defined as follows.

$$\begin{aligned}
M_{\mathcal{T}} &:= g \\
g &:= (\{(E, 0)\}, d) \\
d &:= \mathbb{E}(\text{ld aux}_1 [\text{arg} + 1]), \\
&\quad \mathbb{E}(\text{ld clo } 2), \\
&\quad \mathbb{E}(\text{alloc arg clo}), \\
&\quad \mathbb{E}(\text{sto } [\text{arg} + 0] \text{ clo}), \\
&\quad \mathbb{E}(\text{sto } [\text{arg} + 1] \text{ aux}_1), \\
&\quad \mathbb{E}(\text{lpc aux}_1), \\
&\quad \mathbb{E}(\text{bop } (-) \text{ aux}_1 \text{ aux}_1 8), \\
&\quad \mathbb{E}(\text{ld clo } [\text{aux}_1 + 0]), \\
&\quad \mathbb{E}(\text{jmp } [\text{clo} + 0]) \\
&\vdash M_{\mathcal{T}} : E
\end{aligned}$$

It consists of a single group, whose code pointer table in turn consists of a single entry, mapping E to offset 0 in the data block d . The code starting there first projects the second component of the argument pair and stores it in register aux_1 . Next, it allocates two heap cells for the sum value in the argument register. It writes 2 into the tag component, indicating "in right" (any non-zero value would do), and writes the contents of aux_1 into the value component. It then calculates the address of F 's entry in the label environment by getting the current program counter and subtracting 8 (there are 5 instructions before lpc , as well as the full group header of size 3, at the beginning of which F 's value is stored). Finally, the code loads F 's function value from there into register clo and performs an indirect jump to it (following the calling convention). Note that E does not modify registers ret and env and thus passes them on to F , such that F will return to wherever E was supposed to return to.

4.6.2 Proof

Our goal now is to show that the target module refines the source module relative to the given typing:

$$\Gamma \vdash M_{\mathcal{T}} \lesssim_{\mathcal{TS}} M_{\mathcal{S}} : \Gamma'$$

The first part of this concerns the module level and is very tedious, but we spell it out here for once. The second part concerns the behavior of the actual functions.

4.6.2.1 First Part

Before we construct a suitable local world, let us look at what we know about loaded configurations

$$\begin{aligned}
(c_{\mathcal{T}}^g, c_{\mathcal{T}}^l) &\in \mathbf{load}(M_{\mathcal{T}})(\Psi_{\mathcal{T}})(\{F \mapsto v_{\mathcal{T}}\}) \\
(c_{\mathcal{S}}^g, c_{\mathcal{S}}^l) &\in \mathbf{load}(M_{\mathcal{S}})(\Psi_{\mathcal{S}})(\{F \mapsto v_{\mathcal{S}}\})
\end{aligned}$$

as we will be given in a moment ($v_{\mathcal{T}}$ and $v_{\mathcal{S}}$ are the arbitrary values of the imported function F). By the respective implementations of **load**, we have:

$$\begin{aligned} c_{\mathcal{T}}^g &= (\emptyset, R, (\lambda_.0), [v_{\mathcal{T}} \mapsto w] \sqcup [\Psi_{\mathcal{T}}+2 \mapsto \Psi_{\mathcal{T}}+3]) \\ c_{\mathcal{T}}^l &= (\emptyset, \emptyset, \emptyset, [\Psi_{\mathcal{T}} \mapsto v_{\mathcal{T}}, \Psi_{\mathcal{T}}+2] \sqcup [\Psi_{\mathcal{T}}+3 \mapsto d]) \\ c_{\mathcal{S}}^g &= (\emptyset, \{F \mapsto v_{\mathcal{S}}\} \uplus M_{\mathcal{S}} \uplus \rho) \\ c_{\mathcal{S}}^l &= \emptyset \end{aligned}$$

The global part of \mathcal{T} 's heap consists of two cells. The first cell's address is F 's value ($v_{\mathcal{T}}$) and its contents is an arbitrary code pointer w . The second cell's address is E 's value, namely $\Psi_{\mathcal{T}}+2$, and its code pointer refers to the next cell. This next cell is contained in the local part of \mathcal{T} 's heap and marks the beginning of the data block d . This heap also contains the label environment (cells $\Psi_{\mathcal{T}}$ and $\Psi_{\mathcal{T}}+1$), through which the code accesses F .

For \mathcal{S} , things are simpler. The global configuration merely contains a label environment. It maps F to $v_{\mathcal{S}}$ and E to $\lambda x. F$ (in $x.2$)—and possibly further labels to further values (in ρ) but we do not care about those.

We now construct a local world $w \in \text{LWorld}^{\Omega_{\mathcal{T}\mathcal{S}} \cdot \top}$ such that its configuration relation insists that \mathcal{T} 's heap contains the local heap shown above (in $c_{\mathcal{T}}^l$). Since that heap depends on the load address ($\Psi_{\mathcal{T}}$) and the value of F ($v_{\mathcal{T}}$), which will only be given *after* we have constructed w , we must define w 's state space to range over all such possible values and then later pick the initial state to match whatever we were given. We could define states to be pairs of a load address and a value for F , but a more general pattern is to define states simply as heaps:

$$\begin{aligned} w.S &:= \text{Heap}_{\mathcal{T}} \\ w.\sqsubseteq &:= (=) \\ w.\sqsubseteq_{\text{pub}} &:= (=) \\ w.O &:= \mathbb{N}_{\text{awf}} \\ w.C(-)(-, s) &:= \{((\emptyset, \emptyset, \emptyset, s), \emptyset)\} \\ w.NS &:= \emptyset \\ w.NR(-)(-) &:= \emptyset \end{aligned}$$

As algebraic well-founded set we choose the natural numbers.

The configuration relation $w.C$ does not depend on the global state (or even the global knowledge) and thus w is trivially stable. Note how $w.C$ contains a single pair, where \mathcal{T} 's heap is exactly s .

Since the source program is so simple (*e.g.*, does not involve any memory), w 's states do not need to track any \mathcal{S} -related information. Moreover, the target implementation of E does not involve any other local memory, so the only states we are interested in are of the form $[\Psi_{\mathcal{T}} \mapsto v_{\mathcal{T}}, \Psi_{\mathcal{T}}+2] \sqcup [\Psi_{\mathcal{T}}+3 \mapsto d]$ (for some $\Psi_{\mathcal{T}}$ and $v_{\mathcal{T}}$). Since the code is not self-modifying, this heap won't ever change, which is why we allow no transitions.

Having constructed w , let us now suppose we are given anchors, import values and corresponding loadings $(c_{\mathcal{T}}^g, c_{\mathcal{T}}^l)$ and $(c_{\mathcal{S}}^g, c_{\mathcal{S}}^l)$ as stated and destructed above. We pick the obvious initial state s_0 that matches these configurations:

$$\begin{aligned} s_0 &:= ((s_{\mathcal{T}}^\circ, \emptyset, s_{\mathcal{S}}^\circ), [\Psi_{\mathcal{T}} \mapsto v_{\mathcal{T}}, \Psi_{\mathcal{T}}+2] \sqcup [\Psi_{\mathcal{T}}+3 \mapsto d]) \\ s_{\mathcal{T}}^\circ &:= (R, \{(\text{fun}_{\mathcal{T}} w, v_{\mathcal{T}}), (\text{fun}_{\mathcal{T}} (\Psi_{\mathcal{T}}+3), \Psi_{\mathcal{T}}+2)\}) \\ s_{\mathcal{S}}^\circ &:= \{F \mapsto v_{\mathcal{S}}\} \uplus M_{\mathcal{S}} \uplus \rho \end{aligned}$$

The registry in $s_{\mathcal{T}}^\circ$ contains two function entries, one for $v_{\mathcal{T}}$ with code pointer w (this is F) and one for $\Psi_{\mathcal{T}}+2$ with code pointer $\Psi_{\mathcal{T}}+3$ (this is E). It is easy to check that the following holds for any G , which concludes the first part.

$$\begin{aligned} (c_{\mathcal{T}}^g, c_{\mathcal{S}}^g) &\in \Omega_{\mathcal{TS}}.\mathcal{C}(G(-, [\Psi_{\mathcal{T}} \mapsto v_{\mathcal{T}}, \Psi_{\mathcal{T}}+2] \sqcup [\Psi_{\mathcal{T}}+3 \mapsto d]))(s_{\mathcal{T}}^\circ, \emptyset, s_{\mathcal{S}}^\circ) \\ (c_{\mathcal{T}}^l, c_{\mathcal{S}}^l) &\in w.\mathcal{C}(G)(s_0) \end{aligned}$$

4.6.2.2 Second Part

It remains to reason about the exported functions. By implementation of **vload** we have:

$$\begin{aligned} \Psi_{\mathcal{T}}+2 &\in \mathbf{vload}(M_{\mathcal{T}})(\Psi_{\mathcal{T}})(\{F \mapsto v_{\mathcal{T}}\})(E) \\ \lambda x. F \text{ (inr } x.2) &\in \mathbf{vload}(M_{\mathcal{S}})(\Psi_{\mathcal{S}})(\{F \mapsto v_{\mathcal{S}}\})(E) \end{aligned}$$

Note that we only need to consider future states of s_0 . This is crucial because otherwise $\Psi_{\mathcal{T}}+2$ may not have the meaning that we expect it to have (it might not be registered as a function value, or its code pointer might point to arbitrary code or invalid instructions or even non-existent memory).

So suppose we are given $G \in \mathbf{GK}_{w\uparrow}$ and $s \sqsupseteq s_0$ for which the imported values are related:

$$(v_{\mathcal{T}}, v_{\mathcal{S}}) \in \overline{G}(s)(\text{unit} + \text{nat} \rightarrow \text{nat})$$

We must show that the exported values are then related as well:

$$(\Psi_{\mathcal{T}}+2, \lambda x. F \text{ (inr } x.2)) \in \overline{G}(s)(\text{nat} \times \text{nat} \rightarrow \text{nat})$$

Showing that they count as functions in s , *i.e.*, that they match the fix query, is easy to do by relying on $s \sqsupseteq s^0$, because we explicitly registered $\Psi_{\mathcal{T}}+2$ as a closure in $s_{\mathcal{T}}^\circ$, which is only allowed to grow.

Showing that the functions are related by $G(s)$ takes more work. Since G is a valid global knowledge, we have $\mathbf{F}_{w\uparrow}(G) \subseteq G$. So it suffices to show the functions similar according to \mathbf{F} (the **GoodFuns** condition holds trivially), for which we consider their applications in a future state. Given $s' \sqsupseteq s$, $G' \supseteq G$ and initial continuations $\mathbf{k}_{\mathcal{T}}, \mathbf{k}_{\mathcal{S}}$, it suffices to show that there is an i such that

$$(e_{\mathcal{T}}, e_{\mathcal{S}}) \in \mathbf{E}_{w\uparrow}(i)(\mathbf{k}_{\mathcal{T}}, \mathbf{k}_{\mathcal{S}})(G')(s')(s')(\perp)(\text{nat}).$$

Here, the configurations $e_{\mathcal{T}}$ and $e_{\mathcal{S}}$ are applications of (the values for) E to related values according to \mathbf{U} (recall that these “applications” actually describe the machine *after* the jump or beta reduction):

$$\begin{aligned} e_{\mathcal{T}} &= (\Psi_{\mathcal{T}} + 3, \emptyset, \emptyset, \emptyset) \\ e_{\mathcal{S}} &= (\emptyset, \emptyset, \mathbf{k}_{\mathcal{S}}[F(\text{inr } v'_{\mathcal{S}}.2)]) \\ s'.R(\text{ret}) &= \mathbf{k}_{\mathcal{T}} \\ s'.R(\text{arg}) &= v'_{\mathcal{T}} \\ (v'_{\mathcal{T}}, v'_{\mathcal{S}}) &\in \overline{G'}(s')(\text{nat} \times \text{nat}) \end{aligned}$$

The choice of i somewhat depends on the lemmas that we want to use, so for now we will keep i abstract and pick its value at the end of the proof, when all constraints are known. (Of course we need to be careful that the choice does not depend on anything that got introduced in the meantime).

We start by applying Lemmas 43 and 45. This leaves us having to show

$$(e_{\mathcal{T}} \cdot c_{\mathcal{T}} \cdot \eta_{\mathcal{T}}, e_{\mathcal{S}} \cdot c_{\mathcal{S}} \cdot \eta_{\mathcal{S}}) \in \mathbf{E}_{w\uparrow}(i)(\mathbf{k}_{\mathcal{T}}, \mathbf{k}_{\mathcal{S}})(G')(s')(s')(\eta_{\mathcal{T}}, \eta_{\mathcal{S}})(\text{nat})$$

for any $(c_{\mathcal{T}}, c_{\mathcal{S}}) \in w\uparrow.\mathbf{C}(G')(s')$ and any $\eta_{\mathcal{T}}, \eta_{\mathcal{S}}$, but we may assume that there is $m_{\mathcal{T}} \in \mathbf{real}(e_{\mathcal{T}} \cdot c_{\mathcal{T}} \cdot \eta_{\mathcal{T}})$ (we don't care about a source machine for now).

Next, we apply Lemmas 50 and 49 such that it suffices to show the existence of $n > 0$, j , s'' and $m'_{\mathcal{T}}$ satisfying the following:

1. $m_{\mathcal{T}} \xrightarrow{\iota}^n m'_{\mathcal{T}}$
2. $\exists j' <^{n-1} i$
3. $s'' \sqsupseteq s'$
4. $\forall m_{\mathcal{S}} \in \mathbf{real}(e_{\mathcal{S}} \cdot c_{\mathcal{S}} \cdot \eta_{\mathcal{S}}). \exists m'_{\mathcal{S}}, e'_{\mathcal{T}}, e'_{\mathcal{S}}, c'_{\mathcal{T}}, c'_{\mathcal{S}}.$
 $m_{\mathcal{S}} \xrightarrow{\iota}^+ m'_{\mathcal{S}} \wedge$
 $(m'_{\mathcal{T}}, m'_{\mathcal{S}}) \in \mathbf{cfg}(w\uparrow.\mathbf{C})(G')(s'')(\perp)(e'_{\mathcal{T}}, e'_{\mathcal{S}})(c'_{\mathcal{T}}, c'_{\mathcal{S}})(\eta_{\mathcal{T}}, \eta_{\mathcal{S}}) \wedge$
 $(e'_{\mathcal{T}}, e'_{\mathcal{S}}) \in \mathbf{U}(s'')(G'(s''))(G'(s''))(\mathbf{K}_{w\uparrow}(j)(\mathbf{k}_{\mathcal{T}}, \mathbf{k}_{\mathcal{S}})(G')(s')(s''))(\text{nat})$

(1) First, due to $s' \sqsupseteq s_0$ we know that the heap contains $[\Psi_{\mathcal{T}} + 3 \mapsto d]$. This means that $m_{\mathcal{T}}$'s program counter (which equals that of $e_{\mathcal{T}}$) currently points to the first instruction of E 's code. From $(v'_{\mathcal{T}}, v'_{\mathcal{S}}) \in \overline{G'}(s')(\text{nat} \times \text{nat})$ we know by definition of the value closure and $\Omega_{\mathcal{T}\mathcal{S}}$ that $v'_{\mathcal{T}}$ is registered in the global state's database of \mathcal{T} values, *i.e.*, we have $v'_{\mathcal{T}} \in s'.\rho(\text{pair}_{\mathcal{T}} v_{\mathcal{T}}^1 v_{\mathcal{T}}^2)$ for some machine values $v_{\mathcal{T}}^1, v_{\mathcal{T}}^2$. This, and the connection between s' and $m_{\mathcal{T}}$, tells us that the target heap contains $v_{\mathcal{T}}^2$ at address $v'_{\mathcal{T}} + 1$. Hence, after one (deterministic) step of $m_{\mathcal{T}}$, which performs the first `ld` instruction, register `aux1` holds $v_{\mathcal{T}}^2$ (and of course the program counter has advanced to the second `ld` instruction)—otherwise nothing changes.

Two further steps write the constant 2 into register `clo` and the address of a newly allocated heap block of size 2 into register `arg`. The next two steps write to

this block of memory. The subsequent `lpc` instruction has the effect of writing $\Psi_{\mathcal{T}+8}$ into register `aux1`, so that the `bop` instruction updates `aux1` to $\Psi_{\mathcal{T}}$.

In order to reason about the remaining two instructions, note that due to $s' \sqsupseteq s_0$ we know that the machine heap contains both $[\Psi_{\mathcal{T}} \mapsto v_{\mathcal{T}}]$ and $[v_{\mathcal{T}} \mapsto w]$. Accordingly, the next two steps write $v_{\mathcal{T}}$, the value of F , into register `clo` and set the program counter to its first instruction at address w .

All in all, we have the end-to-end deterministic execution $m_{\mathcal{T}} \xrightarrow{\iota}^9 m'_{\mathcal{T}}$, where

$$\begin{aligned} m'_{\mathcal{T}} &\in \mathbf{real}(e'_{\mathcal{T}} \cdot c'_{\mathcal{T}} \cdot \eta_{\mathcal{T}}) \\ e'_{\mathcal{T}} &= (w, \emptyset, \emptyset, \emptyset) \end{aligned}$$

and $c'_{\mathcal{T}}$ is obtained from $c_{\mathcal{T}}$ by extending its heap with the freshly allocated (and thus disjoint) chunk $[w' \mapsto 2, v_{\mathcal{T}}^2]$ and changing its register file $s'.R$ to

$$s'.R[\mathbf{arg} \mapsto w'][\mathbf{aux}_1 \mapsto \Psi_{\mathcal{T}}][\mathbf{clo} \mapsto v_{\mathcal{T}}].$$

(2) With $n = 9$, it is easy to satisfy $\exists j' <^{n-1} i$ by choosing $i = 8$ and $j' = 0$ (recall that we kept i abstract until now).

(3) We pick s'' basically such that it matches the resulting machine $m'_{\mathcal{T}}$ above. Concretely, we derive s'' from s' by making two changes. First, we set its R component to $s'.R[\mathbf{arg} \mapsto w'][\mathbf{aux}_1 \mapsto \Psi_{\mathcal{T}}][\mathbf{clo} \mapsto v_{\mathcal{T}}]$. Second, we extend the value registry $s'.\rho$ with an entry saying that w' now represents $\text{inr}_{\mathcal{T}} v_{\mathcal{T}}^2$. It is trivial to check that $s'' \sqsupseteq s'$.

(4) Let us focus on the source program. From $(v'_{\mathcal{T}}, v'_{\mathcal{S}}) \in \overline{G'}(s')(\mathbf{nat} \times \mathbf{nat})$ we know by definition of the value closure and $\Omega_{\mathcal{T}\mathcal{S}}$ that $v'_{\mathcal{S}}$ equals $\langle v_{\mathcal{S}}^1, v_{\mathcal{S}}^2 \rangle$ for some source values $v_{\mathcal{S}}^1, v_{\mathcal{S}}^2$ that are related to $v_{\mathcal{T}}^1, v_{\mathcal{S}}^2$ from above. In particular, we have $(v_{\mathcal{T}}^2, v_{\mathcal{S}}^2) \in \overline{G'}(s')(\mathbf{nat})$.

Moreover, from $s' \sqsupseteq s \sqsupseteq s_0$ we know that $c_{\mathcal{S}}$ consists of the environment $s_{\mathcal{S}}^{\circ}$, which maps F to $v_{\mathcal{S}}$. Since $v_{\mathcal{S}}$ is related by $\overline{G'}(s')$ at function type, it must have the form $\text{fix } f(x).e$.

All in all, we thus easily derive $m_{\mathcal{S}} \hookrightarrow^+ m'_{\mathcal{S}}$, where:

$$\begin{aligned} m'_{\mathcal{S}} &\in \mathbf{real}(e'_{\mathcal{S}} \cdot c_{\mathcal{S}} \cdot \eta_{\mathcal{S}}) \\ e'_{\mathcal{S}} &= \mathbf{k}_{\mathcal{S}}[e[v_{\mathcal{S}}/f][\text{inr } v_{\mathcal{S}}^2/x]] \end{aligned}$$

Note that this represents a function call in s'' (in fact, in any state):

$$e'_{\mathcal{S}} \in w \uparrow. \text{cqhb}(s'')(\text{app } v_{\mathcal{S}} (\text{inr } v_{\mathcal{S}}^2) \mathbf{k}_{\mathcal{S}})$$

It is easy to see that $e'_{\mathcal{T}}$ from (1) also represents a function call in s'' :

$$e'_{\mathcal{T}} \in w \uparrow. \text{cqha}(s'')(\text{app } v_{\mathcal{T}} w' \mathbf{k}_{\mathcal{T}})$$

Moreover, using monotonicity, the functions are obviously related by $\overline{G'}(s'')$. The arguments w' and $\text{inr } v_{\mathcal{S}}^2$ are as well, because they represent right-injections of $v_{\mathcal{T}}^2$ and $v_{\mathcal{S}}^2$, respectively, which are related by $\overline{G'}(s')$.

Consequently, the **U** condition in (4) is satisfied if we can show that the continuations $\mathbf{k}_{\mathcal{T}}$ and $\mathbf{k}_{\mathcal{S}}$ are related by $\mathbf{K}_{w\uparrow(j)}(\mathbf{k}_{\mathcal{T}}, \mathbf{k}_{\mathcal{S}})(G')(s')(s'')$ for an arbitrary j of our choosing. Indeed, this is an easy general lemma: the initial continuations are always related (at any public extension of the initial state).

Lemma 52.

$$\forall s \sqsupseteq_{\text{pub}} s_0. (\mathbf{k}_a, \mathbf{k}_b) \in \mathbf{K}_W(i)(\mathbf{k}_a, \mathbf{k}_b)(s_0)(s)(\tau)(\tau)$$

Proof. By unfolding the definitions and immediately using **E**'s **RET** case. \square

To conclude the proof of (4)—and thus the proof of the example—it remains to establish the **cfg** condition, in particular that $c'_{\mathcal{T}}$ and $c'_{\mathcal{S}}$ satisfy the world. This is very easy to check.

4.7 Metatheory

4.7.1 Basics

Definition 26 (Least global knowledge). For $W \in \text{World}_{A,B}$ and a function $g \in W.S \rightarrow \text{VRelF}_{A,B}$, we define $[W]_g \in W.S \rightarrow \text{VRelF}_{A,B}$ as the least global knowledge containing g . Formally it is constructed as a greatest fixed point of a monotone function such that the following holds (using a pointwise union over functions):

$$[W]_g = g \cup \mathbf{F}_W([W]_g) \cup W.\text{NR}([W]_g)$$

We write $[W]$ short for $[W]_{\emptyset}$.

Lemma 53 (Least global knowledge). For any W and g , we have $[W]_g \in \mathbf{GK}_W$ if the following hold:

1. g is monotone: $\forall s', s. s' \sqsupseteq s \implies g(s') \supseteq g(s)$
2. $\forall \mathbf{n} \in W.\text{NS}. \forall s. g(s)(\mathbf{n}) = \emptyset$

. In particular, we always have $[W] \in \mathbf{GK}_W$.

Moreover, $[W]_g \subseteq G$ for any $G \in \mathbf{GK}_W$ with $g \subseteq G$.

Lemma 54 (Composition of **E** and **K**). If

1. $(e_a, e_b) \in \mathbf{E}_W(i)(\mathbf{k}_a, \mathbf{k}_b)(G)(s)(s)(\phi)(\tau)$ and
2. $(\mathbf{k}_a, \mathbf{k}_b) \in \mathbf{K}_W(j)(\mathbf{k}_a^0, \mathbf{k}_b^0)(G)(s_0)(s)(\tau)(\tau_0)$,

then $(e_a, e_b) \in \mathbf{E}_W(i+j)(\mathbf{k}_a^0, \mathbf{k}_b^0)(G)(s_0)(s)(\phi)(\tau_0)$.

Proof. By a fairly straightforward coinduction, at the same time proving a similar property for **K**. \square

Analogously to Section 3.7.1, we define a version of \mathbf{E} (in terms of the original \mathbf{E}) that parameterizes over the value relation by which functions are related in the **CALL** case.

Definition 27.

$$\begin{aligned}
& \mathbf{E}_W^{\mathcal{G}} \in \text{ES}_{A,B}^W \\
& \mathbf{E}_W^{\mathcal{G}}(i)(\mathbf{k}_a^0, \mathbf{k}_b^0)(G)(s_0)(s)(\phi)(\tau) = \{ (e_a, e_b) \mid G \in \mathbf{GK}_W \implies \\
& \quad \forall c_a, c_b, \eta_a, \eta_b. \forall (m_a, m_b) \in \text{cfg}(W.C)(G)(s)(\phi)(e_a, e_b)(c_a, c_b)(\eta_a, \eta_b). \\
& \quad (\text{ERR}) \exists m'_b. m_b \xrightarrow{\iota^*} m'_b \wedge m'_b \in \text{error} \\
& \quad \vee (\text{RET}) \exists m'_b, s', \mathbf{v}_a, \mathbf{v}_b, e'_a, e'_b, c'_a, c'_b. m_b \xrightarrow{\iota^*} m'_b \wedge s' \sqsupseteq s \wedge s' \sqsupseteq_{\text{pub}} s_0 \wedge \\
& \quad (m_a, m'_b) \in \text{cfg}(W.C)(G)(s')(\perp)(e'_a, e'_b)(c'_a, c'_b)(\eta_a, \eta_b) \wedge (\mathbf{v}_a, \mathbf{v}_b) \in \overline{G}(s')(\tau) \wedge \\
& \quad (e'_a, e'_b) \in W.\text{cqh}_a(s')(\text{ret } \mathbf{v}_a \mathbf{k}_a^0) \times W.\text{cqh}_b(s')(\text{ret } \mathbf{v}_b \mathbf{k}_b^0) \\
& \quad \vee (\text{STEP}) m_a \notin \text{halted} \wedge \forall t, m'_a. m_a \xrightarrow{t} m'_a \implies \exists i', e'_a, e'_b, c'_a, c'_b, m'_b, m''_b, s', \phi'. \\
& \quad m_b \xrightarrow{\iota^*} m'_b \wedge s' \sqsupseteq s \wedge (m'_a, m''_b) \in \text{cfg}(W.C)(G)(s')(\phi')(e'_a, e'_b)(c'_a, c'_b)(\eta_a, \eta_b) \wedge \\
& \quad (\text{REC}) (m'_b \xrightarrow{t} m''_b \wedge \vee m'_b = m''_b \wedge t = \iota \wedge i' < i) \wedge \phi' \neq \perp \wedge \\
& \quad (e'_a, e'_b) \in \mathbf{E}_W(i')(\mathbf{k}_a^0, \mathbf{k}_b^0)(G)(s_0)(s')(\phi')(\tau) \\
& \quad \vee (\text{CALL}) m'_b \xrightarrow{t} m''_b \wedge \phi' = \perp \wedge \\
& \quad (e'_a, e'_b) \in \mathbf{U}(s')(\mathcal{G}(s'))(G(s'))(\mathbf{K}_W(i')(\mathbf{k}_a^0, \mathbf{k}_b^0)(G)(s_0)(s'))(\tau) \}
\end{aligned}$$

Lemma 55 (External call). If $G = \mathbf{F}_W(G) \cup W.\text{NR}(G) \cup \mathcal{G}$, then

$$\mathbf{E}_W(i)(\mathbf{k}_a, \mathbf{k}_b)(G) = \mathbf{E}_W^{\mathcal{G}}(i)(\mathbf{k}_a, \mathbf{k}_b)(G).$$

Proof. The proof is fairly similar to the PB version (Lemma 12 in Section 3.7.1). The interesting direction is “ \subseteq ”. Due to the single-step style of PILS, we don’t need to do any induction on steps here, in contrast to the PB proof. As before, the crucial part is the composition of \mathbf{E} and \mathbf{K} in the **CALL** case (here via Lemma 54).

Note that although the difference between \mathbf{E} and $\mathbf{E}^{\mathcal{G}}$ is apparently “lost” if the programs are related via the **REC** case, it can of course be restored by immediately re-applying the lemma. Indeed, we will do so in the proof of Lemma 56. \square

4.7.2 Adequacy

With the external call lemma at hand, we can now approach adequacy. We are only interested in adequacy of $\preceq_{\mathcal{TS}}$, the target-source module refinement. As a lemma, we nevertheless prove adequacy of \mathbf{E} generically for any instance of the typed model (because we can). It states that if e_a and e_b are related by \mathbf{E} at a *consistent* global knowledge and relative to *terminal* continuations, then any completions of e_a and e_b to proper machines m_a and m_b satisfy the standard refinement property. *Consistent* here means that, modulo type names, the “external part” of G is empty—see condition 2 below. *Terminal* here means that as soon as the continuations obtain control, the programs are halted—see conditions 3 and 4 below. Otherwise, the relatedness of e_a and e_b would only account for their behavior until they reach their

continuations, while $\mathbf{behav}(m_a)$ and $\mathbf{behav}(m_b)$ of course include the behavior of these continuations.

Lemma 56 (Generic adequacy of typed **E**). In the context of $W \in \text{World}_{A,B}$, if

1. $(e_a, e_b) \in \mathbf{E}_W(i)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)(\phi)(\tau)$
2. $G \in \mathbf{GK}_W \wedge G = \mathbf{F}(G) \cup W.\mathbf{NR}(G)$
3. $\forall s \sqsupseteq s_0. \forall v_a. \forall e'_a \in W.\text{cqh}_a(s)(\text{ret } v_a \mathbf{k}_a). \forall c'_a. \mathbf{real}(e'_a \cdot c'_a) \subseteq \mathbf{halted}$
4. $\forall s \sqsupseteq s_0. \forall v_b. \forall e'_b \in W.\text{cqh}_b(s)(\text{ret } v_b \mathbf{k}_b). \forall c'_b. \mathbf{real}(e'_b \cdot c'_b) \subseteq \mathbf{halted}$
5. $\eta_a \in \mathbf{frame} \wedge \eta_b \in \mathbf{frame}$
6. $(m_a, m_b) \in \mathbf{cfg}(G)(s)(\phi)(e_a, e_b)(c_a, c_b)(\eta_a, \eta_b)$

then $\mathbf{behav}(m_a) \subseteq \mathbf{behav}(m_b)$.

Proof. By coinduction on m_b 's behavior, with a nested well-founded induction on i . We critically rely on the external call Lemma 55 to rule out any use of the **CALL** disjunct in (1). \square

Theorem 13 (Adequacy of $\lesssim_{\mathcal{TS}}$). If

1. $\epsilon \vdash M_a \lesssim_{\mathcal{TS}} M_b : \Gamma$
2. $\Gamma(F_{\text{main}}) = \text{unit} \rightarrow \tau$
3. $m_a = \text{load}(M_a)$
4. $m_b = \text{load}(M_b)$

then $\mathbf{behav}(m_a) \subseteq \mathbf{behav}(m_b)$.

Proof. We define the following configurations:

$$\begin{aligned} c_a^1 &:= (\emptyset, R, m_a.q, h_1) \\ c_a^2 &:= (\emptyset, \emptyset, \emptyset, h_2) \\ c_b^1 &:= (\emptyset, m_b.\sigma, \emptyset) \\ c_b^2 &:= \emptyset \end{aligned}$$

Here, R is $m_a.R$ and h_1, h_2 are the global and local heap (respectively) returned by load' in the definition of $\text{load}(M_a)$.

With the help of (3) and (4), it is easy to show that these configurations are proper loadings of the given modules (using load address 0 in the case of M_a):

5. $(c_a^1, c_a^2) \in \mathbf{cload}(M_a)(0)(\emptyset)$
6. $(c_b^1, c_b^2) \in \mathbf{cload}(M_b)(1)(\emptyset)$

Using these, we instantiate (1). This yields a local world $w \in \text{LWorld}(\Omega_{\mathcal{TS}})$ and, for $G := [w^\uparrow]$, a state s_0 and values v_a, v_b such that:

7. $(c_a^1, c_b^1) \in \Omega_{\mathcal{TS}}.\mathbf{C}(G(-, \text{snd } s_0))$
8. $(c_a^2, c_b^2) \in w.\mathbf{C}(G)(s_0)$
9. $v_a \in \mathbf{vload}(M_a)(0)(\emptyset)(F_{\text{main}}) \wedge v_b \in \mathbf{vload}(M_b)(1)(\emptyset)(F_{\text{main}})$
10. $(v_a, v_b) \in \overline{G}(s_0)(\text{unit} \rightarrow \tau)$

Note that (10) implies $(v_a, v_b) \in \mathbf{F}_{w^\uparrow}(G)(s_0)(\text{unit} \rightarrow \tau)$ due to Definition 26, and also that $v_b = \text{fix } f(x).e$ (for some f, x, e).

We know that m_a and m_b look as follows (for a particular n):

$$\begin{aligned} m_a &= (n, \emptyset, \emptyset, \emptyset) \cdot c_a^1 \cdot c_a^2 \\ m_b &= (\emptyset, \emptyset, v_b \langle \rangle) \cdot c_b^1 \cdot c_b^2 \end{aligned}$$

Now, let $e_a = (n, \emptyset, \emptyset, \emptyset)$ and $e_b = (\emptyset, \emptyset, e[v_b/f][\langle \rangle/x])$. Then it is easy to see that e_a and e_b represent the following “applications” of v_a and v_b :

$$\begin{aligned} e_a &\in w^\uparrow.\text{cqh}_a(s_0)(\text{app } v_a \ R(\text{arg}) \ R(\text{ret})) \\ e_b &\in w^\uparrow.\text{cqh}_b(s_0)(\text{app } v_b \ \langle \rangle \ \bullet) \end{aligned}$$

Since the arguments are related by $\overline{G}(s_0)(\text{unit})$ trivially and since (7) and (8) compose to $(c_a^1 \cdot c_a^2, c_b^1 \cdot c_b^2) \in w^\uparrow.\mathbf{C}(G)(s_0)$, we can instantiate $(v_a, v_b) \in \mathbf{F}_{w^\uparrow}(G)(s_0)(\text{unit} \rightarrow \tau)$ such that we obtain the following:

$$(e_a, e_b) \in \mathbf{E}_{w^\uparrow}(i)(R(\text{ret}), \bullet)(G)(s_0)(s_0)(\perp)(\tau)$$

It suffices to apply Lemma 56 to this, using empty frames. Note that $R(\text{ret}) = 0$ and thus the target continuation is terminal as required (a target machine cannot step if the program counter is 0). The source continuation \bullet is terminal as well, as a source machine cannot step if its expression is a value. \square

4.7.3 Modularity

We now sketch the proof of modularity, starting with some lemmas about the product of local worlds. Recall its definition, as well as that of $G_{\langle 1 \rangle}^{s_2}$ and $G_{\langle 2 \rangle}^{s_1}$ from Figure 4.14.

4.7.3.1 Generic modularity of \mathbf{E}

Lemma 57 (PILS version of Lemma 9). If $w_1, w_2 \in \text{stable}(\Omega)$ and $w = w_1 \otimes w_2$, then:

1. $s_2 \sqsubseteq_{\text{pub}} s_2^0 \implies$
 $\mathbf{E}_{w_1^\uparrow}(i)(\mathbf{k}_a, \mathbf{k}_b)(G_{\langle 1 \rangle}^{s_2})(s^0, s_1^0)(s, s_1)(\perp) \subseteq$
 $\mathbf{E}_{w^\uparrow}(i, 0)(\mathbf{k}_a, \mathbf{k}_b)(G)(s^0, (s_1^0, s_2^0))(s, (s_1, s_2))(\perp)$

2. $s_1 \sqsupseteq_{\text{pub}} s_1^0 \implies$
 $\mathbf{E}_{w_2\uparrow}(i)(\mathbf{k}_a, \mathbf{k}_b)(G_{\langle 2 \rangle}^{s_1^0})(s^0, s_2^0)(s, s_2)(\perp) \subseteq$
 $\mathbf{E}_{w\uparrow}(0, i)(\mathbf{k}_a, \mathbf{k}_b)(G)(s^0, (s_1^0, s_2^0))(s, (s_1, s_2))(\perp)$
3. $s_2 \sqsupseteq_{\text{pub}} s_2^0 \wedge s' \sqsupseteq s \wedge (c_a, c_b) \in w_2.C(G_{\langle 2 \rangle}^{s_1^0})(s, s_2) \implies$
 $\mathbf{E}_{w_2\uparrow}(i)(\mathbf{k}_a, \mathbf{k}_b)(G_{\langle 1 \rangle}^{s_2^0})(s^0, s_1^0)(s', s_1')(c_a \cdot \eta_a, c_b \cdot \eta_b) \subseteq$
 $\mathbf{E}_{w\uparrow}(i, 0)(\mathbf{k}_a, \mathbf{k}_b)(G)(s^0, (s_1^0, s_2^0))(s', (s_1', s_2'))(\eta_a, \eta_b)$
4. $s_1 \sqsupseteq_{\text{pub}} s_1^0 \wedge s' \sqsupseteq s \wedge (c_a, c_b) \in w_1.C(G_{\langle 1 \rangle}^{s_2^0})(s, s_1) \implies$
 $\mathbf{E}_{w_2\uparrow}(i)(\mathbf{k}_a, \mathbf{k}_b)(G_{\langle 2 \rangle}^{s_1^0})(s^0, s_2^0)(s', s_2')(c_a \cdot \eta_a, c_b \cdot \eta_b) \subseteq$
 $\mathbf{E}_{w\uparrow}(0, i)(\mathbf{k}_a, \mathbf{k}_b)(G)(s^0, (s_1^0, s_2^0))(s', (s_1, s_2'))(\eta_a, \eta_b)$

We actually only care about first two, but need to prove the last two at the same time. Moreover, in contrast to PBs, here we cannot prove (1) and (2) separately, in each relying only on the stability of the respective local world. Why is this? In PBs, due to **GK** being defined in terms of the world's local knowledge, we could prove Lemma 8 a priori, and then prove Lemma 9 on top of that. In PILS, **GK** is mutually recursive with **E**. Proving the analogue of Lemma 8 hence already requires the analogue of Lemma 9. More precisely, the proof of “ $\mathbf{E}_{w_1\uparrow} \subseteq \mathbf{E}_{w\uparrow}$ ” relies on stability of w_2 . In order to make use of this stability, we also need to know part (2) of the lemma below. However, proving that part in turn requires “ $\mathbf{E}_{w_2\uparrow} \subseteq \mathbf{E}_{w\uparrow}$ ”. So we actually require stability of both w_1 and w_2 and have to prove everything simultaneously (also for **F** and **K**).

Lemma 58 (PILS version of Lemma 8). If $w_1, w_2 \in \text{stable}(\Omega)$ and $w = w_1 \otimes w_2$ and $G \in \mathbf{GK}_{w\uparrow}$, then:

1. $\forall s_2 \in w_2.S. G_{\langle 1 \rangle}^{s_2^0} \in \mathbf{GK}_{w_1\uparrow}$
2. $\forall s_1 \in w_1.S. G_{\langle 2 \rangle}^{s_1^0} \in \mathbf{GK}_{w_2\uparrow}$

Lemma 59 (PILS version of part of Lemma 10). If $w_1, w_2 \in \text{stable}(\Omega)$ and $w = w_1 \otimes w_2$, then $w \in \text{stable}(\Omega)$.

4.7.3.2 Modularity of $\lesssim_{\mathcal{TS}}$

Based on these generic results, we now sketch the proof of the key modularity property that we are interested in, namely that $\lesssim_{\mathcal{TS}}$ is preserved under linking:

$$\frac{\Gamma \vdash M_a^1 \lesssim_{\mathcal{TS}} M_b^1 : \Gamma_1 \quad \Gamma, \Gamma_1 \vdash M_a^2 \lesssim_{\mathcal{TS}} M_b^2 : \Gamma_2}{\Gamma \vdash \text{link}(M_a^1, M_a^2) \lesssim_{\mathcal{TS}} \text{link}(M_b^1, M_b^2) : \Gamma_1, \Gamma_2}$$

(Here we drop a few well-formedness side conditions. The full theorem is stated at the end of the section as Theorem 14.) This proof necessarily involves several language details, as otherwise we could have proven a generic version.

Let $M_a := \text{link}(M_a^1, M_a^2)$ and $M_b := \text{link}(M_b^1, M_b^2)$. We are first given a countably infinite set \mathcal{N} of type names, which we split into two disjoint infinite parts \mathcal{N}_1 and \mathcal{N}_2 . Using these to instantiate the two premises gives us stable local worlds w_1 and w_2 , respectively, with disjoint type name sets. As our local world, we choose their product $w := w_1 \otimes w_2$, which is stable as well by Lemma 59.

Next we are given anchors, import lists, and pairs of configurations corresponding to appropriate loadings of the modules:

$$\begin{aligned} (c_a^g, c_a^l) &\in \mathbf{cload}(M_a)(\Psi_a)(\gamma_a) \\ (c_b^g, c_b^l) &\in \mathbf{cload}(M_b)(\Psi_b)(\gamma_b) \end{aligned}$$

Because we know the \mathcal{T} and \mathcal{S} implementations of **cload**, we know what these configurations look like:

$$\begin{aligned} c_a^g &= (\emptyset, R, q, h \sqcup h^g) & c_b^g &= (\emptyset, \gamma_b \uplus M_b \uplus \rho, \emptyset) \\ c_a^l &= (\emptyset, \emptyset, \emptyset, h^l) & c_b^l &= \emptyset \end{aligned}$$

$$\text{where } (h^g, h^l, x) = \text{load}'(M_a)(\Psi_a)(\text{map snd } \gamma_a)$$

We can now take these loadings of M_a and M_b , and cut away some data in order to obtain loadings of M_a^1 and M_b^1 , and similarly of M_a^2 and M_b^2 . Formally, we define

$$\begin{aligned} \Psi_a^1 &:= \Psi_a & \Psi_b^1 &:= 1 \quad (= \Psi_b) \\ \gamma_a^1 &:= \gamma_a & \gamma_b^1 &:= \gamma_b \\ c_a^{g,1} &:= (\emptyset, R, q, h \sqcup h_1^g) & c_b^{g,1} &:= (\emptyset, \gamma_b \uplus M_b^1 \uplus (M_b^2 \uplus \rho), \emptyset) \\ c_a^{l,1} &:= (\emptyset, \emptyset, \emptyset, h_1^l) & c_b^{l,1} &:= \emptyset \\ \\ \Psi_a^2 &:= \Psi_a + \text{size}(h_1^g \sqcup h_1^l) & \Psi_b^2 &:= 1 \\ \gamma_a^2 &:= \gamma_a, x_1 & \gamma_b^2 &:= \gamma_b, M_b^1 \\ c_a^{g,2} &:= (\emptyset, R, q, (h \sqcup h_1^g) \sqcup h_2^g) & c_b^{g,2} &:= (\emptyset, \gamma_b^2 \uplus M_b^2 \uplus \rho, \emptyset) \\ c_a^{l,2} &:= (\emptyset, \emptyset, \emptyset, h_2^l) & c_b^{l,2} &:= \emptyset \end{aligned}$$

$$\begin{aligned} \text{where } (h_1^g, h_1^l, x_1) &= \text{load}'(M_a^1)(\Psi_a^1)(\text{map snd } \gamma_a^1) \\ (h_2^g, h_2^l, x_2) &= \text{load}'(M_a^2)(\Psi_a^2)(\text{map snd } \gamma_a^2) \end{aligned}$$

and are then able to show the following:

$$\begin{aligned} (c_a^{g,1}, c_a^{l,1}) &\in \mathbf{cload}(M_a^1)(\Psi_a^1)(\gamma_a^1) \\ (c_b^{g,1}, c_b^{l,1}) &\in \mathbf{cload}(M_b^1)(\Psi_b^1)(\gamma_b^1) \\ \\ (c_a^{g,2}, c_a^{l,2}) &\in \mathbf{cload}(M_a^2)(\Psi_a^2)(\gamma_a^2) \\ (c_b^{g,2}, c_b^{l,2}) &\in \mathbf{cload}(M_b^2)(\Psi_b^2)(\gamma_b^2) \end{aligned}$$

This may look complicated but is actually quite simple. There aren't any real choices involved here—everything fits together in a canonical way. For instance, Ψ_a^2 is exactly

the address in the heap where $\text{load}'(\text{link}(M_a^1, M_a^2))$ places the environment for and code of M_a^2 .

Note that, assuming well-typedness, (the functions exported by) \mathbf{Mod}_b^1 and M_b^2 are disjoint. Hence c_b^g , $c_b^{g,1}$, and $c_b^{g,2}$ are actually all the same, just written in a different way. Also note that $h_1^g \sqcup h_2^g = h^g$ and $h_1^l \sqcup h_2^l = h^l$ and $x_1, x_2 = x$ by definition of load' .

We now further instantiate the premises with these constructed loadings. The first premise yields an initial state $(s_1^\circ, s_1) \in w_1 \uparrow \mathbf{S}$ such that:

1. $\forall G \in \mathbf{GK}_{w_1 \uparrow}. (c_a^{g,1}, c_b^{g,1}) \in \Omega_{\mathcal{TS}}.C(G(-, s_1))(s_1^\circ)$
2. $\forall G \in \mathbf{GK}_{w_1 \uparrow}. (c_a^{l,1}, c_b^{l,1}) \in w_1.C(G)(s_1^\circ, s_1)$
3. $\Omega_{\mathcal{TS}}.\text{rqh}(s_1^\circ) = \emptyset$
4. $\forall f': \tau' \in \Gamma_1. \exists (v_a, v_b) \in \mathbf{vload}(M_a^1)(\Psi_a^1)(\gamma_a^1)(f') \times \mathbf{vload}(M_b^1)(\Psi_b^1)(\gamma_b^1)(f').$
 $\forall s' \sqsupseteq (s_1^\circ, s_1). \forall G \in \mathbf{GK}_{w_1 \uparrow}. (s_1 \in \text{sat}(G) \wedge \forall f: \tau \in \Gamma. (\gamma_a^1 f, \gamma_b^1 f) \in \overline{G}(s')(\tau)) \implies$
 $(v_a, v_b) \in \overline{G}(s')(\tau')$

Similarly, the second premise yields a state $(s_2^\circ, s_2) \in w_2 \uparrow \mathbf{S}$ such that:

5. $\forall G \in \mathbf{GK}_{w_2 \uparrow}. (c_a^{g,2}, c_b^{g,2}) \in \Omega_{\mathcal{TS}}.C(G(-, s_2))(s_2^\circ)$
6. $\forall G \in \mathbf{GK}_{w_2 \uparrow}. (c_a^{l,2}, c_b^{l,2}) \in w_2.C(G)(s_2^\circ, s_2)$
7. $\Omega_{\mathcal{TS}}.\text{rqh}(s_2^\circ) = \emptyset$
8. $\forall f': \tau' \in \Gamma_2. \exists (v_a, v_b) \in \mathbf{vload}(M_a^2)(\Psi_a^2)(\gamma_a^2)(f') \times \mathbf{vload}(M_b^2)(\Psi_b^2)(\gamma_b^2)(f').$
 $\forall s' \sqsupseteq (s_2^\circ, s_2). \forall G \in \mathbf{GK}_{w_2 \uparrow}. (s_2 \in \text{sat}(G) \wedge \forall f: \tau \in \Gamma, \Gamma_1. (\gamma_a^2 f, \gamma_b^2 f) \in \overline{G}(s')(\tau)) \implies$
 $(v_a, v_b) \in \overline{G}(s')(\tau')$

Note that (4) only cares about values provided for Γ , while (8) cares about values provided for Γ, Γ_1 (recall that linking is asymmetric—the right module may depend on the left but not vice versa).

We now need to come up with a state $(s^\circ, s) \in w \uparrow \mathbf{S}$ such that:

9. $\forall G \in \mathbf{GK}_{w \uparrow}. (c_a^g, c_b^g) \in \Omega_{\mathcal{TS}}.C(G(-, s))(s^\circ)$
10. $\forall G \in \mathbf{GK}_{w \uparrow}. (c_a^l, c_b^l) \in w.C(G)(s^\circ, s)$
11. $\Omega_{\mathcal{TS}}.\text{rqh}(s^\circ) = \emptyset$
12. $\forall f': \tau' \in \Gamma_1, \Gamma_2. \exists (v_a, v_b) \in \mathbf{vload}(M_a)(\Psi_a)(\gamma_a)(f') \times \mathbf{vload}(M_b)(\Psi_b)(\gamma_b)(f').$
 $\forall s' \sqsupseteq (s^\circ, s). \forall G \in \mathbf{GK}_{w \uparrow}. (s \in \text{sat}(G) \wedge \forall f: \tau \in \Gamma. (\gamma_a f, \gamma_b f) \in \overline{G}(s')(\tau)) \implies$
 $(v_a, v_b) \in \overline{G}(s')(\tau')$

Before we define s° and s , we observe that $s_1^\circ \in \Omega_{\mathcal{TS}}.S$ and $s_2^\circ \in \Omega_{\mathcal{TS}}.S$ necessarily look as follows (for some registries κ_1, κ_2):

$$s_1^\circ = ((R, \kappa_1), \emptyset, \gamma_b \uplus M_b \uplus \rho) \quad s_2^\circ = ((R, \kappa_2), \emptyset, \gamma_b \uplus M_b \uplus \rho)$$

The fact that in both states the $\mathsf{T}_{\text{ref}}^{\mathcal{TS}}$ component is empty (\emptyset) follows from condition (3) and (7), respectively. The knowledge about the unary components comes from (1) and (5), which we can instantiate using the least global knowledges $[w_1\uparrow]$ and $[w_2\uparrow]$, respectively.

We define $s^\circ \in \Omega_{\mathcal{TS}}.S$ as follows:

$$s^\circ := ((R, \kappa_1 \cup \kappa_2), \emptyset, \gamma_b \uplus M_b \uplus \rho)$$

Observe that we have $s^\circ \sqsupseteq_{\text{pub}} s_1^\circ$ and $s^\circ \sqsupseteq_{\text{pub}} s_2^\circ$. With the help of mainly (1), (5), and Lemma 58, it is not hard to show condition (9) for any s satisfying $s \sqsupseteq (s_1, s_2)$. For instance, since h_1^g conforms to κ_1 and h_2^g conforms to κ_2 , h^g conforms to $\kappa_1 \cup \kappa_2$.

So what about s ? The obvious choice of $s := (s_1, s_2)$ does not let us prove (10). To see this, suppose $G \in \mathbf{GK}_{w\uparrow}$. Since $c_a^l = c_a^{l,1} \cdot c_a^{l,2}$ (and similarly for c_b^l), the goal $(c_a^l, c_b^l) \in w.C(G)(s^\circ, s)$ would follow directly from

$$(c_a^{l,1}, c_b^{l,1}) \in w_1.C(G_{\langle 1 \rangle}^{s_2^\circ})(s^\circ, s_1) \quad \text{and} \quad (c_a^{l,2}, c_b^{l,2}) \in w_2.C(G_{\langle 2 \rangle}^{s_1^\circ})(s^\circ, s_2)$$

by construction of w as the product of w_1 and w_2 . However, (2) and (6) only tell us

$$(c_a^{l,1}, c_b^{l,1}) \in w_1.C(G_{\langle 1 \rangle}^{s_2^\circ})(s_1^\circ, s_1) \quad \text{and} \quad (c_a^{l,2}, c_b^{l,2}) \in w_2.C(G_{\langle 2 \rangle}^{s_1^\circ})(s_2^\circ, s_2).$$

While $s^\circ \sqsupseteq s_i^\circ$ holds, these configuration relations may of course not be monotone in their state argument.

Fortunately, the solution is easy: we apply stability of w_1 and w_2 to the previous statements. Doing so yields $s'_1 \sqsupseteq_{\text{pub}} s_1$ and $s'_2 \sqsupseteq_{\text{pub}} s_2$ such that

$$(c_a^{l,1}, c_b^{l,1}) \in w_1.C(G_{\langle 1 \rangle}^{s_2^\circ})(s^\circ, s'_1) \quad \text{and} \quad (c_a^{l,2}, c_b^{l,2}) \in w_2.C(G_{\langle 2 \rangle}^{s_1^\circ})(s^\circ, s'_2).$$

Now we can pick $s := (s'_1, s'_2)$ and (10) follows by monotonicity in the global knowledge.

Condition (11) obviously holds because we chose the reference bijection in s° to be empty. It thus remains to show (12). In order to do so, the following will be very helpful:

13. $\forall f': \tau' \in \Gamma_1. \mathbf{vload}(M_a^1)(\Psi_a^1)(\gamma_a^1)(f') \subseteq \mathbf{vload}(\text{link}(M_a^1, M_a^2))(\Psi_a)(\gamma_a)(f')$
14. $\forall f': \tau' \in \Gamma_2. \mathbf{vload}(M_a^2)(\Psi_a^2)(\gamma_a^2)(f') \subseteq \mathbf{vload}(\text{link}(M_a^1, M_a^2))(\Psi_a)(\gamma_a)(f')$
15. $\forall f': \tau' \in \Gamma_1. \mathbf{vload}(M_b^1)(\Psi_b^1)(\gamma_b^1)(f') \subseteq \mathbf{vload}(\text{link}(M_b^1, M_b^2))(\Psi_b)(\gamma_b)(f')$
16. $\forall f': \tau' \in \Gamma_2. \mathbf{vload}(M_b^2)(\Psi_b^2)(\gamma_b^2)(f') \subseteq \mathbf{vload}(\text{link}(M_b^1, M_b^2))(\Psi_b)(\gamma_b)(f')$

They roughly say that retrieving the value of exported function f' from the composed module is the same as retrieving it from the component module in which it is defined (recall that Γ_i describes the exports of M_a^i and M_b^i). Their proofs are straightforward, relying on the exports of the linked modules being disjoint.

So consider $f':\tau' \in \Gamma_1, \Gamma_2$. This means either $f':\tau' \in \Gamma_1$ or $f':\tau' \in \Gamma_2$. In the first case, we use (4) with (13) and (15) to obtain values v_a, v_b for which we know

$$\begin{aligned} & \forall s' \sqsupseteq (s_1^\circ, s_1). \forall G \in \mathbf{GK}_{w_1\uparrow}. \\ & (s_1 \in \text{sat}(G) \wedge \forall f:\tau \in \Gamma. (\gamma_a^1 f, \gamma_b^1 f) \in \overline{G}(s')(\tau)) \implies (v_a, v_b) \in \overline{G}(s')(\tau') \end{aligned}$$

and must show:

$$\begin{aligned} & \forall s' \sqsupseteq (s^\circ, s). \forall G \in \mathbf{GK}_{w\uparrow}. \\ & (s \in \text{sat}(G) \wedge \forall f:\tau \in \Gamma. (\gamma_a f, \gamma_b f) \in \overline{G}(s')(\tau)) \implies (v_a, v_b) \in \overline{G}(s')(\tau') \end{aligned}$$

We know that the given future state s' must have the form $(s^{\circ'}, (s_1'', s_2''))$ with $s^{\circ'} \sqsupseteq s^\circ \sqsupseteq s_1^\circ$ and $s_1'' \sqsupseteq s_1' \sqsupseteq s_1$. Therefore we can instantiate the premise above with $(s^{\circ'}, s_1'') \sqsupseteq (s_1^\circ, s_1)$. It is then easy to check the rest. In particular, if γ_a, γ_b is pointwise related for labels in Γ , then so is γ_a^1, γ_b^1 since it is the same pair by construction.

In the second case, where $f':\tau' \in \Gamma_2$, we use (8) with (14) and (16) to obtain v_a, v_b for which we know

$$\begin{aligned} & \forall s' \sqsupseteq (s_2^\circ, s_2). \forall G \in \mathbf{GK}_{w_2\uparrow}. \\ & (s_2 \in \text{sat}(G) \wedge \forall f:\tau \in \Gamma, \Gamma_1. (\gamma_a^2 f, \gamma_b^2 f) \in \overline{G}(s')(\tau)) \implies (v_a, v_b) \in \overline{G}(s')(\tau') \end{aligned}$$

and must show:

$$\begin{aligned} & \forall s' \sqsupseteq (s^\circ, s). \forall G \in \mathbf{GK}_{w\uparrow}. \\ & (s \in \text{sat}(G) \wedge \forall f:\tau \in \Gamma. (\gamma_a f, \gamma_b f) \in \overline{G}(s')(\tau)) \implies (v_a, v_b) \in \overline{G}(s')(\tau') \end{aligned}$$

The interesting part here is: how do we know that if γ_a, γ_b is pointwise related for labels in Γ , then so is γ_a^2, γ_b^2 for Γ, Γ_1 ? Recall that γ_a^2 and γ_b^2 extend γ_a and γ_b with the values exported by (the concrete loadings of) M_a^1 and M_b^2 , respectively. We have just shown, in the first case above, that these exported values are indeed related!

We conclude this section by stating the full modularity theorem.

Theorem 14 (Modularity of $\lesssim_{\mathcal{TS}}$).

$$\frac{\begin{array}{c} \Gamma \vdash M_b^1 : \Gamma_1 \\ \vdash M_a^1 : |\Gamma_1| \\ \Gamma \vdash M_a^1 \lesssim_{\mathcal{TS}} M_b^1 : \Gamma_1 \end{array} \quad \begin{array}{c} \Gamma, \Gamma_1 \vdash M_b^2 : \Gamma_2 \\ \vdash M_a^2 : |\Gamma_2| \\ \Gamma, \Gamma_1 \vdash M_a^2 \lesssim_{\mathcal{TS}} M_b^2 : \Gamma_2 \end{array}}{\Gamma \vdash \text{link}(M_a^1, M_a^2) \lesssim_{\mathcal{TS}} \text{link}(M_b^1, M_b^2) : \Gamma_1, \Gamma_2}$$

4.8 Proof of Transitivity

4.8.1 Overview

Let us begin with the end, the theorem we wish to prove.

Theorem 15 (Transitivity).

$$\frac{|\Gamma| \vdash M_{\mathcal{T}} \lesssim_{\mathcal{TI}} M_{\mathcal{I}} : |\Gamma'| \quad |\Gamma| \vdash M_{\mathcal{I}} \lesssim_{\mathcal{II}}^* M'_{\mathcal{I}} : |\Gamma'| \quad \Gamma \vdash M'_{\mathcal{I}} \lesssim_{\mathcal{IS}} M_{\mathcal{S}} : \Gamma'}{\Gamma \vdash M_{\mathcal{T}} \lesssim_{\mathcal{TS}} M_{\mathcal{S}} : \Gamma'}$$

We derive this as a corollary of two lemmas, one about composing $\lesssim_{\mathcal{TI}}$ with $\lesssim_{\mathcal{IS}}$ and one about composing $\lesssim_{\mathcal{II}}$ with $\lesssim_{\mathcal{IS}}$:

Lemma 60 (Transitivity, Part 1).

$$\frac{|\Gamma| \vdash M_{\mathcal{T}} \lesssim_{\mathcal{TI}} M_{\mathcal{I}} : |\Gamma'| \quad \Gamma \vdash M_{\mathcal{I}} \lesssim_{\mathcal{IS}} M_{\mathcal{S}} : \Gamma'}{\Gamma \vdash M_{\mathcal{T}} \lesssim_{\mathcal{TS}} M_{\mathcal{S}} : \Gamma'}$$

Lemma 61 (Transitivity, Part 2).

$$\frac{|\Gamma| \vdash M_{\mathcal{I}} \lesssim_{\mathcal{II}} M'_{\mathcal{I}} : |\Gamma'| \quad \Gamma \vdash M'_{\mathcal{I}} \lesssim_{\mathcal{IS}} M_{\mathcal{S}} : \Gamma'}{\Gamma \vdash M_{\mathcal{I}} \lesssim_{\mathcal{IS}} M_{\mathcal{S}} : \Gamma'}$$

Note that we do not need to show transitivity of $\lesssim_{\mathcal{II}}$ itself because, in order to obtain Theorem 15, we can simply iterate Lemma 61.

Thanks to our uniform setup, the proofs of Lemma 60 and Lemma 61 mirror each other. In fact, they are so similar that large parts of their formalization are identical. They also roughly follow the original transitivity proof of PBs (Section 3.9, but of course many details are necessarily different. One key simplification in PILS is that since our intermediate language is untyped, the complexity having to do with abstract types can be avoided. On the other hand, one key complication in PILS is due to **E**'s asymmetric small-step formulation of **E**.

Note that, as an alternative to Lemma 61, we could try to instead show the following:

$$\frac{|\Gamma| \vdash M_{\mathcal{T}} \lesssim_{\mathcal{TI}} M'_{\mathcal{I}} : |\Gamma'| \quad |\Gamma| \vdash M'_{\mathcal{I}} \lesssim_{\mathcal{II}} M_{\mathcal{I}} : |\Gamma'|}{|\Gamma| \vdash M_{\mathcal{T}} \lesssim_{\mathcal{TI}} M_{\mathcal{I}} : |\Gamma'|}$$

Together with Lemma 60, this would also yield Theorem 15. However, we choose to prove Lemma 61 because its statement is closer to that of Lemma 60: in each the second premise refers to an instance of the *typed* model. In the rest of this section we go into details of the proof of Lemma 60 only, since that of Lemma 61 is so similar.

Recall that the proof of transitivity for PBs consists of two parts. In the first, we construct a world W that relates the programs in question but has the wrong shape. In the second, we construct a world of the right shape and prove—with the help of a notion of isomorphism—that it relates the same programs as W . Here, for simplicity, we *directly* construct a world of the correct shape that relates the programs in question, *i.e.*, we skip the indirection through a theory of world morphisms.

We do not spell out all the details here but mainly present our definitions and lemma statements and sketch the proofs of some of the key lemmas.

4.8.2 Constructing the Local World

At the very outset, we are given a set of type names. We use this set, unchanged, to instantiate the second premise ($\Gamma \vdash M_{\mathcal{I}} \lesssim_{\mathcal{IS}} M_{\mathcal{S}} : \Gamma'$). The first premise doesn't require any such set since it is phrased in terms of the untyped model $\lesssim_{\mathcal{TI}}$. Hence the two premises give us local worlds $w_1 \in \text{LWorld}^{\Omega_{\mathcal{TI}}.\mathcal{T}}$ and $w_2 \in \text{LWorld}^{\Omega_{\mathcal{IS}}.\mathcal{T}}$, respectively. Based on these, we construct a local world $w \in \text{LWorld}^{\Omega_{\mathcal{TS}}.\mathcal{T}}$. The difficulty, of course, is to do so such that w satisfies all the required properties (*e.g.*, stability).

Like for PBs (cf. Lemma 33 in Section 3.9), one of these key properties concerns the decomposition of states and global knowledges: given $G \in \mathbf{GK}_{w\uparrow}$ and a state $s \in w\uparrow.\mathcal{S}$, we must be able to decompose G into a global knowledge for $w_1\uparrow$ (say G_1) and one for $w_2\uparrow$ (say G_2) as well as decompose s into s_1 and s_2 such that

$$\overline{G_1}(s_1) \circ \overline{G_2}(s_2)(\tau) = \overline{G}(s)(\tau).$$

We start by defining w 's transition system (the remaining components will be defined later on). Basically, we follow the same approach as in the proof for PBs, which was to define $w.\mathcal{S}$ as the product of $w_1\uparrow.\mathcal{S}$ and $w_2\uparrow.\mathcal{S}$. In PILS, however, the global part of a state contains some information that does not need to end up in w 's state. More concretely, recall that any states $s_1 \in w_1\uparrow.\mathcal{S}$ and $s_2 \in w_2\uparrow.\mathcal{S}$ look as follows:

$$s_1 = (s_1^\circ, s_{\text{lc}}^1) = ((s_1^{\mathcal{T}}, s_{\text{ref}}^1, s_1^{\mathcal{I}}), s_{\text{lc}}^1) \quad s_2 = (s_2^\circ, s_{\text{lc}}^2) = ((s_2^{\mathcal{I}}, s_{\text{ref}}^2, s_2^{\mathcal{S}}), s_{\text{lc}}^2),$$

Here,

- $s_1^{\mathcal{T}}$ is the unary component concerning \mathcal{T} (register file and registry);
- $s_1^{\mathcal{I}}$ and $s_2^{\mathcal{I}}$ are the unary components concerning \mathcal{I} (*i.e.*, carry no information)
- $s_2^{\mathcal{S}}$ is the unary component concerning \mathcal{S} (label environment);
- $s_{\text{ref}}^1 \in \mathcal{T}_{\text{ref}}^{\mathcal{TI}}$ is the (untyped) partial bijection on global references between \mathcal{T} and \mathcal{I} ;
- $s_{\text{ref}}^2 \in \mathcal{T}_{\text{ref}}^{\mathcal{IS}}$ is the (typed) partial bijection on global references between \mathcal{I} and \mathcal{S} ;
- s_{lc}^1 is the local state from the unknown $w_1.\mathcal{S}$;
- s_{lc}^2 is the local state from the unknown $w_2.\mathcal{S}$.

We define w 's states to be products of these states without the unary parts, because $\Omega_{\mathcal{TS}}.\mathcal{S}$ already provides those⁹:

$$w.\mathcal{T} := (\mathcal{T}_{\text{ref}}^{\mathcal{TI}} \times w_1.\mathcal{T}) \times (\mathcal{T}_{\text{ref}}^{\mathcal{IS}} \times w_2.\mathcal{T})$$

⁹It does, of course, not provide the unary part specific to the intermediate program, but that carries no information anyway. If it would be non-trivial, we would have to include it in $w.\mathcal{S}$.

Like for PBs, it is crucial that we embed $w_1\uparrow$ and $w_2\uparrow$'s global reference information in $w.S$.

Given $s_1 \in w_1\uparrow$ and $s_2 \in w_2\uparrow$, we compose them to a $w\uparrow$ state, written $s_1 \diamond s_2$, in the following way:

Definition 28.

$$\begin{aligned} (-) \diamond (-) &\in w_1\uparrow.S \rightarrow w_2\uparrow.S \rightarrow w\uparrow.S \\ ((s_1^T, s_{\text{ref}}^1, s_1^I), s_{\text{lc}}^1) \diamond ((s_2^T, s_{\text{ref}}^2, s_2^S), s_{\text{lc}}^2) &:= ((s_1^T, s_{\text{ref}}, s_2^S), ((s_{\text{ref}}^1, s_{\text{lc}}^1), (s_{\text{ref}}^2, s_{\text{lc}}^2))) \\ \text{where } s_{\text{ref}} &:= \{ (\tau, x_1, x_3) \mid \exists x_2. (x_1, x_2) \in s_{\text{ref}}^1 \wedge (\tau, x_2, x_3) \in s_{\text{ref}}^2 \} \end{aligned}$$

To decompose a $w\uparrow$ state, we ignore $\Omega_{\mathcal{T}\mathcal{S}}$'s information on global references:

Definition 29.

$$\begin{aligned} (-)_{(1)} &\in w\uparrow.S \rightarrow w_1\uparrow.S \\ ((s_1^{\circ}, s_{\text{ref}}, s_2^{\circ}), ((s_{\text{ref}}^1, s_1), (s_{\text{ref}}^2, s_2)))_{(1)} &:= ((s_1^{\circ}, s_{\text{ref}}^1, 1), s_1) \\ (-)_{(2)} &\in w\uparrow.S \rightarrow w_2\uparrow.S \\ ((s_1^{\circ}, s_{\text{ref}}, s_2^{\circ}), ((s_{\text{ref}}^1, s_1), (s_{\text{ref}}^2, s_2)))_{(2)} &:= ((1, s_{\text{ref}}^2, s_2^{\circ}), s_2) \end{aligned}$$

It is easy to see that $(s_1 \diamond s_2)_{(1)} = s_1$ and $(s_1 \diamond s_2)_{(2)} = s_2$ for any $s_1 \in w_1\uparrow$ and $s_2 \in w_2\uparrow$.

These definitions immediately yield the following nice properties, independent of how exactly the other components of w will be defined.

Lemma 62.

- $w\uparrow.\text{vqh}_a(s_1 \diamond s_2) = w_1\uparrow.\text{vqh}_a(s_1)$
- $w\uparrow.\text{vqh}_b(s_1 \diamond s_2) = w_2\uparrow.\text{vqh}_b(s_2)$
- $w\uparrow.\text{cqh}_a(s_1 \diamond s_2) = w_1\uparrow.\text{cqh}_a(s_1)$
- $w\uparrow.\text{cqh}_b(s_1 \diamond s_2) = w_2\uparrow.\text{cqh}_b(s_2)$
- $w\uparrow.\text{rgh}(s_1 \diamond s_2)(\tau) = w_1\uparrow.\text{rgh}(s_1) \circ w_2\uparrow.\text{rgh}(s_2)(\tau)$

With the help of these operations, we now define the decomposition of a global knowledge for the yet-to-be-defined w . The idea is the same as in the PB transitivity proof: for each tuple τ, v_1, v_3 in the global knowledge, we craft a value $\mathbf{B}(\tau, v_1, v_3)$ “in the middle” (here: an \mathcal{I} value) that serves three purposes:

1. It uniquely encodes the tuple (*i.e.*, \mathbf{B} is injective).
2. It violates the **GoodFuns** condition (if τ is a function type), so it cannot possibly be related by \mathbf{F} .

3. It rules out certain cases in the proof of a transitivity property for **E** (Lemma 66). For instance, an \mathcal{I} machine representing a call to such a function must not be able to produce an error.

B is defined using an arbitrary (but fixed) injective function $\mathbf{I} \in \text{CTyF} \times \mathcal{T}.\text{Val} \times \mathcal{S}.\text{Val} \rightarrow \mathcal{I}.\text{Val}$ that always returns a natural number value. We wrap **I**'s result depending on the type τ such that $\mathbf{B}(\tau, v_1, v_3)$ has the right shape according to the value query handler for \mathcal{I} from Figure 4.18. For instance, if $\tau = \tau_1 \rightarrow \tau_2$, then it is important that $\mathbf{B}(\tau, v_1, v_3)$ is a proper function value, *i.e.*, a closure with a fix term inside.

Definition 30.

$$\begin{aligned} \mathbf{B} &\in \text{CTyF} \times \mathcal{T}.\text{Val} \times \mathcal{S}.\text{Val} \rightarrow \mathcal{I}.\text{Val} \\ \mathbf{B}(\tau \rightarrow \tau', v_1, v_3) &:= \langle \epsilon; \text{fix } f(y, k). \text{halt } n \rangle \quad \text{where } n := \mathbf{I}(\tau \rightarrow \tau', v_1, v_3) \\ \mathbf{B}(\forall \alpha. \tau, v_1, v_3) &:= \langle \epsilon; \Lambda k. \text{halt } n \rangle \quad \text{where } n := \mathbf{I}(\forall \alpha. \tau, v_1, v_3) \\ \mathbf{B}(\mathbf{n}, v_1, v_3) &:= \langle \epsilon; n \rangle \quad \text{where } n := \mathbf{I}(\mathbf{n}, v_1, v_3) \end{aligned}$$

Using this, we first construct $G_{\{1\}}^{s_2}$ and $G_{\{2\}}^{s_1}$, and then close them over **F** and $w_1.\text{NR}$ (respectively $w_2.\text{NR}$) to ensure that they can be shown to be proper global knowledges later.

Definition 31. For $G \in w \uparrow. \mathcal{S} \rightarrow \text{VRelF}_{\mathcal{T}, \mathcal{S}}$ and $s_2 \in w_2 \uparrow. \mathcal{S}$, we define $G_{(1)}^{s_2} \in w_1 \uparrow. \mathcal{S} \rightarrow \text{VRelF}_{\mathcal{T}, \mathcal{I}}$ (untyped):

$$\begin{aligned} G_{\{1\}}^{s_2}(s_1) &:= \{ (v_1, \mathbf{B}(\tau, v_1, v_3)) \mid \tau \notin w_2.\text{NS} \wedge \\ &\quad \exists s'_1 \sqsubseteq s_1. \exists s'_2 \sqsubseteq s_2. (v_1, v_3) \in G(s'_1 \diamond s'_2)(\tau) \} \\ G_{(1)}^{s_2} &:= [w_1 \uparrow]_{G_{\{1\}}^{s_2}} \end{aligned}$$

Definition 32. For $G \in w \uparrow. \mathcal{S} \rightarrow \text{VRelF}_{\mathcal{T}, \mathcal{S}}$ and $s_1 \in w_1 \uparrow. \mathcal{S}$, we define $G_{(2)}^{s_1} \in w_2 \uparrow. \mathcal{S} \rightarrow \text{VRelF}_{\mathcal{I}, \mathcal{S}}$ (typed).

$$\begin{aligned} G_{\{2\}}^{s_1}(s_2)(\tau) &:= \{ (\mathbf{B}(\tau, v_1, v_3), v_3) \mid \tau \notin w_2.\text{NS} \wedge \\ &\quad \exists s'_1 \sqsubseteq s_1. \exists s'_2 \sqsubseteq s_2. (v_1, v_3) \in G(s'_1 \diamond s'_2)(\tau) \} \\ G_{(2)}^{s_1} &:= [w_2 \uparrow]_{G_{\{2\}}^{s_1}} \end{aligned}$$

We can now fully define w . As its order we choose the lexicographic product of w_1 and w_2 's order. (This will be justified in the proof of Lemma 66.)

Definition 33.

$$\begin{aligned} s_{\text{ref}}^1 \setminus_{[1]} s_{\text{ref}} &:= \{ (v_1, v_2) \in s_{\text{ref}}^1 \mid \nexists \tau, v_3. (\tau, v_1, v_3) \in s_{\text{ref}} \} \\ s_{\text{ref}}^2 \setminus_{[2]} s_{\text{ref}} &:= \{ (\tau, v_2, v_3) \in s_{\text{ref}}^2 \mid \nexists v_1. (\tau, v_1, v_3) \in s_{\text{ref}} \} \end{aligned}$$

Definition 34.

$$\begin{aligned}
w.\top &:= (\top_{\text{ref}}^{\mathcal{TI}} \times w_1.\top) \times (\top_{\text{ref}}^{\mathcal{IS}} \times w_2.\top) \\
w.\mathbf{O} &:= w_2.\mathbf{O} \# w_1.\mathbf{O} \\
w.\mathbf{C}(G)(s) &:= \{ (c_{\text{ref}}^{\mathcal{T}} \cdot c^{\mathcal{T}}, d_{\text{ref}}^{\mathcal{S}} \cdot d^{\mathcal{S}}) \mid \\
&\quad \exists s_1, s_2, s_{\text{ref}}. s = s_1 \diamond s_2 \wedge s = ((-, s_{\text{ref}}, -), -) \wedge \\
&\quad \exists s_1^{\mathcal{T}}, s_{\text{ref}}^1, s_1^{\mathcal{I}}, s_{\text{lc}}^1. s_1 = ((s_1^{\mathcal{T}}, s_{\text{ref}}^1, s_1^{\mathcal{I}}), s_{\text{lc}}^1) \wedge \\
&\quad \exists s_2^{\mathcal{T}}, s_{\text{ref}}^2, s_2^{\mathcal{I}}, s_{\text{lc}}^2. s_2 = ((s_2^{\mathcal{T}}, s_{\text{ref}}^2, s_2^{\mathcal{I}}), s_{\text{lc}}^2) \wedge \\
&\quad \exists (c_{\text{ref}}^{\mathcal{T}}, c_{\text{ref}}^{\mathcal{I}}) \in C_{\text{ref}}^{\mathcal{TI}}(s_{\text{ref}}^1 \setminus_{[1]} s_{\text{ref}})(\overline{G_{(1)}^{s_2}}(s_1)). \\
&\quad \exists (d_{\text{ref}}^{\mathcal{I}}, d_{\text{ref}}^{\mathcal{S}}) \in C_{\text{ref}}^{\mathcal{IS}}(s_{\text{ref}}^2 \setminus_{[2]} s_{\text{ref}})(\overline{G_{(2)}^{s_1}}(s_2)). \\
&\quad \exists (c^{\mathcal{T}}, c^{\mathcal{I}}) \in w_1.\mathbf{C}(\overline{G_{(1)}^{s_2}})(s_1). \\
&\quad \exists (d^{\mathcal{I}}, d^{\mathcal{S}}) \in w_2.\mathbf{C}(\overline{G_{(2)}^{s_1}})(s_2). \\
&\quad c_{\text{ref}}^{\mathcal{I}} \cdot c^{\mathcal{I}} = d_{\text{ref}}^{\mathcal{I}} \cdot d^{\mathcal{I}} \wedge \\
&\quad \forall h \neq \perp. \text{dom}(h) \subseteq \{l_2 \mid \exists l_1. (l_1, l_2) \in s_{\text{ref}}^1\} \implies \\
&\quad \forall c \in \mathbf{core}. \exists m \in \mathbf{real}(c \cdot (h, \emptyset) \cdot c^{\mathcal{I}}) \}
\end{aligned}$$

$$\begin{aligned}
w.\mathbf{NS} &:= w_2.\mathbf{NS} \\
w.\mathbf{NR}(G)(s)(\mathbf{n}) &:= \{ (v_1, v_3) \in \overline{G_{(1)}^{s_2}}(s_1) \circ \overline{G_{(2)}^{s_1}}(s_2)(\mathbf{n}) \mid \mathbf{n} \in w.\mathbf{NS} \}
\end{aligned}$$

The construction of its configuration relation $w.\mathbf{C}$ closely follows that of the heap relation in the PB proof. The last condition basically asserts that $c^{\mathcal{I}}$ is disjoint from the intermediate locations in s_{ref}^1 and from any core.

Lemma 63. If $G \in \mathbf{GK}_{w\uparrow}$, then

- $\forall s_2. G_{(1)}^{s_2} \in \mathbf{GK}_{w_1\uparrow}$, and
- $\forall s_1. G_{(2)}^{s_1} \in \mathbf{GK}_{w_2\uparrow}$.

Lemma 64. If $G \in \mathbf{GK}_{w\uparrow}$, then

$$\overline{G_{(1)}^{s_2}}(s_1) \circ \overline{G_{(2)}^{s_1}}(s_2)(\tau) = \overline{G}(s_1 \diamond s_2)(\tau).$$

Proof. The \supseteq direction is fairly straightforward by induction on the value closure (which is a least fixed point). One interesting point has to do with type names \mathbf{n} that lie outside the control of w , i.e., $\mathbf{n} \notin w.\mathbf{NS}$. If values v_1, v_3 are related at such a type, then we establish

$$(v_1, v_3) \in \overline{G_{(1)}^{s_2}}(s_1) \circ \overline{G_{(2)}^{s_1}}(s_2)(\tau)$$

by showing

$$(v_1, \mathbf{B}(\mathbf{n}, v_1, v_3)) \in \overline{G_{\{1\}}^{s_2}}(s_1) \quad \wedge \quad (\mathbf{B}(\mathbf{n}, v_1, v_3), v_3) \in \overline{G_{\{2\}}^{s_1}}(s_2)(\mathbf{n}).$$

In order for this to work, it is crucial that the *untyped* value closure operator also includes a “type name” case (cf. Figure 4.13 in Section 4.3), and that this case does *not* restrict the value on the target side (here v_1) as we know nothing about it.

The \subseteq direction, like for PBs, is proven by induction on the left value closure. Formally, the statement that we show by induction is

$$\overline{G_{(1)}^{s_2}} \subseteq X,$$

where $X(s_1) := \{ (v_1, v_2) \mid \forall \tau, v_3. (v_2, v_3) \in \overline{G_{(2)}^{s_1}}(s_2)(\tau) \implies \overline{G}(s_1 \diamond s_2)(\tau) \}$.

There is an interesting PILS twist here, too. Consider the case where inverting the value closure’s generating function yields the “unit” case. Then we know $v_1 \in \Omega_{\mathcal{I}\mathcal{L}}.\text{vqh}_a(-)(\text{unit})$ and $v_2 \in \Omega_{\mathcal{I}\mathcal{L}}.\text{vqh}_b(-)(\text{unit})$, *i.e.*, v_1 is an arbitrary machine word and v_2 is $\langle \rangle$. However, since the $\mathcal{I}\mathcal{L}$ model is untyped, we know nothing about τ yet. So when we invert $\overline{G_{(2)}^{s_1}}(s_2)(\tau)$, there are many possible cases to deal with—and several of them cannot be ruled out. For instance, $(v_2, v_3) \in \overline{G_{(2)}^{s_1}}(s_2)(\tau)$ may hold due to the “roll” case where $v_2 \in \Omega_{\mathcal{I}\mathcal{L}}.\text{vqh}_b(-)(\text{roll } v'_2)$. Note that this is not in contradiction to $v_2 = \langle \rangle$ because \mathcal{I} does not have an explicit *roll* construct. Hence all we know is $v_2 = v'_2$ and $(v'_2, v'_3) \in \overline{G_{(2)}^{s_1}}(s_2)(\tau)$, so we are going in circles. For this reason, unlike in the PB proof, we actually have to do a second induction, this time on $\overline{G_{(2)}^{s_1}}$.

Unrelatedly, and similar to PBs, the proof of this lemma is the only place where we exploit the **GoodFuns** restriction on **F** to infer that

$$\mathbf{F}_{w_1\uparrow}(G_{(1)}^{s_2})(s_1) \circ G_{\{2\}}^{s_1}(s_2)(\tau)$$

and

$$G_{\{1\}}^{s_2}(s_1) \circ \mathbf{F}_{w_2\uparrow}(G_{(2)}^{s_1})(s_2)(\tau)$$

are empty (for function types). □

The next lemma talks about the configuration relations. We did not have such a lemma explicitly in the PB proof, but it was hidden in establishing the world isomorphism. Here we have an additional condition, though, which states that the \mathcal{I} configuration c_2 can be realized whenever composed with a core. This should be thought of as a sanity check.

Lemma 65. If $G \in \mathbf{GK}_{w\uparrow}$, then:

$$\begin{aligned} (c_1, c_3) \in w\uparrow.\mathbf{C}(G)(s_1 \diamond s_2) &\iff \exists c_2. \\ &(\forall e \in \mathbf{core}. \exists m \in \mathbf{real}(e \cdot c_2)) \wedge \\ &(c_1, c_2) \in w_1\uparrow.\mathbf{C}(G_{(1)}^{s_2})(s_1) \wedge (c_2, c_3) \in w_2\uparrow.\mathbf{C}(G_{(2)}^{s_1})(s_2) \end{aligned}$$

Of course, this property does not hold for arbitrary $w\uparrow$ states but only for those of the form $s_1 \diamond s_2$.

One of the most important and also most involved parts is the proof of transitivity of **E**. We first prove the following lemma, concerning the case where the ϕ argument is non-trivial. The other case, Lemma 67, follows from that.

		err	ret	step		
				call	rec	
					log	phys
err	0	ζ_1	ζ_1	ζ_1	ζ_1	ζ_1
	+	err	ζ_2	ζ_3	(ind)	(ind)
ret	0	err	ret*	ζ_4	err	err
	+	err	ζ_5	ζ_6	(ind)	(ind)
step	0	call	err	ζ	call*	err
		rec	err	log*	log*	log*
			err	ζ	ζ	log*
	+	phys	err	ζ	ζ	phys*
			err	ζ	(ind)	(ind)

Figure 4.23: Overview of the proof of transitivity of \mathbf{E} (Lemma 66)

Definition 35.

$$\mathbf{B}_k := \langle \epsilon; \text{cont } y. \text{halt } 0 \rangle$$

Lemma 66. If

- $(e_1, e_2) \in \mathbf{E}_{w_1 \uparrow}(i)(\mathbf{k}_1, \mathbf{B}_k)(G_{(1)}^{s'_2})(s_1)(s'_1)(\phi)$
- $(e_2, e_3) \in \mathbf{E}_{w_2 \uparrow}(j)(\mathbf{B}_k, \mathbf{k}_3)(G_{(2)}^{s'_1})(s_2)(s'_2)(\phi)(\tau)$
- $\text{real}(e_2) \neq \emptyset$
- $s'_1 \sqsupseteq s_1$
- $s'_2 \sqsupseteq s_2$
- $\phi \neq \perp$

then $(e_1, e_3) \in \mathbf{E}_{w \uparrow}(\langle j, i \rangle)(\mathbf{k}_1, \mathbf{k}_3)(G)(s_1 \diamond s_2)(s'_1 \diamond s'_2)(\phi)(\tau)$.

The lemma is stated in terms of a very particular continuation for the intermediate program, namely \mathbf{B}_k . It is defined above as a continuation that is immediately halted (after one step). This is critical in ruling out certain cases in the proof.

Note that the Lemma's third condition is harmless, because, due to $\phi \neq \perp$, we are only interested in fully configured e_2 's, which clearly are realizable. Without this condition, however, the first two assumptions might hold trivially, while the goal might not hold trivially (e_2 might not be realizable, while e_1 and e_3 might).

Proof. By coinduction. Suppose $G \in \mathbf{GK}_{w\uparrow}$ and

$$(m_1, m_3) \in \mathbf{cfg}(w\uparrow.\mathbf{C})(G)(s'_1 \diamond s'_2)(\phi)(e_1, e_3)(\emptyset, \emptyset)(\eta_1, \eta_3).$$

We can decompose this using Lemma 65 (with an empty frame in the middle) such that for any m_2 and e_2 with $m_2 \in \mathbf{real}(e_2)$ we have:

$$(m_1, m_2) \in \mathbf{cfg}(w_1\uparrow.\mathbf{C})(G_{(1)}^{s'_2})(s'_1)(\phi)(e_1, e_2)(\emptyset, \emptyset)(\eta_1, \emptyset) \quad (4.1)$$

$$(m_2, m_3) \in \mathbf{cfg}(w_2\uparrow.\mathbf{C})(G_{(2)}^{s'_1})(s'_2)(\phi)(e_2, e_3)(\emptyset, \emptyset)(\emptyset, \eta_3) \quad (4.2)$$

Note that by construction we have $G_{(1)}^{s'_2} = \mathbf{F}_{w_1\uparrow}(G_{(1)}^{s'_2}) \cup w_2\uparrow.\mathbf{NR}(G_{(1)}^{s'_2}) \cup G_{\{1\}}^{s'_2}$ (and similarly for $G_{(2)}^{s'_1}$). Hence we can apply the external call lemma (Lemma 55) to the two main premises to obtain the following:

$$(e_1, e_2) \in \mathbf{E}_{w_1\uparrow}^{G_{\{1\}}^{s'_2}}(i)(\mathbf{k}_1, \mathbf{B}_\mathbf{k})(G_{(1)}^{s'_2})(s_1)(s'_1)(\phi) \quad (4.3)$$

$$(e_2, e_3) \in \mathbf{E}_{w_2\uparrow}^{G_{\{2\}}^{s'_1}}(j)(\mathbf{B}_\mathbf{k}, \mathbf{k}_3)(G_{(2)}^{s'_1})(s_2)(s'_2)(\phi)(\tau) \quad (4.4)$$

We now instantiate 4.3 with 4.1 and do a case analysis on the result. The possible cases make up the vertical axis of Figure 4.23. There are three main cases, but in each we make further distinctions. Moreover, we also (instantiate and) analyze 4.4 each time, which yields the additional cases shown on the horizontal axis.

In each of the three main cases of 4.3 (**ERR**, **RET**, **STEP** on the vertical axis), we are given a silent reduction sequence $m_2 \hookrightarrow^* m'_2$ and do an induction on it. “0” in the figure refers to the base case, where the sequence is empty and thus $m_2 = m'_2$. “+” refers to the case where the sequence is non-empty and where the inductive hypothesis plays a role (“ind”). Stars (*e.g.*, **RET**^{*}) indicate that the argument in the respective case makes use of the coinductive hypothesis. \nmid indicates impossible situations, where we derive a contradiction.

Here we sketch mainly the (vertical) **ERR** and **RET** cases. We use the **STEP** case to explain the choice of the well-founded order $w.\mathbf{O}$.

Case ERR. The intermediate program is erroneous: $m_2 \hookrightarrow^* m'_2$ with $m'_2 \in \mathbf{error}$.

We show the goal by also using the **ERR** clause, *i.e.*, we show $m_3 \hookrightarrow^* m'_3$ for some $m'_3 \in \mathbf{error}$.

We do so by induction on the reduction sequence $m_2 \hookrightarrow^* m'_2$. Note that $m_2 = m'_2$ is ruled out by 4.2, because an erroneous machine is not realizable (\nmid_1). So there is at least one step: $m_2 \hookrightarrow m''_2 \hookrightarrow^* m'_2$.

Let us now instantiate 4.4 with 4.2 and do another case analysis on the result.

Subcase ERR. The source machine m_3 is also erroneous, so we are done by appeal to **ERR**.

Subcase RET. The intermediate machine m_2 represents a return to “bad” continuation \mathbf{B}_k . Therefore, after one step, it reaches a machine that is halted, and moreover, not an error. This contradicts that m_2 eventually produces an error (\nmid_2).

Subcase STEP. We apply $m_2 \hookrightarrow m_2''$ and must deal with two further cases.

In **REC**, m_3 steps to or equals a machine that, by the inductive hypothesis, eventually produces an error. The goal follows immediately.

In **CALL**, m_2'' realizes a function call and thus $m_2'' \notin \mathbf{error}$. Moreover, by construction of $G_{\{2\}}^{s'_1}$, we know that the called function is “bad”, so m_2'' is actually halted and thus $m_2'' = m_2'$. This contradicts $m_2' \in \mathbf{error}$ (\nmid_3).

Case RET. The programs return to their continuations, with related values. In particular, $m_2 \hookrightarrow^* m_2'$ where m_2' returns some value v_2 to the “bad” continuation \mathbf{B}_k . We show the goal by induction on this reduction sequence, eventually using either the **ERR** or **RET** clause.

Suppose $m_2 = m_2'$. Unfortunately, this does not contradict anything. So we instantiate 4.4 with 4.2 and do another case analysis on the result.

Subcase ERR. As before.

Subcase RET. Using Lemmas 64 and 65, it is easy to compose all the information we have at this point and establish the goal using **RET**.

Subcase STEP. Since m_2 is a return to \mathbf{B}_k , it takes one step to a particular halted machine, say $m_2 \hookrightarrow m_2^\circ$. For this step, the subcase yields two further cases.

In **REC**, m_2° is again related to some source program by $\mathbf{E}_{w_2\uparrow}^{G_{\{2\}}^{s'_1}}$. It is easy to show that this can only hold due to the **ERR** case: $m_2^\circ \in \mathbf{halted}$ directly contradicts the **STEP** case; it also contradicts the **RET** case, because an \mathcal{I} machine realizing a return is never halted.

In **CALL**, m_2° realizes a function call. It is easy to see that this contradicts m_2° being halted due to \mathbf{B}_k (\nmid_4).

This was the base case of the induction. Now suppose there is at least one step: $m_2 \hookrightarrow m_2'' \hookrightarrow^* m_2'$. Here too, we instantiate 4.4 with 4.2 and do another case analysis on the result.

Subcase ERR. As before.

Subcase RET. Then m_2 is a return to \mathbf{B}_k . Its immediate successor m_2'' is therefore halted by construction of \mathbf{B}_k . Hence m_2'' must be m_2' , which is also a return. This is a contradiction: an \mathcal{I} machine cannot realize a return and be halted at the same time (\nmid_5).

Subcase STEP. We apply $m_2 \hookrightarrow m_2''$ and must deal with two further cases.

In **REC**, we can easily conclude with the help of the inductive hypothesis.

In **CALL**, m_2'' realizes a call to a “bad” function and thus is halted. This is in contradiction to m_2' realizing a return, as in \mathcal{I} a return is never halted ($\not\downarrow_6$).

Case STEP. Along the same lines as the previous two, but more tedious. As can be seen in Figure 4.23, we also need to further distinguish between stuttering and non-stuttering. Instead of going into the details, let us use this case to illustrate our choice of $w.O$ as the lexicographic product of $w_2.O$ and $w_1.O$. While it is natural that we choose a product of the two, the two obvious questions are: why lexicographic, and why $w_2.O$ before $w_1.O$?

To answer this, let us consider the **REC** case and assume that there is no stuttering, *i.e.*, the recursive occurrence of **E** uses an *unknown* $i' \in w_1.O$. Further assume that 4.4 yields the **REC** case as well, but that there stuttering happens. Hence the recursive occurrence of **E** necessarily uses a $j' \in w_2.O$ that is *smaller than* the initial j (from the lemma statement). Obviously we want to show the goal by appealing to the **REC** ourselves. And due to the stuttering in 4.4, we are forced to stutter as well. Fortunately, we can do so just fine by picking the budget $\langle j', i' \rangle$ since $j' < j$ implies $\langle j', i' \rangle < \langle j, i \rangle$. Note that no other product would work because i' may be greater than i .

If the intermediate program *does* stutter, witnessed by some $i' < i$, there is no reason to look into 4.4 at all. We can simply show the goal by stuttering immediately, leaving the source program untouched. This means choosing the new budget to be $\langle j, i' \rangle$, which is of course smaller than $\langle j, i \rangle$ due to $i' < i$ in that case.

□

Lemma 67. If

- $(e_1, e_2) \in \mathbf{E}_{w_1 \uparrow}(i)(\mathbf{k}_1, \mathbf{B}_k)(G_{(1)}^{s_2'})(s_1)(s_1')(\perp)$
- $(e_2, e_3) \in \mathbf{E}_{w_2 \uparrow}(j)(\mathbf{B}_k, \mathbf{k}_3)(G_{(2)}^{s_1'})(s_2)(s_2')(\perp)(\tau)$
- $e_2 \in \mathbf{core}$
- $s_1' \sqsupseteq s_1$
- $s_2' \sqsupseteq s_2$

then $(e_1, e_3) \in \mathbf{E}_{w \uparrow}(\langle j, i \rangle)(\mathbf{k}_1, \mathbf{k}_3)(G)(s_1 \diamond s_2)(s_1' \diamond s_2')(\perp)(\tau)$.

Note that the third condition ($e_2 \in \mathbf{core}$) is harmless, as we are only interested in such e_2 's. Without this condition, however, the first two assumptions might hold trivially, while the goal might not hold trivially (e_2 might not be a core while e_1 and e_3 might).

Proof. Straightforward, using Lemmas 65 and 66.

□

The key thing left to do, is proving stability of w . This is indeed very tedious but closely follows its PB counterpart (there are no new ideas needed). With that, we can easily establish the final property.

Theorem 16.

$$\frac{|\Gamma| \vdash M_1 \lesssim_{\mathcal{TI}} M_2 : |\Gamma'| \quad \Gamma \vdash M_2 \lesssim_{\mathcal{IS}} M_3 : \Gamma'}{\Gamma \vdash M_1 \lesssim_{\mathcal{TI}} M_3 : \Gamma'}$$

4.8.3 Discussion

As for PBs, one of the main complexities in the PILS transitivity proof lies in dealing with an ambiguity regarding reference allocation: while in one of the two given proofs, an allocation of the middle program may be treated as public (extending the global state), the same allocation may be treated as private (extending the local state) in the other proof. This is a result of transitivity being proven for completely arbitrary local worlds! One might wonder if we could not simplify matters significantly by resorting to an instrumentation of the IL that makes the choice of public vs. private allocation explicit in the program code. It seems to us that this approach only makes sense if one is willing to a priori decide on all subsequent optimizations. The issue is that a later added optimization might, for instance, figure out that some reference is never used and therefore can be removed. Such can only be proven correct if the reference was allocated as private, which it may not have been. For the sake of modularity, we therefore believe it is better to bite the bullet and deal with the ambiguity issue semantically, *e.g.*, in the way we did.

Chapter 5

The Pilsner Compiler and Its Verification

Contents

5.1	Overview	178
5.2	From S to \mathcal{I}: CPS Transformation	180
5.2.1	Definition	180
5.2.2	Verification	185
5.3	Infrastructure for Optimizations	187
5.3.1	Relating Open Expressions	188
5.3.2	Annotating Expressions With Transformations	193
5.4	The Commute Pass	197
5.4.1	Transformation	197
5.4.2	Verification	198
5.4.3	Alternative Implementation	199
5.5	The Dedup Pass	199
5.5.1	Transformation	199
5.5.2	Verification	199
5.6	The Hoist Pass	200
5.6.1	Transformation	200
5.6.2	Verification	201
5.7	The Dead Code Elimination Pass	202
5.7.1	Transformation	202
5.7.2	Verification	202
5.8	The Inline Pass	206
5.8.1	Transformation	206
5.8.2	Verification	206
5.8.3	Freshening	210

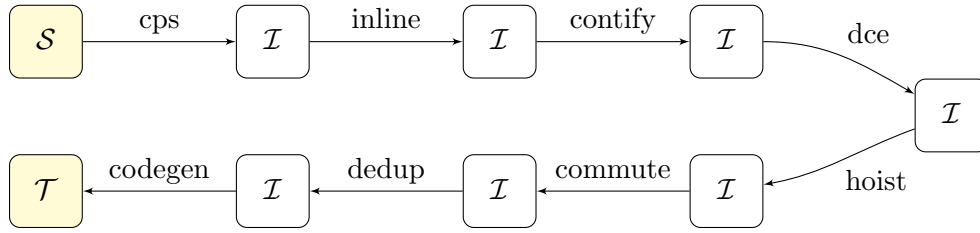


Figure 5.1: Structure of the Pilsner compiler

5.9 The Contify Pass	210
5.9.1 Transformation	210
5.9.2 Verification	213
5.10 The Codegen Pass	213
5.10.1 Transformation	213
5.10.2 Verification	221
5.11 The Full Pilsner Compiler	229
5.12 The Zwikel Compiler	230
5.12.1 Transformation	231
5.12.2 Verification	232
5.13 The Self-Modifying Awkward Example	236
5.14 Mechanization and Extraction	238
5.15 Putting It All Together	239

5.1 Overview

Using PILS, we have proven *in Coq* the correctness of two compilers from \mathcal{S} to \mathcal{T} : Pilsner and Zwikel. Pilsner’s structure is depicted in Figure 5.1. It uses \mathcal{I} as its CPS-based intermediate language and performs several code transformations. Zwikel, on the other hand, is more simplistic: it directly translates \mathcal{S} code into \mathcal{T} code in a straightforward way, similar to Hur and Dreyer’s one-pass compiler [33]. In particular, Zwikel neither uses an intermediate language nor performs any CPS transformation. In this chapter, we focus mostly on Pilsner, because it is by far the more interesting compiler.

Given a source module, Pilsner first translates it to \mathcal{I} via a CPS transformation. It also takes care to alpha-rename all bound variables such that in the resulting \mathcal{I} module, every variable is bound at most once. This *uniqueness condition* simplifies the implementation of most of the subsequent transformation passes, as one does not

have to worry about accidental variable capturing when rearranging code. Another nice characteristic of the produced intermediate code (not of \mathcal{I} per se) is that continuations are used in an affine fashion [39], *i.e.*, called at most once. This property is preserved by all other transformations at the intermediate level and enables a more efficient treatment of continuation variables compared to ordinary variables in the code generation pass.

At the intermediate level, Pilsner performs six “optimizations”. It first inlines selected top-level functions. For instance, if a module defines $F = \text{fix } f(y, k). e$, then a call to F inside a subsequently defined function will be rewritten as follows:

$$F \ x \ k' \rightsquigarrow e[F/f][x/y][k'/k]$$

(If the function is recursive, *i.e.*, if f is used in e , this is essentially just an unrolling.) Since inlining destroys the uniqueness property of bound variables, we immediately follow it with a “freshening” pass that re-establishes uniqueness.

Next comes contification, which converts certain functions to continuations. Subsequently, Pilsner performs a simple dead code (and variable) elimination, rewriting $\text{let } x = a \text{ in } e$ to e whenever x does not occur in e . This is justified because, in our IL, evaluation of a pure expression a does not have any observable side effects. In the same manner, it also eliminates unused `read` operations and unused allocations (but not `write` operations because that would be unsound). Following DCE, it hoists let-bindings out of function and continuation definitions, subject to some syntactic constraints. For example:

$$\begin{aligned} & \text{let } f = (\text{fix } f(y, k). \text{let } z = x.1 \text{ in } e) \text{ in } e' \\ \rightsquigarrow & \text{let } z = x.1 \text{ in let } f = (\text{fix } f(y, k). e) \text{ in } e' \end{aligned}$$

if x is none of f, y, k . This avoids recomputation of the projection each time f is called.

Next comes a pass that commutes let-bindings (where possible) in order to group together bindings that assign names to the same expression. For instance:

$$\begin{aligned} & \text{let } x = a \text{ in let } y = b \text{ in let } z = a \text{ in } e \\ \rightsquigarrow & \text{let } x = a \text{ in let } z = a \text{ in let } y = b \text{ in } e \end{aligned}$$

The last IL transformation, deduplication, gets rid of such consecutive duplicate bindings by rewriting the above expression as follows:

$$\rightsquigarrow \text{let } x = a \text{ in let } y = b \text{ in } e[x/z]$$

This can be seen as a common subexpression elimination.

Code generation, the final pass in the chain, translates to the machine language \mathcal{T} . Recall that there are three kinds of “variables” in \mathcal{I} : term variables x , continuation variables k , and labels F . Labels are translated to absolute addresses according to the import table. Term variable accesses are translated to lookups (based on position) in a linked list on the heap, pointed to by the `env` register. Functions are converted to

closures, *i.e.*, pairs of environment and code pointer (module-level functions simply have an empty environment), which live on the heap. A closure's environment is loaded into the `env` register when the function is called. Finally, continuations are allocated on the stack. Accordingly, continuation variable accesses are translated to lookups (based on position) on the stack, with the side effect that the continuation in question, as well as all more-recently defined ones (above it on the stack), are popped. This is safe because the affinity property mentioned earlier ensures that they won't be needed anymore.

5.2 From \mathcal{S} to \mathcal{I} : CPS Transformation

5.2.1 Definition

The CPS transformation translates from \mathcal{S} to \mathcal{I} . Figure 5.2 show its definition. The module-level translation *cps* is shown right at the top. It is defined in terms of an auxiliary function *mcps*, which recursively uses *fcps* to translate top-level functions. *fcps* itself relies on *ecps* to translate expressions, whose definition is continued in Figures 5.3 and 5.4.

Besides performing a fairly standard CPS conversion, our transformation also makes sure that in the resulting expression no variable is bound twice by always choosing a fresh one. To do so, it uses injections *tvar* and *kvar* that map numbers to variables. The various translation functions take as input a number *n* and assume that any variable corresponding to a number greater or equal *n* is unused. Whenever they introduce a bound variable, they will pick one of these unused variables. To maintain the set of unused ones, the translations return an updated number along with the converted program.

When the expression translation *ecps* introduces a bound variable *y* in correspondence to a bound variable *x* in the source program, this correspondence is remembered in the ϕ argument such that, in the recursion on subexpressions, the translation knows what to do when it hits upon a free occurrence of *x*.

The last argument of *ecps* is a meta-level continuation κ . It takes a number (the next free variable) and a term variable or label (the input to the continuation) and returns an \mathcal{I} expression and an updated variable number. We use meta-level functions rather than object-level continuations in order to avoid creating too many redexes (this could be optimized further, tough) [39].

Let us look at a few cases.

- The translation of a label *F* simply passes the label to the meta-level continuation κ .
- The translation of a source variable *x* does the same—except that it first renames the variable according to ϕ .
- The translation of a number *m* essentially just passes the number to κ . To do so, it introduces a let-binding, for which it uses the fresh term variable

$$\begin{aligned}
& \text{cps} \in \mathcal{S}.\mathbf{Mod} \rightarrow \mathcal{I}.\mathbf{Mod} \\
& \text{cps}(M) = \text{mcps}(0)(M) \\
\\
& \text{mcps} \in \mathbb{N} \rightarrow \mathcal{S}.\mathbf{Mod} \rightarrow \mathcal{I}.\mathbf{Mod} \\
& \text{mcps}(n)(\epsilon) = \epsilon \\
& \text{mcps}(n)((F, \mathbf{v}), M) = \text{let } (n', v') := \text{fcps}(n)(\mathbf{v}) \text{ in } (F, v'), \text{mcps}(n')(M) \\
\\
& \text{fcps} \in \mathbb{N} \rightarrow \mathcal{S}.\mathbf{Val} \rightarrow \mathbf{Exp}_{\mathcal{I}} \\
& \text{fcps}(n)(\text{fix } f(y). e) = \text{let } f' := \text{tvar}(n) \text{ in} \\
& \quad \text{let } y' := \text{tvar}(n+1) \text{ in} \\
& \quad \text{let } k := \text{kvar}(n+2) \text{ in} \\
& \quad \text{let } \phi' := \text{id}[f'/f][y'/y] \text{ in} \\
& \quad \text{let } (n_1, e') := \text{ecps}(e)(n+3)(\phi')(\lambda n_2, x. (n_2, k \ x)) \text{ in} \\
& \quad (n_1, \text{fix } f'(y', k). e') \\
& \text{fcps}(n)(\Lambda. e) = \text{let } k := \text{kvar}(n) \text{ in} \\
& \quad \text{let } (n_1, e') := \text{ecps}(e)(n+1)(\phi)(\lambda n_2, x. (n_2, k \ x)) \text{ in} \\
& \quad (n_1, \Lambda k. e') \\
\\
& \text{ret } (f) (n, e) = (n, f(e)) \\
\\
& \text{ecps} \in \mathbf{Exp}_{\mathcal{S}} \rightarrow \mathbb{N} \rightarrow (\mathbf{Var} \rightarrow \mathbf{TVar} \cup \mathbf{Lbl}) \rightarrow (\mathbb{N} \rightarrow \mathbf{TVar} \cup \mathbf{Lbl} \rightarrow \mathbb{N} \times \mathbf{Exp}_{\mathcal{I}}) \rightarrow \mathbb{N} \times \mathbf{Exp}_{\mathcal{I}} \\
& \text{ecps}(F)(n)(\phi)(\kappa) = \kappa(n)(F) \\
& \text{ecps}(x)(n)(\phi)(\kappa) = \kappa(n)(\phi(x)) \\
& \text{ecps}(\langle \rangle)(n)(\phi)(\kappa) = \text{ret } (\lambda e. \text{let } \text{tvar}(n) = \langle \rangle \text{ in } e) (\kappa(n+1)(\text{tvar}(n))) \\
& \text{ecps}(m)(n)(\phi)(\kappa) = \text{ret } (\lambda e. \text{let } \text{tvar}(n) = m \text{ in } e) (\kappa(n+1)(\text{tvar}(n))) \\
& \text{ecps}(\text{input})(n)(\phi)(\kappa) = \text{ret } (\lambda e. \text{tvar}(n) \leftarrow \text{input}; e) (\kappa(n+1)(\text{tvar}(n))) \\
& \text{ecps}(\text{roll } e)(n)(\phi)(\kappa) = \text{ecps}(e)(n)(\phi)(\kappa) \\
& \text{ecps}(\text{unroll } e)(n)(\phi)(\kappa) = \text{ecps}(e)(n)(\phi)(\kappa) \\
& \text{ecps}(\text{pack } e)(n)(\phi)(\kappa) = \text{ecps}(e)(n)(\phi)(\kappa)
\end{aligned}$$

Figure 5.2: CPS transformation (part 1 of 3).

$$\begin{aligned}
ecps(\text{unpack } y \text{ as } e_1 \text{ in } e_2)(n)(\phi)(\kappa) &= \\
&ecps(e_1)(n)(\phi)(\lambda(n', x_1). \\
&\quad ecps(e_2)(n')(\phi[x_1/y])(\kappa)) \\
ecps(e.1)(n)(\phi)(\kappa) &= \\
&ecps(e)(n)(\phi)(\lambda(n', x). \\
&\quad \text{ret } (\lambda e'. \text{ let } tvar(n') = x.1 \text{ in } e') (\kappa(n' + 1)(tvar(n')))) \\
ecps(e.2)(n)(\phi)(\kappa) &= \\
&ecps(e)(n)(\phi)(\lambda(n', x). \\
&\quad \text{ret } (\lambda e'. \text{ let } tvar(n') = x.2 \text{ in } e') (\kappa(n' + 1)(tvar(n')))) \\
ecps(\text{inl } e)(n)(\phi)(\kappa) &= \\
&ecps(e)(n)(\phi)(\lambda(n', x). \\
&\quad \text{ret } (\lambda e'. \text{ let } tvar(n') = \text{inl } x \text{ in } e') (\kappa(n' + 1)(tvar(n')))) \\
ecps(\text{inr } e)(n)(\phi)(\kappa) &= \\
&ecps(e)(n)(\phi)(\lambda(n', x). \\
&\quad \text{ret } (\lambda e'. \text{ let } tvar(n') = \text{inr } x \text{ in } e') (\kappa(n' + 1)(tvar(n')))) \\
ecps(\text{ref } e)(n)(\phi)(\kappa) &= \\
&ecps(e)(n)(\phi)(\lambda(n', x). \\
&\quad \text{ret } (\lambda e'. tvar(n') \leftarrow \text{ref } x; e') (\kappa(n' + 1)(tvar(n')))) \\
ecps(!e)(n)(\phi)(\kappa) &= \\
&ecps(e)(n)(\phi)(\lambda(n', x). \\
&\quad \text{ret } (\lambda e'. tvar(n') \leftarrow !x; e') (\kappa(n' + 1)(tvar(n')))) \\
ecps(\text{output } e)(n)(\phi)(\kappa) &= \\
&ecps(e)(n)(\phi)(\lambda(n', x). \\
&\quad \text{ret } (\lambda e'. \text{ let } tvar(n') = \langle \rangle \text{ in output } x; e') (\kappa(n' + 1)(tvar(n')))) \\
ecps(\langle e_1, e_2 \rangle)(n)(\phi)(\kappa) &= \\
&ecps(e_1)(n)(\phi)(\lambda(n', x_1). \\
&\quad ecps(e_2)(n')(\phi)(\lambda(n'', x_2). \\
&\quad \quad \text{ret } (\lambda e. \text{ let } tvar(n'') = \langle x_1, x_2 \rangle \text{ in } e) (\kappa(n'' + 1)(tvar(n''))))) \\
ecps(e_1 \odot e_2)(n)(\phi)(\kappa) &= \\
&ecps(e_1)(n)(\phi)(\lambda(n', x_1). \\
&\quad ecps(e_2)(n')(\phi)(\lambda(n'', x_2). \\
&\quad \quad \text{ret } (\lambda e. \text{ let } tvar(n'') = x_1 \odot x_2 \text{ in } e) (\kappa(n'' + 1)(tvar(n''))))) \\
ecps(e_1 := e_2)(n)(\phi)(\kappa) &= \\
&ecps(e_1)(n)(\phi)(\lambda(n', x_1). \\
&\quad ecps(e_2)(n')(\phi)(\lambda(n'', x_2). \\
&\quad \quad \text{ret } (\lambda e. \text{ let } tvar(n'') = \langle \rangle \text{ in } x_1 := x_2; e) (\kappa(n'' + 1)(tvar(n'')))))
\end{aligned}$$

Figure 5.3: CPS transformation (part 2 of 3).

$$\begin{aligned}
& ecps(e_1 == e_2)(n)(\phi)(\kappa) = \\
& \quad ecps(e_1)(n)(\phi)(\lambda(n', x_1). \\
& \quad \quad ecps(e_2)(n')(\phi)(\lambda(n'', x_2). \\
& \quad \quad \quad ret (\lambda e. let tvar(n'') = x_1 == x_2 in e) (\kappa(n'' + 1)(tvar(n''))))) \\
\\
& ecps(ifnz e then e_1 else e_2)(n)(\phi)(\kappa) = \\
& \quad ecps(e)(n)(\phi)(\lambda(n', x). \\
& \quad \quad let (n'', e'_1) := ecps(e_1)(n')(\phi)(\kappa) in \\
& \quad \quad \quad ret (\lambda e'_2. ifnz x then e'_1 else e'_2) (ecps(e_2)(n'')(\phi)(\kappa))) \\
\\
& ecps(e_1 e_2)(n)(\phi)(\kappa) = \\
& \quad ecps(e_1)(n)(\phi)(\lambda(n', x_1). \\
& \quad \quad ecps(e_2)(n')(\phi)(\lambda(n'', x_2). \\
& \quad \quad \quad ret (\lambda e'. let kvar(n'') = cont tvar(n'' + 1). e' in x_1 x_2 kvar(n'')) \\
& \quad \quad \quad \quad (\kappa(n'' + 2)(tvar(n'' + 1))))) \\
\\
& ecps(e \square)(n)(\phi)(\kappa) = \\
& \quad ecps(e)(n)(\phi)(\lambda(n', x). \\
& \quad \quad ret (\lambda e'. let kvar(n') = cont tvar(n' + 1). e' in x \square kvar(n')) \\
& \quad \quad \quad (\kappa(n' + 2)(tvar(n' + 1))))) \\
\\
& ecps(case e (y. e_1) (y. e_2))(n)(\phi)(\kappa) = \\
& \quad ecps(e)(n)(\phi)(\lambda(n', x). \\
& \quad \quad let \phi' := \phi[tvar(n')/y] in \\
& \quad \quad \quad let (n'', e'_1) := ecps(e_1)(n' + 1)(\phi')(\kappa) in \\
& \quad \quad \quad \quad ret (\lambda e'_2. case x (tvar(n'). e'_1) (tvar(n'). e'_2)) (ecps(e_2)(n'')(\phi')(\kappa))) \\
\\
& ecps(fix f(y). e)(n)(\phi)(\kappa) = \\
& \quad let f' := tvar(n + 1) in \\
& \quad let y' := tvar(n + 2) in \\
& \quad let k := kvar(n + 3) in \\
& \quad let \phi' := \phi[f'/f][y'/y] in \\
& \quad let (n_1, e') := ecps(e)(n + 4)(\phi')(\lambda n_2, x. (n_2, k x)) in \\
& \quad \quad ret (\lambda e''. let tvar(n) = fix f'(y', k). e' in e'') (\kappa(n_1)(tvar(n))) \\
\\
& ecps(\Lambda. e)(n)(\phi)(\kappa) = \\
& \quad let k := kvar(n + 1) in \\
& \quad let (n_1, e') := ecps(e)(n + 2)(\phi')(\lambda n_2, x. (n_2, k x)) in \\
& \quad \quad ret (\lambda e''. let tvar(n) = \Lambda k. e' in e'') (\kappa(n_1)(tvar(n)))
\end{aligned}$$

Figure 5.4: CPS transformation (part 3 of 3).

$tvar(n)$. It passes this variable on as input to the continuation, together with the updated variable counter $n + 1$ in order to ensure that $tvar(n)$ will not be chosen again.

- The translation of `pack e` simply translates e and forgets about the packing since there is no such concept in \mathcal{I} .
- The translation of `inl e` translates e with a special continuation. This continuation introduces a let-binding for inl_x , where x is the variable (either term variable or label) holding the value of e 's translation. The body of the let-binding is determined by the meta-level continuation κ .
- The translation of `fix $f(y). e$` creates a new let-binding for the function definition. The body of the let-binding is determined by κ . The body of the function is the translation of e relative to the meta-level continuation $\lambda n_2, x. (n_2, k \ x)$, which is a reified version of the function's formal continuation argument k . All bound variables are chosen to be fresh, and the renaming ϕ' used in translation of the function body knows which fresh variables correspond to occurrences of f and y in e . Note that *fcps* necessarily translates top-level functions slightly differently.
- The translation of $e_1 \ e_2$ translates e_1 with a continuation that translates e_2 , in turn with a continuation that issues the actual function call. The original continuation κ is reified into an object-level continuation by introducing a continuation let-binding. The continuation variable $kvar(n'')$ is then passed to the function x_1 in the call.

5.2.1.1 Example

Consider a source module M_S consisting of a single function labelled G :

$$G = \text{fix } f(x). \langle F \ 42, !x \rangle$$

It returns a pair whose first component is the result of calling the imported function F on 42. Its second component is the value stored in the argument reference x .

The \mathcal{I} module $cps(M_S)$, *i.e.*, the result of transforming M_S , consists of the following single function (up to the exact choice of fresh variables):

$$\begin{aligned} G = & \text{fix } f'(y, k). \\ & \text{let } y_1 = 42 \text{ in} \\ & \text{let } k_1 = (\text{cont } y_2. \ y_3 \leftarrow !y; \\ & \quad \text{let } y_4 = \langle y_2, y_3 \rangle \text{ in} \\ & \quad \quad k_1 \ y_4) \text{ in} \\ & F \ y_1 \ k_1 \end{aligned}$$

5.2.2 Verification

It is easy to show that the CPS transformation of a well-typed \mathcal{S} module is a well-formed \mathcal{I} module with the same imports and exports. Recall that \mathcal{I} well-formedness includes the affinity of continuation variables (Section 4.2.4.2). It does not include the uniqueness of bound variables, but that does hold as well. These facts are important for the correctness of the subsequent transformations performed by Pilsner.

Theorem 17.

$$\frac{\Gamma \vdash M : \Gamma'}{|\Gamma| \vdash \text{cps}(M) : |\Gamma'| \quad \text{uniqmod}(\text{cps}(M))}$$

The main correctness statement is the following.

Theorem 18.

$$\frac{\Gamma \vdash M : \Gamma'}{\Gamma \vdash \text{cps}(M) \lesssim_{\mathcal{IS}} M : \Gamma'}$$

For the proof of this, we pick the *empty local world* $w^\emptyset \in \text{World}^{\Omega_{\mathcal{IS}} \cdot \top}$ because the effects of the CPS transformation on memory are trivial. We define the empty world generically, such that it can be reused. Note the use of *gwf* for flexibility.

Definition 36 (Empty local world).

$$\begin{aligned} w^\emptyset &\in \text{LWorld}_{A,B}^T \\ w^\emptyset.\mathsf{T.S} &:= 1 \\ w^\emptyset.\mathsf{T}.\sqsubseteq &:= 1 \times 1 \\ w^\emptyset.\mathsf{T}.\sqsubseteq_{\text{pub}} &:= 1 \times 1 \\ w^\emptyset.\mathsf{C}(G)(s) &:= \{(\emptyset, \emptyset)\} \\ w^\emptyset.\mathsf{N.NS} &:= \emptyset \\ w^\emptyset.\mathsf{N.NR}(G)(s)(\mathbf{n}) &:= \emptyset \\ w^\emptyset.\mathsf{O} &:= \text{gwf} \end{aligned}$$

The proof of Theorem 18 boils down to finding a relation $\lesssim_{\mathcal{IS}}^{\text{cps}}$ on expressions and their translations that satisfies

$$\frac{\Gamma \vdash e : \tau}{\Gamma \vdash \text{ecps}(e) \lesssim_{\mathcal{IS}}^{\text{cps}} e : \tau}$$

and at the same time is strong enough to let us derive the module-level correctness. We now present and justify our choice of $\lesssim_{\mathcal{IS}}^{\text{cps}}$.

Definition 37.

$$\begin{aligned}
\Gamma \vdash f \lesssim_{\mathcal{IS}} e : \tau := & \\
& \exists i. \forall \delta, n, \phi, G, s_0, s, \sigma, \gamma, j, \mathbf{k}_a, \mathbf{k}_b, \tau', \kappa, K, e'. \\
& (_, e') = f(n)(\phi)(\kappa) \wedge \\
& G \in \mathbf{GK}_{w\emptyset\uparrow} \wedge s_0 \sqsupseteq_{\text{pub}} s \wedge \\
& (\forall F:\tau'' \in \Gamma. \exists \mathbf{v}_a, \mathbf{v}_b. \sigma(F) = \mathbf{v}_a \wedge s.\rho(F) = \mathbf{v}_b \wedge (\mathbf{v}_a, \mathbf{v}_b) \in \overline{G}(s)(\delta\tau'')) \wedge \\
& (\forall x:\tau'' \in \Gamma. \exists \mathbf{v}_a, \mathbf{v}_b. \sigma(\phi(x)) = \mathbf{v}_a \wedge \gamma(x) = \mathbf{v}_b \wedge (\mathbf{v}_a, \mathbf{v}_b) \in \overline{G}(s)(\delta\tau'')) \wedge \\
& \text{uniquars}(e') \wedge \text{bv}(e') \cap \text{dom}(\sigma) = \emptyset \wedge \\
& ((\sigma, \kappa), K) \in \mathfrak{K}(j)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)(\delta\tau)(\tau') \\
\implies & ((\emptyset, (\sigma, e'), (\emptyset, \emptyset, K[\gamma(e)])) \in \mathbf{E}_{w\emptyset\uparrow}(i+j)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)(\perp)(\tau'))
\end{aligned}$$

Here, f is a meta-level function of type $\mathbb{N} \rightarrow (\text{Var} \rightarrow \text{TVar} \cup \text{Lbl}) \rightarrow (\mathbb{N} \rightarrow \text{TVar} \cup \text{Lbl} \rightarrow \mathbb{N} \times \text{Exp}_{\mathcal{I}}) \rightarrow \mathbb{N} \times \text{Exp}_{\mathcal{I}}$ and e an \mathcal{S} expression. Roughly, we want to say that f (think: $\text{ecps}(e)$) is related to e by the well-known \mathbf{E} relation. Of course, for this to make sense, we must first turn them into suitable configurations.

First we need to quantify over the arguments to f , namely n , ϕ , and κ . Following the shown definition, let us call the resulting expression e' . Now we have two open expressions, e' and e . To close these with respect to labels—as listed in Γ —we quantify over an environment σ for e' and a state s containing a (label-only) environment $s.\rho$ for e . These should of course not be arbitrary, but related by (the value closure of) the global knowledge, over which we quantify as well, at the type indicated by Γ . Since that type might contain free variables as well, we first close it with a substitution δ mapping type variables to arbitrary closed types (which may contain type names).

Besides labels, Γ and thus e can also contain variables (recall that $\Gamma \vdash e : \tau$). To close e with respect to those, we quantify over a value substitution γ . These \mathcal{S} variables correspond to \mathcal{I} term variables and labels, as given by the renaming ϕ used in the CPS transformation. We require that the values in γ are related to the values in σ accordingly, *i.e.*, if $x:\tau'' \in \Gamma$ then the value in σ for the renaming of x is related to $\gamma(x)$. In addition to closing e using γ , we also put it in an evaluation context K , since the expression may occur anywhere in the source program.

The key remaining condition enforces a connection between this evaluation context K and the meta-level continuation κ used to generate e' . But what is \mathfrak{K} ? Certainly we don't want to describe their relationship syntactically. Can we not just say that these continuations (or some conversions thereof) are related by PILS's \mathbf{K} relation? The issue is that the way in which e' and e “return” to these continuations is very different from how functions return to their continuations and thus is not compatible with \mathbf{K} . This is not specific to the CPS transformation but rather is a general observation: the “continuations” inherent in the translation of inductively-defined expressions may be very different (in shape and/or behavior) from the continuations associated with function calls that are exposed in our language specification and configuration queries. This will become even more apparent when we get to code generation in Section 5.10.

So let us look at \mathfrak{R} , defined as follows.

Definition 38.

$$\begin{aligned} \mathfrak{R}(i)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)(\tau)(\tau') := & \{ ((\sigma, \kappa), K) \mid \\ & \forall G', s', \mathbf{v}_a, \mathbf{v}_b, x, \sigma', n, e''. \\ & (-, e'') = \kappa(n)(x) \wedge \\ & G' \supseteq G \wedge s' \sqsupseteq_{\text{pub}} s_0 \wedge s' \sqsupseteq s \wedge \\ & (\mathbf{v}_a, \mathbf{v}_b) \in \overline{G'}(s')(\tau) \wedge \\ & \sigma \subseteq \sigma' \wedge \sigma'(x) = \mathbf{v}_a \wedge \\ & \text{uniquvars}(e'') \wedge \text{dom}(\sigma') \cap \text{bv}(e') = \emptyset \\ \implies & ((\emptyset, (\sigma', e'')), (\emptyset, \emptyset, K[\mathbf{v}_b])) \in \mathbf{E}_{w\emptyset\uparrow}(i)(\mathbf{k}_a, \mathbf{k}_b)(G')(s_0)(s')(\perp)(\tau) \} \end{aligned}$$

First, it gives arguments n and x to κ in order to generate the actual expression e'' . Then it roughly says that e'' in an environment σ' that maps the term variable or label x to a value \mathbf{v}_a is related to $K[\mathbf{v}_b]$ by the usual \mathbf{E} relation if \mathbf{v}_a and \mathbf{v}_b are related (by the value closure of the global knowledge, as usual). Without any further restrictions on σ' , however, we would have little hope of showing continuations related by \mathfrak{R} . We must insist that σ' be an extension of the earlier σ in which we started executing e' . Note that x may or may not have already existed in σ . Also note that we quantify over a larger global knowledge G' because e' and e (from $\lesssim_{\mathcal{IS}}$) may make external calls before the continuations are reached.

A word on the use of budgets: $\lesssim_{\mathcal{IS}}^{\text{cps}}$ requires that there exists a budget i (intuitively: for f) such that the constructed configurations are related by \mathbf{E} at budget $i + j$, where j is the continuation budget. This is very natural, as the configuration consists of both the continuation and the expression. The assumption that all continuations are related at the same budget j does not impose a restriction because one can always choose j to be large enough for all continuations in question¹.

As intended, we can prove the following key lemma about $\lesssim_{\mathcal{IS}}^{\text{cps}}$.

Lemma 68.

$$\frac{\Gamma \vdash e : \tau}{\Gamma \vdash \text{ecps}(e) \lesssim_{\mathcal{IS}}^{\text{cps}} e : \tau}$$

Proof. By induction on e , with a nested coinduction on \mathbf{E} when e is a recursive function $\text{fix } f(y, k). e'$. \square

This lemma is then used in the proof of Theorem 18, where it is applied to the body of top-level functions.

5.3 Infrastructure for Optimizations

After CPS conversion and thus translation into \mathcal{I} , Pilsner performs several transformations at the intermediate level, *i.e.*, transformations from \mathcal{I} to \mathcal{I} . In preparation of verifying these, we develop some common infrastructure.

¹The \mathbf{E} relation, and thus \mathfrak{R} , is closed under budget increase.

5.3.1 Relating Open Expressions

Somewhat similar to $\lesssim_{\mathcal{IS}}^{\text{cps}}$, we define a relation between open \mathcal{I} expressions. This is a fairly straightforward lifting of **E**:

Definition 39.

$$\begin{aligned} \Gamma \vdash e_a \lesssim_w^\phi e_b := & \\ & \exists i. \forall G, s_0, s, \sigma_a, \sigma_b, j, \mathbf{k}_a, \mathbf{k}_b. \\ & G \in \mathbf{GK}_{w\uparrow} \wedge s \sqsupseteq_{\text{pub}} s_0 \wedge \\ & (\forall x \in \Gamma. \exists \mathbf{v}_a, \mathbf{v}_b. \sigma_a(\phi(x)) = \mathbf{v}_a \wedge \sigma_b(x) = \mathbf{v}_b \wedge (\mathbf{v}_a, \mathbf{v}_b) \in \overline{G}(s)) \wedge \\ & (\forall k \in \Gamma. \exists \mathbf{k}'_a, \mathbf{k}'_b. \sigma_a(\phi(k)) = \mathbf{k}'_a \wedge \sigma_b(k) = \mathbf{k}'_b \wedge \\ & \quad (\mathbf{k}'_a, \mathbf{k}'_b) \in \mathbf{K}_{w\uparrow}(j)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)) \\ \implies & ((\emptyset, (\sigma_a, e_a)), (\emptyset, (\sigma_b, e_b))) \in \mathbf{E}_{w\uparrow}(i+j)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)(\perp) \end{aligned}$$

While $\lesssim_{\mathcal{IS}}^{\text{cps}}$ was defined in terms of the empty local world, here we parameterize the relation by an arbitrary local world $w \in \text{LWorld}^{\Omega_{\text{IT}}.\top}$ for flexibility (we will need to use a non-trivial world in some proofs). The relation roughly says that the \mathcal{I} expressions e_a and e_b must be related by **E** when paired with related environments σ_a and σ_b . Related environments must map term variables and labels x to values related by (the value closure) of the global knowledge. They must also map continuation variables k to continuations related by **K**. In contrast to $\lesssim_{\mathcal{IS}}^{\text{cps}}$, here we do not need to define a custom continuation relation because the notion of internal continuation coincides with the “official” notion of continuation, and so does their calling convention.

Note how we allow any variable in the source e_b to have been renamed in the target e_a by applying a renaming ϕ before the lookup in σ_a . This renaming is an index of the relation.

Besides control expressions e , there are pure expressions a in \mathcal{I} . For these we define a similar relation (in fact, we overload the notation). This relation roughly says that, given related environments, if the target expression a_b successfully evaluates to a value, then so does the source expression a_a , and the resulting values are related. Regarding the environments, the relation only cares about term variables and labels in Γ , not about continuation variables—recall that a well-formed pure expression does not refer to any continuation variables (Section 4.2.4.2).

Definition 40.

$$\begin{aligned} \Gamma \vdash a_a \lesssim_w^\phi a_b := & \\ & \forall G, s, \sigma_a, \sigma_b, \mathbf{v}_b. \\ & G \in \mathbf{GK}_{w\uparrow} \wedge \\ & (\forall x \in \Gamma. \exists \mathbf{v}'_a, \mathbf{v}'_b. \sigma_a(\phi(x)) = \mathbf{v}'_a \wedge \sigma_b(x) = \mathbf{v}'_b \wedge (\mathbf{v}'_a, \mathbf{v}'_b) \in \overline{G}(s)) \wedge \\ & \llbracket a_b \rrbracket_{\sigma_b} = \mathbf{v}_b \\ \implies & \exists \mathbf{v}_a. \llbracket a_a \rrbracket_{\sigma_a} = \mathbf{v}_a \wedge (\mathbf{v}_a, \mathbf{v}_b) \in \overline{G}(s) \end{aligned}$$

For these relations, we prove typical compatibility lemmas, which state that each relation is preserved by the respective forms of \mathcal{I} expressions. These lemmas are very helpful in verifying transformations where subexpressions stay unchanged (as we will see in the coming sections). They are shown in Figures 5.5 and 5.6. For each well-formedness rule from Figure 4.5, we have a corresponding lemma here.

In contrast to the compatibility lemmas that we have seen for PBs in Section 3.7.3, here we allow variables of the source expression to be renamed in the target expression. As a consequence of this, lemmas for constructs that introduce bound variables require a side condition of the form $y_a \triangleleft_Z^\phi y_b$, where y_a is the variable bound by the left-side term and y_b is the variable bound by the right-side term. This relation is defined as follows.

Definition 41.

$$\begin{aligned} y_a \triangleleft_Z^\phi y_b &:= \forall y'_b \in Z. \phi(y'_b) = y_a \implies y'_b = y_b \\ k_a \triangleleft_Z^\phi k_b &:= \forall k'_b \in Z. \phi(k'_b) = k_a \implies k'_b = k_b \end{aligned}$$

This side condition enforces that a variable y_b can be renamed to y_a only if no other variable has been renamed to y_a . A special case of these lemmas is obtained when not renaming anything, *i.e.*, when $\phi = id$. In that case, all these side conditions hold trivially.

The lemmas are straightforward but tedious to show. Only the one about recursive functions requires a proof by coinduction, as one would expect.

Note that the lemma about allocation ($y \leftarrow \text{ref } x; e$) requires stability of w : in the proof of the lemma, when we extend the global state by adding a new pair of references to the partial bijection, we need to know that the local world's configuration relation can still be satisfied. None of the other lemmas requires stability.

Combining these lemmas with a straightforward induction yields the following lemma.

Lemma 69.

$$\frac{\Gamma \vdash a \quad w \in \text{stable}(\Omega_{II}) \quad \forall z \in \Gamma. \phi(z) \in \text{bv}(e) \implies \phi(z) = z}{\Gamma \vdash \phi(a) \lesssim_w^\phi a}$$

$$\frac{\Gamma \vdash e \quad w \in \text{stable}(\Omega_{II}) \quad \forall z \in \Gamma. \phi(z) \in \text{bv}(e) \implies \phi(z) = z}{\Gamma \vdash \phi(e) \lesssim_w^\phi e}$$

In the conclusions, $\phi(a)$ and $\phi(e)$ denote the application of the renaming ϕ to each *free* variable in a and e , respectively. The side condition on ϕ rules out that any of these is renamed to a variable captured by a binder in the expression. Naturally, if the renaming is the identity function, this holds again trivially and thus the lemma becomes a (conditional) reflexivity property:

$$\begin{array}{c}
\frac{}{Z \vdash \langle \rangle \lesssim_w^\phi \langle \rangle} \quad \frac{}{Z \vdash n \lesssim_w^\phi n} \quad \frac{x_b \in Z \quad x_a = \phi(x_b)}{Z \vdash x_{a.1} \lesssim_w^\phi x_{b.1}} \quad \frac{x_b \in Z \quad x_a = \phi(x_b)}{Z \vdash x_{a.2} \lesssim_w^\phi x_{b.2}} \\
\\
\frac{x_b \in Z \quad x_a = \phi(x_b)}{Z \vdash \text{inl } x_a \lesssim_w^\phi \text{inl } x_b} \quad \frac{x_b \in Z \quad x_a = \phi(x_b)}{Z \vdash \text{inr } x_a \lesssim_w^\phi \text{inr } x_b} \\
\\
\frac{x_b^1, x_b^2 \in Z \quad x_a^1 = \phi(x_b^1) \quad x_a^2 = \phi(x_b^2)}{Z \vdash \langle x_a^1, x_a^2 \rangle \lesssim_w^\phi \langle x_b^1, x_b^2 \rangle} \quad \frac{x_b^1, x_b^2 \in Z \quad x_a^1 = \phi(x_b^1) \quad x_a^2 = \phi(x_b^2)}{Z \vdash x_a^1 \odot x_a^2 \lesssim_w^\phi x_b^1 \odot x_b^2} \\
\\
\frac{x_b^1, x_b^2 \in Z \quad x_a^1 = \phi(x_b^1) \quad x_a^2 = \phi(x_b^2)}{Z \vdash x_a^1 == x_a^2 \lesssim_w^\phi x_b^1 == x_b^2} \\
\\
\frac{Z^\dagger, f_b, y_b, k_b \vdash e_a \lesssim_w^{\phi'} e_b \quad \phi' = \phi[f_a/f_b][y_a/y_b][k_a/k_b] \quad f_a = y_a \implies f_b = y_b \quad f_a \triangleleft_Z^\phi f_b \quad y_a \triangleleft_Z^\phi y_b}{Z \vdash \text{fix } f_a(y_a, k_a). e_a \lesssim_w^\phi \text{fix } f_b(y_b, k_b). e_b} \\
\\
\frac{Z^\dagger, k_b \vdash e_a \lesssim_w^{\phi'} e_b \quad \phi' = \phi[k_a/k_b]}{Z \vdash \Lambda k_a. e_a \lesssim_w^\phi \Lambda k_b. e_b} \\
\\
\text{.....} \\
\\
\frac{\Gamma \vdash a_a \lesssim_w^\phi a_b \quad \Gamma, y_b \vdash e_a \lesssim_w^{\phi'} e_b \quad \phi' = \phi[y_a/y_b] \quad y_a \triangleleft_Z^\phi y_b}{\Gamma \vdash \text{let } y_a = a_a \text{ in } e_a \lesssim_w^\phi \text{let } y_b = a_b \text{ in } e_b} \\
\\
\frac{\Gamma, y_b \vdash e_a \lesssim_w^{\phi'} e_b \quad \phi' = \phi[y_a/y_b] \quad y_a \triangleleft_Z^\phi y_b \quad \Gamma, k_b \vdash e'_a \lesssim_w^{\phi''} e'_b \quad \phi'' = \phi[k_a/k_b] \quad k_a \triangleleft_Z^\phi k_b}{\Gamma \vdash \text{let } k_a = \text{cont } y_a. e_a \text{ in } e'_a \lesssim_w^\phi \text{let } k_b = \text{cont } y_b. e_b \text{ in } e'_b} \\
\\
\frac{\Gamma, y_b \vdash e_a \lesssim_w^{\phi'} e_b \quad \phi' = \phi[y_a/y_b] \quad y_a \triangleleft_Z^\phi y_b}{\Gamma \vdash y_a \leftarrow \text{input}; e_a \lesssim_w^\phi y_b \leftarrow \text{input}; e_b}
\end{array}$$

Figure 5.5: Compatibility lemmas for \lesssim_w^ϕ

$$\begin{array}{c}
\frac{x_b \in Z \quad x_a = \phi(x_b) \quad \Gamma \vdash e_a \lesssim_w^\phi e_b}{\Gamma \vdash \text{output } x_a; e_a \lesssim_w^\phi \text{output } x_b; e_b} \\
\\
\frac{x_b \in Z \quad x_a = \phi(x_b) \quad \Gamma \vdash e_a^1 \lesssim_w^\phi e_b^1 \quad \Gamma \vdash e_a^2 \lesssim_w^\phi e_b^2}{\Gamma \vdash \text{ifnz } x_a \text{ then } e_a^1 \text{ else } e_a^2 \lesssim_w^\phi \text{ifnz } x_b \text{ then } e_b^1 \text{ else } e_b^2} \\
\\
\frac{x_b \in Z \quad x_a = \phi(x_b) \quad \Gamma, y_b \vdash e_a^1 \lesssim_w^{\phi'} e_b^1 \quad \Gamma, y_b \vdash e_a^2 \lesssim_w^{\phi'} e_b^2 \quad \phi' = \phi[y_a/y_b] \quad y_a \triangleleft_Z^\phi y_b}{\Gamma \vdash \text{case } x_a (y_a. e_a^1) (y_a. e_a^2) \lesssim_w^\phi \text{case } x_b (y_b. e_b^1) (y_b. e_b^2)} \\
\\
\frac{x_b^1, x_b^2, k_b \in Z \quad x_a^1 = \phi(x_b^1) \quad x_a^2 = \phi(x_b^2) \quad k_a = \phi(k_b)}{Z \vdash x_a^1 x_a^2 k_a \lesssim_w^\phi x_b^1 x_b^2 k_b} \\
\\
\frac{x_b, k_b \in Z \quad x_a = \phi(x_b) \quad k_a = \phi(k_b)}{Z \vdash x_a \sqcap k_a \lesssim_w^\phi x_b \sqcap k_b} \quad \frac{k_b, x_b \in Z \quad k_a = \phi(k_b) \quad x_a = \phi(x_b)}{Z \vdash k_a x_a \lesssim_w^\phi k_b x_b} \\
\\
\frac{x_b \in Z \quad x_a = \phi(x_b) \quad \Gamma, y_b \vdash e_a \lesssim_w^{\phi'} e_b \quad \phi' = \phi[y_a/y_b] \quad y_a \triangleleft_Z^\phi y_b \quad w \in \text{stable}(\Omega_{II})}{\Gamma \vdash y_a \leftarrow \text{ref } x_a; e_a \lesssim_w^\phi y_b \leftarrow \text{ref } x_b; e_b} \\
\\
\frac{x_b \in Z \quad x_a = \phi(x_b) \quad \Gamma, y_b \vdash e_a \lesssim_w^{\phi'} e_b \quad \phi' = \phi[y_a/y_b] \quad y_a \triangleleft_Z^\phi y_b}{\Gamma \vdash y_a \leftarrow !x_a; e_a \lesssim_w^\phi y_b \leftarrow !x_b; e_b} \\
\\
\frac{x_b^1, x_b^2 \in Z \quad x_a^1 = \phi(x_b^1) \quad x_a^2 = \phi(x_b^2) \quad \Gamma \vdash e_a \lesssim_w^\phi e_b}{\Gamma \vdash x_a^1 := x_a^2; e_a \lesssim_w^\phi x_b^1 := x_b^2; e_b}
\end{array}$$

Figure 5.6: Compatibility lemmas for \lesssim_w^ϕ (cntd.)

Lemma 70.

$$\frac{\Gamma \vdash a \quad w \in \text{stable}(\Omega_{II})}{\Gamma \vdash a \lesssim_w a} \quad \frac{\Gamma \vdash e \quad w \in \text{stable}(\Omega_{II})}{\Gamma \vdash e \lesssim_w e}$$

Here we write \lesssim_w short for \lesssim_w^{id} .

With this and a little extra work, we can also lift the identity-renaming versions of the compatibility lemmas in Figures 5.5 and 5.6 from \lesssim_w to its reflexive-transitive closure \lesssim_w^* . These will be used to establish the correctness of the framework presented in the next section. Let us just show and prove one of these rules.

Lemma 71.

$$\frac{\Gamma \vdash a_a \lesssim_w^* a_b \quad \Gamma \vdash a_b \quad \Gamma, y \vdash e_a \lesssim_w^* e_b \quad \Gamma, y \vdash e_b \quad w \in \text{stable}(\Omega_{II})}{\Gamma \vdash \text{let } y = a_a \text{ in } e_a \lesssim_w^* \text{let } y = a_b \text{ in } e_b}$$

Proof. We do an induction on the first reflexive-transitive closure in the premise.

Case $a_a = a_b$. We do another induction, this time on the second reflexive-transitive closure in the premise.

Case $e_a = e_b$. Using the extra conditions on well-formedness and stability, we are done by Lemma 70.

Case $\Gamma^\dagger, k_b \vdash e_a \lesssim_w e$ **and** $\Gamma^\dagger, k_b \vdash e \lesssim_w^* e_b$. By the inner inductive hypothesis we have

$$\Gamma \vdash \text{let } y = a_a \text{ in } e \lesssim_w^* \text{let } y = a_b \text{ in } e_b,$$

so it suffices to show

$$\Gamma \vdash \text{let } y = a_a \text{ in } e_a \lesssim_w \text{let } y = a_a \text{ in } e.$$

Since $\Gamma \vdash a_a \lesssim_w a_a$ by Lemma 70, this follows by the compatibility lemma for let-bindings.

Case $\Gamma \vdash a_a \lesssim_w a$ **and** $\Gamma \vdash a \lesssim_w^* a_b$. We do a case analysis (not an induction) on $\Gamma, y \vdash e_a \lesssim_w^* e_b$.

Case $e_a = e_b$. By the outer inductive hypothesis we have

$$\Gamma \vdash \text{let } y = a \text{ in } e_a \lesssim_w^* \text{let } y = a_b \text{ in } e_b,$$

so it suffices to show

$$\Gamma \vdash \text{let } y = a_a \text{ in } e_a \lesssim_w \text{let } y = a \text{ in } e_a.$$

Since $e_a = e_b$, we know $\Gamma, y \vdash e_a$ and can use Lemma 70 to get $\Gamma, y \vdash e_a \lesssim_w e_a$. The goal follows by the compatibility lemma for let-bindings.

Case $\Gamma, y \vdash e_a \lesssim_w e$ **and** $\Gamma, y \vdash e \lesssim_w^* e_b$. Applying the inductive hypothesis to $\Gamma, y \vdash e \lesssim_w^* e_b$ (not to $\Gamma, y \vdash e_a \lesssim_w^* e_b$) yields

$$\Gamma \vdash \text{let } y = a \text{ in } e \lesssim_w^* \text{let } y = a_b \text{ in } e_b,$$

so it suffices to show

$$\Gamma \vdash \text{let } y = a_a \text{ in } e_a \lesssim_w \text{let } y = a \text{ in } e.$$

This follows immediately by the compatibility lemma for let-bindings.

□

5.3.2 Annotating Expressions With Transformations

In order to simplify the definition and verification of Pilsner’s intermediate transformations (transformations from \mathcal{I} to \mathcal{I}), we developed a simple framework of *transformations as expression annotations*. The idea is to split module-level transformations into two parts: (1) an analysis that is applied to each top-level function and annotates selected subexpressions with to-be-performed micro-transformations (but does not actually rewrite the code); and (2) the micro-transformations themselves, together with their correctness proofs. Given these, we automatically produce a verified module transformation that analyzes the input module and performs transformations according to the produced annotations in a bottom-up manner.

A micro-transformation is a partial function on expressions—it must fail if the preconditions for its correctness do not hold. It is typically non-recursive. For instance, here is the (only) micro-transformation used in the `commute`-pass of Pilsner (Section 5.4):

$$\begin{aligned} \text{commute} &\in \text{Exp} \rightarrow \text{Exp} \\ \text{commute}(e) &:= \text{let } y_2 = a_2 \text{ in let } y_1 = a_1 \text{ in } e_0 \\ &\quad \text{if } e \text{ is } (\text{let } y_1 = a_1 \text{ in let } y_2 = a_2 \text{ in } e_0) \text{ and } x \notin \text{fv}(a_2) \end{aligned}$$

In the case that a micro-transformation fails (for which the analysis is to blame), the subexpression that was being transformed simply stays unchanged, or, alternatively, the whole module transformation—and thus the compiler—fails. In either case, if the module transformation succeeds, the output module is guaranteed to correctly implement the input module. This means that the analysis does not need to be verified—in the worst case, the transformation doesn’t optimize the code.

We now turn to the details.

5.3.2.1 Annotations and Transformers

Definition 42 (Annotated expressions and their erasure). We write Exp^A for the set of \mathcal{I} expressions that are annotated with elements l from A , defined recursively

in the obvious way (here we show only two forms but the others are analogous).

$$\begin{aligned} a &\in \text{BExp}^A ::= \dots \mid [l] \text{ fix } f(y, k). e \mid \dots \\ e &\in \text{Exp}^A ::= \dots \mid [l] \text{ let } y = a \text{ in } e \mid \dots \end{aligned}$$

(A can be arbitrary. If it is a singleton set, then Exp^A is isomorphic to Exp .)

We write $|e|$ for the \mathcal{I} expression obtained by erasing all annotations from e .

Definition 43. A *micro-transformation* is any function in $\text{Exp} \rightarrow \text{Exp}$. Given a set A of annotations, a *transformer for A* is a function $\beta \in A \rightarrow \text{Exp} \rightarrow \text{Exp}$, *i.e.*, it maps each annotation l to a micro-transformation $\beta(l)$. An *analysis for A* is a function $\alpha \in \text{Exp} \rightarrow \text{Exp}^A$ satisfying $|\alpha(e)| = e$ for any e , *i.e.*, it annotates expressions with elements from A but does not otherwise modify the program.

Given a transformer β and an analysis α for the same set of annotations A , we can generate a module-level transformation tf_α^β as follows. For each top-level function definition a , we run the analysis and thereby annotate the function. Then we recursively transform it according to these annotations with the help of atf^β and etf^β .

Definition 44.

$$\begin{aligned} tf_\alpha^\beta &\in \mathbf{Mod} \rightarrow \mathbf{Mod} \\ tf_\alpha^\beta &:= \text{map } (\lambda(F, a). (F, atf_\alpha^\beta(\alpha(a)))) \\ \\ atf^\beta &\in \text{BExp}^A \rightarrow \text{BExp} \\ atf^\beta([l] \text{ fix } f(y, k). e) &:= \\ &\quad \text{let } e' := \text{fix } f(y, k). etf^\beta(e) \text{ in} \\ &\quad \begin{cases} e' & \text{if } \beta(l)(e') = \perp \\ e'' & \text{if } \beta(l)(e') = e'' \end{cases} \\ &\dots \\ \\ etf^\beta &\in \text{Exp}^A \rightarrow \text{Exp} \\ etf^\beta([l] \text{ let } y = a \text{ in } e) &:= \\ &\quad \text{let } e' := (\text{let } y = atf^\beta(a) \text{ in } etf^\beta(e)) \text{ in} \\ &\quad \begin{cases} e' & \text{if } \beta(l)(e') = \perp \\ e'' & \text{if } \beta(l)(e') = e'' \end{cases} \\ &\dots \end{aligned}$$

(We are showing only one case each for atf and etf . The others are analogous.)

Let us illustrate how this works by looking at the case of (annotated) let-bindings:

$$etf^\beta([l] \text{ let } y = a \text{ in } e)$$

First, it recursively transforms the bound term a and the body e and puts them back together in the form of an ordinary (non-annotated) let-binding e' . Now it looks up the micro-transformation for the given annotation l and transforms e' accordingly. If this fails, it simply returns e' , in effect not changing the program at all. Otherwise it returns the resulting expression (which may no longer be a let-binding).

Before moving on to the correctness of tf , let us define two convenient lifting operation on transformers.

Definition 45 (Optional Annotations). Given a transformer β for some A , we define a transformer β_\perp for A_\perp that behaves in the obvious way: where the annotation is \perp , it does nothing; elsewhere it does the same as β .

$$\begin{aligned} \beta_\perp &\in A_\perp \rightarrow \text{Exp} \rightarrow \text{Exp} \\ \beta_\perp(\perp)(e) &:= e \\ \beta_\perp(l)(e) &:= \begin{cases} e & \text{if } \beta(l)(e) = \perp \\ e' & \text{if } \beta(l)(e) = e' \end{cases} \end{aligned}$$

Definition 46 (List Annotations). Given a transformer β for some A , we define a transformer $\bar{\beta}$ for $List(A)$. It performs transformations according to the annotations in the given list from right to left until one fails, in which case the input to that failing transformation is returned. This means that $\bar{\beta}$ itself never fails.

$$\begin{aligned} \bar{\beta} &\in List(A) \rightarrow \text{Exp} \rightarrow \text{Exp} \\ \bar{\beta}(\epsilon)(e) &:= e \\ \bar{\beta}(L, l)(e) &:= \begin{cases} e & \text{if } \beta(l)(e) = \perp \\ \bar{\beta}(L)(e') & \text{if } \beta(l)(e) = e' \end{cases} \end{aligned}$$

Note that one can imagine similar list transformers with slightly different semantics, *e.g.*, one that fails itself if one of the sub-transformations fails.

5.3.2.2 Correctness

Theorem 19 states the correctness of tf : if a well-formed module M with unique variables is transformed according to a correct transformer β , then the resulting module is equally well-formed, also has unique variables, and, crucially, refines M according to the reflexive-transitive closure of \lesssim_{II} .

Theorem 19.

$$\frac{\Gamma \vdash M : \Gamma' \quad \text{uniqmod}(M) \quad \text{correct}(\beta) \quad \forall e. |\alpha(e)| = e}{\Gamma \vdash tf_\alpha^\beta(M) : \Gamma' \quad \text{uniqmod}(tf_\alpha^\beta(M)) \quad \Gamma \vdash tf_\alpha^\beta(M) \lesssim_{II}^* M : \Gamma'}$$

Recall that uniqueness is initially provided by the CPS transformation.

Of course, this theorem relies on the transformer β being *correct*. We define this to be the conjunction of two properties:

1. Syntactic correctness: each transformation preserves well-formedness, which includes affinity of continuation variables, and variable uniqueness (see the comment below).

$$\forall l, \Gamma, e, e'. \quad \Gamma \vdash e \wedge \text{uniq}(\text{bv}(e), \Gamma) \wedge \beta(l)(e) = e' \implies \\ \Gamma \vdash e' \wedge \text{bv}(e') \subseteq \text{bv}(e)$$

2. Semantic correctness: each transformation preserves the behavior.

$$\forall l, \Gamma, e, e', w. \quad \Gamma \vdash e \wedge \text{uniq}(\text{bv}(e), \Gamma) \wedge \beta(l)(e) = e' \wedge w \in \text{stable}(\Omega_{II}) \implies \\ \Gamma \vdash e' \lesssim_w^* e$$

The relation in the conclusion is the one from the previous section.

(It makes sense to separate these conditions because it is typically convenient to prove them separately.)

Because our variable uniqueness property is not compositional, the syntactic correctness criterion above cannot simply require the preservation of uniqueness. Instead, it requires something much stronger (and compositional), namely that the list of variables bound in e' is a sublist of that of e . This is simple and effective but also quite restrictive, and we will have to relax it later on for one of Pilsner's optimizations (Section 5.9).

The semantic part of $\text{correct}(\beta)$ is also fairly strict in that it requires preservation with respect to *any* (stable) local world. This is sufficient for intermediate transformations that do not touch the memory, which are the majority in Pilsner. It is not sufficient, *e.g.*, for dead code elimination (Section 5.7). We leave a suitable generalization of the correctness requirement for future work. Roughly, one would like to turn the universal quantification over w into an existential one, but one must impose some additional constraints in order to be able to find an initial state in the proof of Theorem 19.

For transformers that are correct in the above sense, we can prove the correctness of etf .

Lemma 72.

$$\frac{\Gamma \vdash e \quad |e| = e \quad \text{uniq}(\text{bv}(e), \Gamma) \quad \text{correct}(\beta)}{\Gamma \vdash \text{etf}^\beta(e) \lesssim_{w^\emptyset}^* e}$$

Proof. By induction on e . □

This lemma is strong enough to let us derive the correctness of the module-level transformation tf (Theorem 19).

Finally, recall the transformer liftings defined at the end of the previous section. It should come as no surprise that they preserve correctness:

Theorem 20 (Correctness of transformer liftings).

- $\text{correct}(\beta) \implies \text{correct}(\beta_\perp)$

- $correct(\beta) \implies correct(\bar{\beta})$

We now move on to discuss Pilsner’s individual IL transformations (for pedagogical reasons not in the order in which they are performed).

5.4 The Commute Pass

5.4.1 Transformation

Let us start with the pass that commutes let-expressions in order to enable deduplication. Recall that the goal of the commute pass is to shuffle let-bindings around such that bindings of the same expression end up being “adjacent” (if possible). For instance, in

$$\text{let } x = a \text{ in let } y = b \text{ in let } z = a \text{ in } e$$

we would like to group the two bindings of a together, because then deduplication (Section 5.5) can get rid of one of them.

We achieve this very naively by partially sorting chains of let-bindings according to the expression to which they assign a name. This is implemented with the help of the annotation framework from Section 5.3 as follows.

First, we have a single and very simple micro-transformation:

$$\begin{aligned} ecommute &\in \text{Exp} \rightarrow \text{Exp} \\ ecommute(e) &:= \text{let } y = b \text{ in let } x = a \text{ in } e_0 \\ &\quad \text{if } e = (\text{let } x = a \text{ in let } y = b \text{ in } e_0) \text{ and } x \notin \text{fv}(b) \\ ecommute(e) &:= \text{let } k_2 = \text{cont } y_2. e_2 \text{ in let } k_2 = \text{cont } y_2. e_2 \text{ in } e_0 \\ &\quad \text{if } e = (\text{let } k_1 = \text{cont } y_1. e_1 \text{ in let } k_2 = \text{cont } y_2. e_2 \text{ in } e_0) \text{ and } k_1 \notin \text{fv}(e_2) \end{aligned}$$

Note that the side condition on x (respectively k_1) is crucial for correctness (we will sketch the proof of correctness in a moment).

Second, we have an analysis α that marks any let-binding that needs to be swapped with its succeeding one. More formally, it places annotations that are elements of 1_\perp and the transformer that consumes them is the lifting β_\perp of the following:

$$\begin{aligned} \beta &\in 1 \rightarrow \text{Exp} \rightarrow \text{Exp} \\ \beta(1)(e) &:= ecommute(e) \end{aligned}$$

Accordingly, the annotation \perp means “do nothing here”, whereas the annotation 1 means “apply *ecommute* here”.

The details of the analysis are not very interesting (and importantly, as far as semantics preservation is concerned, they are irrelevant). The analysis recursively descends into the given function expression, for which the interesting case is, of course, a let-binding. Let us consider a regular let-binding:

$$\text{let } y_1 = a_1 \text{ in } e_1$$

We recursively analyze a_1 and e_1 , yielding a_1 and e_1 . Now we want to know whether $etf^{\beta\perp}(e_1)$ (*i.e.*, the result of transforming e_1), is another let-binding $\text{let } y_2 = a_2 \text{ in } e_2$ and if so, produce a \perp or 1 annotation depending on whether a_2 is “less”² than a_1 .

The full commute pass $\text{commute} \in \mathbf{Mod} \rightarrow \mathbf{Mod}$ is then defined as $\text{commute} := tf_{\alpha}^{\beta\perp}$.

5.4.2 Verification

In order to establish the correctness of $tf_{\alpha}^{\beta\perp}$, it suffices by Theorems 19 and 20 to show the correctness of β . The syntactic aspect of this is straightforward. We sketch the semantic aspect, *i.e.*, the proof of the following:

$$\forall l, \Gamma, e, e', w. \quad \Gamma \vdash e \wedge \text{uniq}(\text{bv}(e), \Gamma) \wedge \text{commute}(e) = e' \wedge w \in \text{stable}(\Omega_{II}) \implies \\ \Gamma \vdash e' \lesssim_w^* e$$

Here we consider the case of let-bindings for term variables (the argument for continuations is analogous). So suppose $e = (\text{let } x = a \text{ in let } y = b \text{ in } e_0)$ with $x \notin \text{fv}(b)$ and $e' = (\text{let } y = b \text{ in let } x = a \text{ in } e_0)$. It suffices to show³ $((\emptyset, (\sigma', e')), (\emptyset, (\sigma, e))) \in \mathbf{E}$, where the environments σ' and σ contain related values for variables in Γ .

If the evaluation of a or that of b in the source fails, then the source produces an error and there is nothing more to show. So let us assume that a evaluates in σ to v_a and that b evaluates in $\sigma, x \mapsto v_a$ to v_b . By the well-formedness of e we know $\Gamma, x \vdash b$. But since $x \notin \text{fv}(b)$, we can strengthen this to $\Gamma \vdash b$, for which we get $\Gamma \vdash b \lesssim_w b$ by Lemma 70. Note that since $x \notin \Gamma$ (by the uniqueness condition), the environments σ' and $\sigma, x \mapsto v_a$ are still related for Γ . Thus we learn that the target’s evaluation of b in σ' evaluates to some v'_b that is related to v_b , extending σ' to $\sigma', y \mapsto v'_b$.

Now we use $\Gamma \vdash a$ to get $\Gamma \vdash a \lesssim_w a$. Much as before, note that since $y \notin \Gamma$, we know that $\sigma', y \mapsto v'_b$ and σ are related for Γ as well. Hence the target program’s evaluation of a in $\sigma', y \mapsto v'_b$ evaluates to some v'_a that is related to v_a . Using Lemmas 43, 47, 48, 44, it remains to show the resulting programs related:

$$((\emptyset, ((\sigma', y \mapsto v'_b, x \mapsto v'_a), e_0)), (\emptyset, ((\sigma, x \mapsto v_a, y \mapsto v_b), e_0))) \in \mathbf{E}$$

It is easy to see that the two environments above are related for Γ, x, y , because σ' and σ are related, v'_a and v_a are related, v'_b and v_b are related, and moreover $x \neq y$ by uniqueness. Therefore we are done by instantiating $\Gamma, x, y \vdash e_0 \lesssim_w e_0$, which we get from $\Gamma, x, y \vdash e_0$.

²In the Coq development, we do not fix an order on expressions but parameterize over a comparison operator. In the extracted Ocaml code, we implement the comparison with the help of a hash table, thus creating the order on the fly.

³Here and in subsequent proof sketches, we often gloss over states, worlds, etc. for the sake of readability.

5.4.3 Alternative Implementation.

There is a somewhat simpler way to implement the commute pass than what we described in Section 5.4.1. The idea is to move the expression comparison into the micro-transformation:

$$\begin{aligned} ecommute' &\in \text{exp} \rightarrow \text{exp} \\ ecommute'(e) &:= \text{let } x_2 = a_2 \text{ in let } x_1 = a_1 \text{ in } e_2 \\ &\quad \text{if } e = (\text{let } x_1 = a_1 \text{ in let } x_2 = a_2 \text{ in } e_2) \text{ and } x \notin \text{fv}(b) \text{ and } a_2 < a_1 \end{aligned}$$

This is of course still correct, but now the analysis can be trivial: mark *every* expression (so we can use 1 instead of 1_\perp).

We prefer the first version because it is more general and might be useful in other contexts.

5.5 The Dedup Pass

5.5.1 Transformation

Deduplication is also implemented using the annotation framework. We use the following micro-transformation:

$$\begin{aligned} ededup &\in \text{Exp} \rightarrow \text{Exp} \\ ededup(e) &:= \text{let } x = a \text{ in } e_0[x/y] \\ &\quad \text{if } e = (\text{let } x = a \text{ in let } y = a \text{ in } e_0) \\ ededup(e) &:= \text{let } k_1 = \text{cont } y. e_1 \text{ in } e_0[k_1/k_2] \\ &\quad \text{if } e = (\text{let } k_1 = \text{cont } y. e_1 \text{ in let } k_2 = \text{cont } y. e_1 \text{ in } e_0) \end{aligned}$$

As annotations we choose the trivial set (1) and, accordingly, the trivial transformer β and the trivial analysis α . The complete deduplication pass then is defined as $dedup := tf_\alpha^\beta$, *i.e.*, we apply the micro-transformation wherever possible.

5.5.2 Verification

In order to establish the correctness of tf_α^β , it suffices by Theorem 19 to show the correctness of β . The syntactic aspect of this is straightforward. We sketch the semantic aspect, *i.e.*, the proof of the following:

$$\forall l, \Gamma, e, e', w. \quad \Gamma \vdash e \wedge \text{uniq}(\text{bv}(e), \Gamma) \wedge ededup(e) = e' \wedge w \in \text{stable}(\Omega_{\mathcal{IT}}) \implies \Gamma \vdash e' \lesssim_w^* e$$

Here we consider the case of let-bindings for term variables (the argument for continuations is analogous). So suppose $e = (\text{let } x = a \text{ in let } y = a \text{ in } e_0)$ and $e' = (\text{let } x = a \text{ in } e_0[x/y])$. It suffices to show $((\emptyset, (\sigma', e')), (\emptyset, (\sigma, e))) \in \mathbf{E}$, where the environments σ' and σ contain related values for variables in Γ .

If the first or second evaluation of a in the source fails, then the source produces an error and there is nothing more to show. So let us assume that a first evaluates in σ to v_1 and then in $\sigma, x \mapsto v_1$ to v_2 .

From $\Gamma \vdash a$ we get $\Gamma \vdash a \lesssim_w a$ by Lemma 70. Since σ' and σ are related for Γ , we thus know that the target evaluation of a in σ' results in a value v' related to v_1 . But note that σ' and $\sigma, x \mapsto v_1$ are also related for Γ because $x \notin \Gamma$ by uniqueness, hence v' is also related to v_2 (evaluation is deterministic). Using Lemmas 43, 47, 48, 44, it remains to show the resulting programs related:

$$((\emptyset, ((\sigma', x \mapsto v'), e_0[x/y])), (\emptyset, ((\sigma, x \mapsto v_1, y \mapsto v_2), e_0))) \in \mathbf{E}$$

Now we apply Lemma 69 to $\Gamma, x, y \vdash e_0$ (not Lemma 70), which yields $\Gamma, x, y \vdash e_0[x/y] \lesssim_w^{[x/y]} e_0$. To show the goal, it thus suffices to show the environments above related for Γ, x, y relative to the renaming $[x/y]$. This holds essentially because v' is related to both v_1 and v_2 .

5.6 The Hoist Pass

5.6.1 Transformation

Hoisting is also implemented using the annotation framework. We define the complete pass as tf_α^β , where—like for deduplication— α is the trivial analysis for 1 and β the trivial transformer $\lambda l. \text{hoist}$. The underlying micro-transformation is defined as follows:

$$\begin{aligned} \text{ehoist} &\in \text{Exp} \rightarrow \text{Exp} \\ \text{ehoist}(e) &:= \text{let } y_2 = a \text{ in let } y_1 = (\text{fix } f(y_3, k). e_1) \text{ in } e_2 \\ &\quad \text{if } e = (\text{let } y_1 = (\text{fix } f(y_3, k). \text{let } y_2 = a \text{ in } e_1) \text{ in } e_2) \text{ and } f, y_3, k \notin \text{fv}(a) \\ &\quad \text{and } a \notin \{x_1 \odot x_2, x_1 == x_2, x.1, x.2\} \\ \text{ehoist}(e) &:= \text{let } y_2 = a \text{ in let } y_1 = (\Lambda k. e_1) \text{ in } e_2 \\ &\quad \text{if } e = (\text{let } y_1 = (\Lambda k. \text{let } y_2 = a \text{ in } e_1) \text{ in } e_2) \text{ and } k \notin \text{fv}(a) \\ &\quad \text{and } a \notin \{x_1 \odot x_2, x_1 == x_2, x.1, x.2\} \end{aligned}$$

The restriction of a to certain forms is a syntactic criterion ensuring that the evaluation of a will succeed (yield a value). It is easy to see that *ehoist* would not be correct in general if the evaluation of a could fail.

Here we do not bother to hoist out of continuations because continuations are only executed at most once anyways.

Note that we cannot not hoist continuation definitions (only regular let-bindings) out of functions. This is because the resulting program would generally be ill-formed due to functions not being able to access continuations from a surrounding scope (Section 4.2.4.2).

5.6.2 Verification

In order to establish the correctness of tf_α^β , it suffices by Theorem 19 to show the correctness of β . The syntactic aspect of this is straightforward. We sketch the semantic aspect, *i.e.*, the proof of the following:

$$\forall l, \Gamma, e, e', w. \quad \Gamma \vdash e \wedge \text{uniq}(\text{bv}(e), \Gamma) \wedge \text{ehoist}(e) = e' \wedge w \in \text{stable}(\Omega_{II}) \implies \Gamma \vdash e' \lesssim_w^* e$$

Here we consider the second case in the definition of *ehoist*, *i.e.*, we hoist out of a generalization (the first case is analogous). So suppose that:

$$\begin{aligned} e &= (\text{let } y_1 = (\Lambda k. \text{let } y_2 = a \text{ in } e_1) \text{ in } e_2) \\ e' &= (\text{let } y_2 = a \text{ in let } y_1 = (\Lambda k. e_1) \text{ in } e_2) \end{aligned}$$

It suffices to show $((\emptyset, (\sigma', e')), (\emptyset, (\sigma, e))) \in \mathbf{E}$ for environments σ' and σ related relative to Γ .

Regarding the source program e , we know that $\Lambda k. \text{let } y_2 = a \text{ in } e_1$ evaluates in σ to $v_g := \langle \sigma; \Lambda k. \text{let } y_2 = a \text{ in } e_1 \rangle$. Regarding the target program e' , we know that a evaluates in σ' to some v' thanks to the syntactic restriction of a . Afterwards, $\Lambda k. e_1$ evaluates to $v'_g := \langle \sigma', y_2 \mapsto v'; \Lambda k. e_1 \rangle$. Using Lemmas 43, 47, 48, 44, it remains to show the resulting programs related:

$$((\emptyset, ((\sigma', y_2 \mapsto v', y_1 \mapsto v'_g), e_2)), (\emptyset, ((\sigma, y_1 \mapsto v_g), e_2))) \in \mathbf{E}$$

From $\Gamma, y_1 \vdash e_2$ we get $\Gamma, y_1 \vdash e_2 \lesssim_w e_2$ by Lemma 70. To show the goal, it thus suffices to show the environments above related for Γ, y_1 . This boils down to showing the closures v'_g and v_g related.

Considering applications of those to continuations v'_k and v_k , respectively, we must show:

$$((\emptyset, ((\sigma', y_2 \mapsto v', k \mapsto v'_k), e_1)), (\emptyset, ((\sigma, k \mapsto v_k), \text{let } y_2 = a \text{ in } e_1))) \in \mathbf{E}(v'_k, v_k)$$

By Lemmas 43, 47, 44 this reduces to

$$((\emptyset, ((\sigma', y_2 \mapsto v', k \mapsto v'_k), e_1)), (\emptyset, ((\sigma, k \mapsto v_k, y_2 \mapsto v), e_1))) \in \mathbf{E}(v'_k, v_k)$$

where v is the result of evaluating a in $\sigma, k \mapsto v_k$.

From $\Gamma^\dagger, k, y_2 \vdash e_1$ we get $\Gamma^\dagger, k, y_2 \vdash e_1 \lesssim_w e_1$ by Lemma 70. To show the goal, it thus suffices to show the above environments related for Γ^\dagger, k, y_2 and initial continuations v'_k and v_k . Relying on uniqueness, this means showing the following:

1. σ' and σ are related environments for Γ^\dagger with initial continuations v'_k, v_k .
2. v' and v are related values.
3. v'_k and v_k are related continuations relative to themselves.

Goal (1) follows from their relatedness for Γ because Γ^\dagger does not contain any continuations. Goal (3) holds by Lemma 52. It remains to show (2).

From $\Gamma, k \vdash a$ and $k \notin \text{fv}(a)$ we know $\Gamma \vdash a$ and thus get $\Gamma \vdash a \lesssim_w a$ by Lemma 70. Since σ' is related to $\sigma, k \mapsto v_k$ for Γ ($k \notin \Gamma$ by uniqueness), we know that v' , the result of evaluating a in σ' , is related to v .

5.7 The Dead Code Elimination Pass

Pilsner's dead code elimination applies to let-bindings, read operations, and allocations. It removes these operations if the variable that they introduce does not occur in the rest of the program. This is safe because these operations have no observable side effects (even though some of them concern memory).

5.7.1 Transformation

Dead code elimination (DCE) is implemented directly, not using the annotation framework. This is because the framework was developed only after dead code elimination was already finished. In its current form, the framework cannot be used to implement DCE because proving correct the elimination of allocations requires a custom local world (cf. the comment in Section 5.3.2.2).

Figure 5.7 shows the definition of the module-level transformation dce in terms of a recursive helper function $mdce$ that applies $adce$ to each top-level function. $adce$ in turn is defined in Figure 5.8 over the structure of pure expressions, in mutual recursion with $edce$ for control expressions. The interesting cases are the first four of $edce$ (those in Figure 5.7), all others are merely structural.

$edce$ determines whether to perform an elimination by checking if the introduced variable occurs freely in the rest of the program. In order for the transformation to be idempotent, it is crucial, however, that we first transform the rest of the program, *i.e.*, that we do the occurrence check on the free variables of the *transformed* rest of the program (which may be fewer). For efficiency, $adce$ and $edce$ additionally return (an over-approximation of) this set, so that it doesn't need to be recomputed all the time: if $(e', Z) = edce(e)$, then $\text{fv}(e') \subseteq Z$ (and similarly for $adce$).

5.7.2 Verification

We use the following local world, whose states are heaps on the source side.

$w.T.S$	$:=$	Heap
$w.T.\sqsubseteq$	$:=$	Heap \times Heap
$w.T.\sqsubseteq_{\text{pub}}$	$:=$	Heap \times Heap
$w.C(G)(s_g, h)$	$:=$	$\{ (\emptyset, (h, \emptyset)) \}$
$w.O$	$:=$	gwf

The states are only relevant for the elimination of allocations. The idea is very simple: we need to consider any allocation in the source program that got eliminated

$$\begin{aligned}
& dce \in \mathbf{Mod} \rightarrow \mathbf{Mod} \\
& dce(\epsilon) = \epsilon \\
& dce((F, a), M) = \text{let } (a', -) := \text{adce}(a) \text{ in } (F, a'), dce(M) \\
\\
& edce \in \mathbf{Exp} \rightarrow \mathbf{Exp} \times \mathbf{Var} \\
\\
& edce(\text{let } y = a \text{ in } e) = \\
& \quad \text{let } (e', Z_1) := edce(e) \text{ in} \\
& \quad \text{if } y \in Z_1 \text{ then} \\
& \quad \quad \text{let } (a', Z_2) := \text{adce}(a) \text{ in} \\
& \quad \quad (\text{let } y = a' \text{ in } e', Z_1 \cup Z_2) \\
& \quad \text{else } (e', Z_1) \\
\\
& edce(\text{let } k = \text{cont } y. e_1 \text{ in } e_2) = \\
& \quad \text{let } (e'_2, Z_1) := edce(e_2) \text{ in} \\
& \quad \text{if } k \in Z_1 \text{ then} \\
& \quad \quad \text{let } (e'_1, Z_2) := edce(e_1) \text{ in} \\
& \quad \quad (\text{let } k = \text{cont } y. e'_1 \text{ in } e'_2, Z_1 \cup Z_2) \\
& \quad \text{else } (e'_2, Z_1) \\
\\
& edce(y \leftarrow \text{ref } x; e) = \\
& \quad \text{let } (e', Z) := edce(e) \text{ in} \\
& \quad \text{if } y \in Z \text{ then} \\
& \quad \quad (y \leftarrow \text{ref } x; e', \{x\} \cup Z) \\
& \quad \text{else } (e', Z) \\
\\
& edce(y \leftarrow !x; e) = \\
& \quad \text{let } (e', Z) := edce(e) \text{ in} \\
& \quad \text{if } y \in Z \text{ then} \\
& \quad \quad (y \leftarrow !x; e', \{x\} \cup Z) \\
& \quad \text{else } (e', Z)
\end{aligned}$$

Figure 5.7: DCE transformation (part 1 of 2).

$$\begin{aligned}
edce(y \leftarrow \text{input}; e) &= \\
&\quad \text{let } (e', Z) := edce(e) \text{ in} \\
&\quad (y \leftarrow \text{input}; e', Z) \\
\\
edce(\text{output } x; e) &= \\
&\quad \text{let } (e', Z) := edce(e) \text{ in} \\
&\quad (\text{output } x; e', \{x\} \cup Z) \\
\\
edce(\text{ifnz } x \text{ then } e_1 \text{ else } e_2) &= \\
&\quad \text{let } (e'_1, Z_1) := edce(e_1) \text{ in} \\
&\quad \text{let } (e'_2, Z_2) := edce(e_2) \text{ in} \\
&\quad (\text{ifnz } x \text{ then } e'_1 \text{ else } e'_2, \{x\} \cup Z_1 \cup Z_2) \\
\\
edce(\text{case } x \text{ (y. } e_1) \text{ (y. } e_2)) &= \\
&\quad \text{let } (e'_1, Z_1) := edce(e_1) \text{ in} \\
&\quad \text{let } (e'_2, Z_2) := edce(e_2) \text{ in} \\
&\quad (\text{case } x \text{ (y. } e'_1) \text{ (y. } e'_2), \{x\} \cup Z_1 \cup Z_2) \\
\\
edce(x_1 \ x_2 \ k) &= \\
&\quad (x_1 \ x_2 \ k, \{x_1, x_2, k\}) \\
\\
edce(x \ [] \ k) &= \\
&\quad (x \ [] \ k, \{x, k\}) \\
\\
edce(k \ x) &= \\
&\quad (k \ x, \{k, x\}) \\
\\
edce(x_1 := x_2; e) &= \\
&\quad \text{let } (e', Z) := edce(e) \text{ in} \\
&\quad (x_1 := x_2; e', \{x_1, x_2\} \cup Z) \\
\\
adce \in \text{BExp} \rightarrow \text{BExp} \times \text{Var} \\
adce(\langle \rangle) &= (\langle \rangle, \emptyset) \\
adce(n) &= (n, \emptyset) \\
adce(x.1) &= (x.1, \{x\}) \\
adce(x.2) &= (x.2, \{x\}) \\
adce(\text{inl } x) &= (\text{inl } x, \{x\}) \\
adce(\text{inr } x) &= (\text{inr } x, \{x\}) \\
adce(\langle x_1, x_2 \rangle) &= (\langle x_1, x_2 \rangle, \{x_1, x_2\}) \\
adce(x_1 \odot x_2) &= (x_1 \odot x_2, \{x_1, x_2\}) \\
adce(x_1 == x_2) &= (x_1 == x_2, \{x_1, x_2\}) \\
adce(\text{fix } f(y, k). e) &= \text{let } (e', Z) := edce(e) \text{ in } (\text{fix } f(y, k). e', Z) \\
adce(\Lambda k. e) &= \text{let } (e', Z) := edce(e) \text{ in } (\Lambda k. e', Z)
\end{aligned}$$

Figure 5.8: DCE transformation (part 2 of 2).

in the target program as being local, because otherwise we would be forced to provide a target location corresponding to the freshly allocated source location. For convenience we allow arbitrary transitions between the states—this is okay precisely because the heap locations in these states won't ever be used by the programs.

The correctness proof for *dce* boils down to showing

$$\Gamma \vdash e' \wedge \text{uniq}(\text{bv}(e), \Gamma) \wedge (e', _) = \text{edce}(e) \implies \Gamma \vdash e' \lesssim_w e$$

(and similarly for *adce*). We do this by induction on the structure of the input program. Note that we assume the well-formedness of e' in Γ , not that of e in Γ . This is somewhat unusual but important for the induction to go through, as we will see in a moment.

Let us sketch the allocation case here. The other interesting cases are analogous but a little simpler because they do not involve any important state transitions. The purely structural cases are trivial to handle using the compatibility lemmas.

Suppose $e = (y \leftarrow \text{ref } x; e_0)$. We must show $\Gamma \vdash e' \lesssim_w e$. Let $(e'_0, Z_0) := \text{edce}(e_0)$. If the allocation is not eliminated ($y \in Z_0$), the goal equals

$$\Gamma \vdash (y \leftarrow \text{ref } x; e'_0) \lesssim_w (y \leftarrow \text{ref } x; e_0)$$

and we are done using the inductive hypothesis and the compatibility lemma for allocation.

Now suppose $y \notin Z_0$ such that the elimination is performed. Then the goal equals $\Gamma \vdash e'_0 \lesssim_w (y \leftarrow \text{ref } x; e_0)$. After unfolding the definitions and using Lemma 43, we must show

$$((h', (\sigma', e'_0)), (h, (\sigma, (y \leftarrow \text{ref } x; e_0)))) \in \mathbf{E}(s_g, s)$$

where the environments σ' and σ are related for Γ and the heaps consist of a related global part, a related local part, and a frame part. By construction of w , the local part of h' is empty and that of h equals the local state s .

We now use Lemma 47 to take a step on the source side so that it remains to show

$$((h', (\sigma', e'_0)), (h \sqcup [l \mapsto \sigma(x)], ((\sigma, y \mapsto l), e_0))) \in \mathbf{E}(s_g, s)$$

(for some fresh location l).

Now we use Lemma 50 to advance the world's state (s_g, s) to $(s_g, s \sqcup [l \mapsto \sigma(x)])$ so that it matches the new heap. That is, we interpret the change to h as a change to its local part (which happens to be s) and pick a new local state accordingly. This enables us to use Lemma 44 such that the goal becomes

$$((\emptyset, (\sigma', e'_0)), (\emptyset, ((\sigma, y \mapsto l), e_0))) \in \mathbf{E}(s_g, s \sqcup [l \mapsto \sigma(x)]).$$

Since $\Gamma, y \vdash e'_0$ and $y \notin Z_0 \supseteq \text{fv}(e'_0)$, we know $\Gamma \vdash e'_0$. The inductive hypothesis then yields $\Gamma \vdash e'_0 \lesssim_w e_0$. To instantiate this and establish the goal, we must show that σ' is related to $\sigma, y \mapsto l$ for Γ , which follows from relatedness of σ' to σ and uniqueness ($y \notin \Gamma$). We are done.

What if we had started out with $\Gamma \vdash e$ instead of $\Gamma \vdash e'$? Then we know $\Gamma, y \vdash e_0$ but we cannot not derive $\Gamma \vdash e_0$ because we only know that y is not free in e'_0 —it may very well be free in e_0 (inside dead code). So instead of $\Gamma \vdash e'_0 \lesssim_w e_0$, we would only get the weaker $\Gamma, y \vdash e'_0 \lesssim_w e_0$, which we would not be able to instantiate because y does not exist in the target program.

5.8 The Inline Pass

5.8.1 Transformation

Pilsner’s inlining pass is defined in Figure 5.9. The auxiliary functions *inline* and *ainline* operate on control expressions and basic expressions, respectively. They take a mapping ψ containing the functions that are to be inlined (mapping their identifiers to their code). The more interesting one is *inline*: when it sees a call to one of the functions in ψ , it performs the actual inlining by replacing the call with the function’s body and substituting the respective variables (note that this works even for recursive functions). Note that inlining is intentionally not an idempotent transformation.

minline lifts these operations to modules in a straightforward way. In the final transformation, *inline*, we decide for simplicity to just inline all top-level functions of a module by picking $\psi := M$. This is of course somewhat unrealistic but does not reflect a real limitation: the “analysis” that computes ψ can be made arbitrarily complex without affecting the correctness of the transformation—similar to the analyses in the previous transformations. Also, an extension of the transformation and its verification to local functions (functions not defined at the top-level) is expected to be straightforward.

5.8.2 Verification

Unfortunately, the tools from the previous sections do not suffice for verifying the inlining pass. Consider the correctness proof of *inline*. Along the lines of the previous verifications, we would have to show something like

$$\forall e, \psi. \Gamma \vdash e \implies \Gamma \vdash \text{inline}(\psi)(e) \lesssim_w^* e$$

But this is clearly not possible. Suppose e gets executed in an environment σ and $\text{inline}(\psi)(e)$ gets executed in a related environment σ' . If *inline* inlines a function in e , it does so by replacing functions calls according to ψ . However, there is no connection between this inlining information ψ and the environment σ in which the original program will lookup the function.

To address this problem, we slightly generalize the \lesssim relations by considering only environments that match given *descriptions* ψ_a and ψ_b , respectively (of the same type as ψ above).

$$\begin{aligned}
& \text{inline} \in \mathbf{Mod} \rightarrow \mathbf{Mod} \\
& \text{inline}(M) = \text{freshen}(\text{minline}(M))(M) \\
\\
& \text{minline} \in (\mathbf{Var} \rightarrow \mathbf{BExp}) \rightarrow \mathbf{Mod} \rightarrow \mathbf{Mod} \\
& \text{minline}(\psi)(\epsilon) = \epsilon \\
& \text{minline}(\psi)((F, a), M) = (F, \text{ainline}(\psi)(a), \text{minline}(\psi)(M)) \\
\\
& \text{einline} \in (\mathbf{Var} \rightarrow \mathbf{BExp}) \rightarrow \mathbf{Exp} \rightarrow \mathbf{Exp} \\
& \text{einline}(\psi)(x_1 \ x_2 \ k) := \begin{cases} e[x_1/f][x_2/y][k/k'] & \text{if } \psi(x_1) = \text{fix } f(y, k'). e \\ x_1 \ x_2 \ k & \text{otherwise} \end{cases} \\
& \text{einline}(\psi)(x \ \square \ k) := \begin{cases} e[k/k'] & \text{if } \psi(x) = \Lambda k'. e \\ x \ \square \ k & \text{otherwise} \end{cases} \\
& \text{einline}(\psi)(\text{let } y = a \text{ in } e) := \text{let } y = \text{ainline}(\psi)(a) \text{ in } \text{einline}(\psi)(e) \\
& \dots \\
\\
& \text{ainline} \in (\mathbf{Var} \rightarrow \mathbf{BExp}) \rightarrow \mathbf{BExp} \rightarrow \mathbf{BExp} \\
& \text{ainline}(\psi)(\text{fix } f(y, k). e) := \text{fix } f(y, k). \text{einline}(\psi)(e) \\
& \dots
\end{aligned}$$

Figure 5.9: Inlining of top-level functions.

Definition 47.

$$\begin{aligned}
& \Gamma; \psi_a; \psi_b \vdash e_a \lesssim_w^\phi e_b := \\
& \quad \exists i. \forall G, s_0, s, \sigma_a, \sigma_b, j, \mathbf{k}_a, \mathbf{k}_b. \\
& \quad \quad G \in \mathbf{GK}_{w\uparrow} \wedge s \sqsupseteq_{\text{pub}} s_0 \wedge \text{respects}(\sigma_a, \psi_a, \Gamma) \wedge \text{respects}(\sigma_b, \psi_b, \Gamma) \wedge \\
& \quad \quad (\forall x \in \Gamma. \exists \mathbf{v}_a, \mathbf{v}_b. \sigma_a(\phi(x)) = \mathbf{v}_a \wedge \sigma_b(x) = \mathbf{v}_b \wedge (\mathbf{v}_a, \mathbf{v}_b) \in \overline{G}(s)) \wedge \\
& \quad \quad (\forall k \in \Gamma. \exists \mathbf{k}'_a, \mathbf{k}'_b. \sigma_a(\phi(k)) = \mathbf{k}'_a \wedge \sigma_b(k) = \mathbf{k}'_b \wedge \\
& \quad \quad \quad (\mathbf{k}'_a, \mathbf{k}'_b) \in \mathbf{K}_{w\uparrow}(j)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)) \\
& \quad \implies ((\emptyset, (\sigma_a, e_a)), (\emptyset, (\sigma_b, e_b))) \in \mathbf{E}_{w\uparrow}(i+j)(\mathbf{k}_a, \mathbf{k}_b)(G)(s_0)(s)(\perp)
\end{aligned}$$

Definition 48.

$$\begin{aligned}
& \Gamma; \psi_a; \psi_b \vdash a_a \lesssim_w^\phi a_b := \\
& \quad \forall G, s, \sigma_a, \sigma_b, \mathbf{v}_b. \\
& \quad \quad G \in \mathbf{GK}_{w\uparrow} \wedge \text{respects}(\sigma_a, \psi_a, \Gamma) \wedge \text{respects}(\sigma_b, \psi_b, \Gamma) \wedge \\
& \quad \quad (\forall x \in \Gamma. \exists \mathbf{v}'_a, \mathbf{v}'_b. \sigma_a(\phi(x)) = \mathbf{v}'_a \wedge \sigma_b(x) = \mathbf{v}'_b \wedge (\mathbf{v}'_a, \mathbf{v}'_b) \in \overline{G}(s)) \wedge \\
& \quad \quad \llbracket a_b \rrbracket_{\sigma_b} = \mathbf{v}_b \\
& \quad \implies \exists \mathbf{v}_a. \llbracket a_a \rrbracket_{\sigma_a} = \mathbf{v}_a \wedge (\mathbf{v}_a, \mathbf{v}_b) \in \overline{G}(s)
\end{aligned}$$

The differences to the earlier definitions are highlighted. The ψ_a and ψ_b components can be thought of as static descriptions of (parts of) the runtime environments. Naturally, we obtain the original relations when plugging in empty ψ_a and ψ_b . For

the concrete purpose of verifying inlining, ψ_a will be empty but ψ_b will be the function mapping used by *inline* and *ainline*. We will see a use case for a non-empty ψ_a in Section 5.9.

A runtime environment σ *respects* such a description, written $\text{respects}(\sigma, \psi, \Gamma)$, if it contains matching closures:

Definition 49.

$$\begin{aligned} \text{respects}(\sigma, \psi, \Gamma) := & \forall z, e. \psi(z) = e \implies \exists \sigma'. \\ & z \in \Gamma \wedge \sigma(z) = \langle \sigma'; e \rangle \wedge \\ & \forall z' \in \text{defBefore}(z, \Gamma). \sigma'(z') = \sigma(z') \end{aligned}$$

This says that whenever ψ maps a variable z to an expression e , the actual runtime environment σ maps z to a closure value whose code component is e . Moreover, the environment σ' stored in that closure must agree with σ on all variables *defined before* z , *i.e.*, those variables that syntactically occur to the left of z in Γ . (In our concrete case of inlining, σ will always be an extension of any such σ' .)

It is easy to see that the new relation generalizes the old one:

Lemma 73.

$$\frac{\Gamma \vdash e_a \lesssim_w^\phi e_b}{\Gamma; \psi_a; \psi_b \vdash e_a \lesssim_w^\phi e_b} \quad \frac{\Gamma; \emptyset; \emptyset \vdash e_a \lesssim_w^\phi e_b}{\Gamma \vdash e_a \lesssim_w^\phi e_b}$$

With its help we can establish the following correctness statement about the function call case of *inline* (and similarly for $x \parallel k$).

Lemma 74.

$$\frac{\begin{array}{c} x_1, x_2, k \in \Gamma \\ \psi(x_1) = \text{fix } f(y, k'). e \\ \{x_1, x_2, k\} \cap \text{bv}(e) = \emptyset \\ \text{defBefore}(x_1, \Gamma) \vdash \text{fix } f(y, k'). e \\ \text{stable}(w) \end{array}}{\Gamma; \epsilon; \psi \vdash \text{inline}(\psi)(x_1 \ x_2 \ k) \lesssim_w^* x_1 \ x_2 \ k}$$

Observe that the inlining information ψ used to transform the call is now connected to whatever runtime environment will be given for the execution of $x_1 \ x_2 \ k$. Also note that we require the well-formedness of the function not just in Γ but in the restriction $\text{defBefore}(x_1, \Gamma)$. Intuitively, this is the environment in which the function was defined, and using it here ensures compatibility with whatever runtime environment the closure of f will contain (in the source program execution). The proof sketch goes as follows.

Proof. After unfolding the definitions, the goal is to show $(\sigma_a, e[x_1/f][x_2/y][k/k'])$ related to $(\sigma_b, x_1 \ x_2 \ k)$ by **E** (ignoring some uninteresting details here). Since $\text{respects}(\sigma_b, \psi, \Gamma)$, we know that $\sigma_b(x_1) = \langle \sigma'_b; \text{fix } f(y, k'). e \rangle$ for some closure environment σ'_b that agrees with σ_b on $\text{defBefore}(x_1, \Gamma)$.

Using Lemmas 43, 47, 44, we can take a step on the source side and thus change $(\sigma_b, x_1 \ x_2 \ k)$ to $(\sigma'_b[f, y, k' \mapsto \sigma_b(x_1), \sigma_b(x_2), \sigma_b(k)], e)$ in the goal.

Now we make use of the original (stronger) \lesssim relation: we apply Lemma 69 to $\text{defBefore}(x_1, \Gamma)^\dagger, f, y, k' \vdash e$, yielding

$$\text{defBefore}(x_1, \Gamma)^\dagger, f, y, k' \vdash \phi(e) \lesssim_w^\phi e$$

for any renaming ϕ .

Note that this immediately leads to the goal if we pick $\phi := [x_1/f][x_2/y][k/k']$ and then show the environments related modulo ϕ . That is, it remains to show the relatedness of $\sigma_a(\phi(z))$ and $(\sigma'_b[f, y, k' \mapsto \sigma_b(x_1), \sigma_b(x_2), \sigma_b(k)])(z)$ for any $z \in \text{defBefore}(x_1, \Gamma)^\dagger, f, y, k'$.

For $z = k'$, this reduces to the relatedness of $\sigma_a(k)$ and $\sigma_b(k)$, which holds due to the assumed relatedness of σ_a and σ_b for Γ . Similarly for $z = y$ and $z = f$. For z in $\text{defBefore}(x_1, \Gamma)^\dagger$, we must show $\sigma_a(z)$ related to $\sigma'_b(z)$. Fortunately we know from $\text{respects}(\sigma_b, \psi, \Gamma)$ that $\sigma'_b(z) = \sigma_b(z)$ for such z , hence we can use the same argument here as well and are done. \square

The correctness statements for the purely structural cases (including *ainline*) are simpler, and easy to obtain with the help of compatibility lemmas for the generalized \lesssim relations. Unfortunately, these lemmas are currently not generalizations of those for the original \lesssim relations (it wasn't obvious how to achieve this due to technical difficulties related to variables). Nevertheless, they look fairly similar and are omitted here. Suffice it to say that for simplicity none of them modifies the ψ_a and ψ_b components. For the purpose of verifying inlining of *local* functions, though, it will most likely be useful to state the lemma for let-bindings of functions such that in the premise the mappings are extended with the newly defined function.

Combining these correctness lemmas with a simple induction leads to the following familiar-looking lemma (omitting some technical side conditions concerning variable freshness and uniqueness of e and the functions inside ψ).

Lemma 75.

$$\frac{\Gamma \vdash e \quad \text{stable}(w) \quad \dots}{\Gamma; \epsilon; \psi \vdash \text{einline}(\psi)(e) \lesssim_w^* e}$$

Based on this, we can now verify the module-level transformation *inline*. Naturally this is the point where we must demonstrate that the runtime environments that arise from loading a module actually respect the module, so that we can discharge the respective assumptions hidden inside the conclusion of Lemma 75. All this is not very hard but tedious. In the end, we arrive at the following correctness property.

Lemma 76 (Correctness of Inlining).

$$\frac{\Gamma \vdash M : \Gamma' \quad \text{uniqmod}(M)}{\Gamma \vdash \text{inline}(M) : \Gamma' \quad \Gamma \vdash \text{inline}(M) \lesssim_{\mathcal{IL}}^* M : \Gamma'}$$

Note that we cannot prove $uniqmod(M')$ because it usually doesn't hold—inlining can obviously destroy the uniqueness property. For this reason, the inlining pass is immediately followed by a very simple *freshening* pass whose sole purpose is to restore this property.

5.8.3 Freshening

The freshening pass renames all bound variables of a module in order to (re-)establish the variable uniqueness property. It is run after inlining and, as we will see in the next section, after contification as well. As shown in Figure 5.10, freshening is a very simple transformation and easy to verify thanks to \lesssim 's support of renaming, hence we keep its presentation very short.

The functions *efreshen* and *afreshen* take a renaming (mapping from variables to variables) and an index indicating the next unused variable. The result is the transformed expression and an updated index. *freshen* lifts these functions to the module level, passing through the current index, which is initially 0.

Regarding semantic correctness, the key lemma is the following (and similarly for *afreshen*). This follows easily with the help of the compatibility lemmas from Section 5.3.

$$\frac{\Gamma \vdash e \quad (e', n') = \text{efreshen}(\phi)(n)(e) \quad (\forall z \in \Gamma. \phi(z) < n) \quad \text{stable}(w)}{\Gamma \vdash e' \lesssim_w e}$$

On top of that, we can derive *freshen*'s semantics preservation. The overall correctness statement includes preservation of well-formedness and—the reason for defining freshening in the first place—enforcement of the uniqueness property that other passes rely on.

$$\frac{\Gamma \vdash M : \Gamma' \quad M' = \text{freshen}(M)}{\Gamma \vdash M' : \Gamma' \quad uniqmod(M') \quad \Gamma \vdash M' \lesssim_{II}^* M : \Gamma'}$$

5.9 The Contify Pass

Pilsner includes a (very simplistic) contification pass. Contification [24] is an optimization that turns a function into a continuation when the function is only ever called with the same continuation argument. This makes control flow more explicit, thus potentially enabling subsequent optimizations. In Pilsner, it has the additional benefit that continuations don't need to be heap-allocated.

5.9.1 Transformation

Contification in Pilsner uses a slightly generalized version of the expression transformation framework from Section 5.3.2, *i.e.*, it first runs an untrusted analysis that

$$\begin{aligned}
& \text{freshen} \in \mathbf{Mod} \rightarrow \mathbf{Mod} \\
& \text{freshen}(M) = \text{mfreshen}(0)(M) \\
\\
& \text{mfreshen} \in \mathbb{N} \rightarrow \mathbf{Mod} \rightarrow \mathbf{Mod} \\
& \text{mfreshen}(n)(\epsilon) = \epsilon \\
& \text{mfreshen}(n)((F, a), M) = \\
& \quad \text{let } (a', n') := \text{afreshen}(\text{id})(n) \text{ in} \\
& \quad (F, a'), \text{mfreshen}(n')(M) \\
\\
& \text{afreshen} \in (\mathbf{Var} \rightarrow \mathbf{Var}) \rightarrow \mathbb{N} \rightarrow \mathbf{BExp} \rightarrow \mathbf{BExp} \times \mathbb{N} \\
& \text{afreshen}(\phi)(n)(\Lambda k. e) = \\
& \quad \text{let } (e', n') := \text{efreshen}(\phi[\text{kvar}(n)/k])(n+1)(e) \text{ in} \\
& \quad (\Lambda \text{kvar}(n). e', n') \\
& \dots \\
\\
& \text{efreshen} \in (\mathbf{Var} \rightarrow \mathbf{Var}) \rightarrow \mathbb{N} \rightarrow \mathbf{Exp} \rightarrow \mathbf{Exp} \times \mathbb{N} \\
& \text{efreshen}(\phi)(n)(y \leftarrow !x; e) = \\
& \quad \text{let } (e', n') := \text{efreshen}(\phi[\text{tvar}(n)/y])(n+1)(e) \text{ in} \\
& \quad (\text{tvar}(n) \leftarrow !\phi(x); e', n') \\
& \dots
\end{aligned}$$

Figure 5.10: Freshening.

annotates places in the module where the expression-level contification should happen. This saves a lot of work because only the expression-level contification needs to be verified.

Expression-level contification consists of two steps. Consider a non-recursive function binding

$$\text{let } y = (\text{fix } _ (y_1, k). e') \text{ in } e$$

that is contifiable, *i.e.*, where all invocations of y in e are done using the same continuation k' . We first extend the binding with the definition of a new continuation, namely the contification of y (w.r.t. to k'):

$$\begin{aligned}
& \text{let } y = (\text{fix } _ (y_1, k). e') \text{ in} \\
& \text{let } k_y = \text{cont } y_2. e'[y_2/y_1][k'/k]e \text{ in}
\end{aligned}$$

In the second step, all invocations of y in e are turned into calls of the newly added continuation k_y (*e.g.*, $y \ x \ k' \rightsquigarrow k_y \ x$). Note that contification does not purge the definition of y , but leaves this to the dead code elimination pass. We exclude recursive functions from contification because our language does not support recursive continuations.

Our analysis is fairly simple. Whenever it finds a function binding that is contifiable (according to the conservative syntactic criterion above), it annotates the

binding as follows:

$$\alpha(\text{let } y = (\text{fix } f(y', k). e') \text{ in } e) := [\text{Subst}_{k'}, \text{Add}_{k'}] \text{let } y = ([\epsilon] \text{fix } f(y', k). \alpha(e')) \text{ in } \alpha(e) \\ (f \notin \text{fv}(e') \wedge \text{contifiable}(y, e) = k')$$

Here, the annotation $\text{Add}_{k'}$ corresponds to the operation of adding a new contification of the function with fixed continuation argument k' (step 1), and $\text{Subst}_{k'}$ corresponds to the operation of replacing all calls of y that use k' with corresponding continuation calls (step 2). By having k' be part of both annotations, the implementations of the respective operations do not need to re-inspect the code in order to find out k' . The analysis uses the list lifting of annotations from Section 5.3.2, which performs transformations from right to left.

Non-contifiable functions as well as other expression forms receive an empty list annotation (if annotatable):

$$\alpha(\text{let } y = (\text{fix } f(y', k). e') \text{ in } e) := \\ [\epsilon] \text{let } y = ([\epsilon] \text{fix } f(y', k). \alpha(e')) \text{ in } \alpha(e) \quad (f \in \text{fv}(e') \vee \text{contifiable}(y, e) = \perp) \\ \alpha(y \leftarrow !x; e) := \\ y \leftarrow !x; \alpha(e) \\ \dots$$

If our analysis (or a different one) is incorrect, then either the expression-level contification will fail, or it will succeed but the later dead code elimination won't be able to remove the original function binding. For instance, the first situation would arise if the analysis were to mark as contifiable anything other than function bindings, and the second situation would arise if the analysis were to mark as contifiable a function that is actually called with different continuation arguments. Neither situation can result in a violation of semantics preservation.

A note on variable uniqueness. Contification, as presented above, basically duplicates existing code (the function body). Clearly this does not preserve the uniqueness property of variables, which was painfully established by CPS and which later compiler passes rely on. The obvious solution to this issue lies in freshening the new copy of the function body with the help of *efreshen* from Section 5.8.3, in the implementation of the Add_k transformer. This immediately raises two further issues, though: (i) We can't tell which variables *efreshen* is allowed to use (*i.e.*, which n to pass), and (ii) the transformation framework's concrete syntactic correctness criterion actually forbids the introduction of new bound variables (Section 5.3.2.2).

For this reason, Pilsner's contification is based on a slight generalization of the transformation framework from Section 5.3.2. In this version, an expression transformation additionally takes the index n of the next unused variable as an argument (which, for instance, it can pass on to *efreshen*). Moreover, the restriction on bound variables is relaxed: a transformation may introduce new bound variables, as long as their numeric values lie between the given n and some n' that the transformation

must output (of course each of them must be bound only once). In this way, the framework can provide an appropriate n for each transformation and ensure that the newly introduced variables are disjoint from any existing ones and from any that get introduced by other transformations.

5.9.2 Verification

The first transformation, the transformer for Add_k is very easy to prove correct since all it does is introduce an unused binding (the reverse of dead code elimination). Recall the expressions from the beginning of this section. All we have to do is basically apply the compatibility lemmas for let and fix , take a step on the target side (binding the continuation k_y), stutter on the source side, and finally apply reflexivity of e (Lemma 70).

The proof of the second transformation, the transformer for Subst_k , mainly boils down to a lemma relating the new continuation calls to the old function calls, which looks roughly as follows:

$$\frac{\psi_a(k_y) = \text{cont } y_2. e[y_2/y_1][k'/k] \quad \psi_b(y) = \text{fix } f(y_1, k). e}{\Gamma; \psi_a; \psi_b \vdash k_y x \lesssim_w^* y x k'}$$

Here we make crucial use of the "context-sensitive" version of \lesssim that we introduced in Section 5.8.2. It lets us express the connection between the two programs's environments via the static descriptions ψ_a and ψ_b : the code of k_y in the transformed program must be precisely the contification of the code of y in the original program (glossing over the aforementioned freshening). Without such an assumption, the connection would be lost and there would be no hope of proving the calls related.

5.10 The Codegen Pass

Code generation, Pilsner's final pass, translates \mathcal{I} code to the machine language \mathcal{T} .

5.10.1 Transformation

5.10.1.1 Compiling Variable Lookups

We first explain how variables are looked up at run-time. Recall that there are three kinds of "variables" in \mathcal{I} : labels F , term variables y , and continuation variables k . For each we define a helper function that generates the code for lookups, as shown in Figure 5.11. In either case, the lookup is based on the variable's position in the well-formedness context Γ of the expression being compiled (we do not store any variable names in memory). Hence these functions (and thus code generation itself) all take Γ as an input. Let us look at the details. (In an attempt to improve readability, we will often omit parentheses around arguments to meta-level functions here.)

```

lookupLbl ∈ Lbl → Word → [Var] → Reg → [Instr]
lookupLbl F n Γ r :=
  let i := index F (Γ|Lbl) in
  lpc r,
  bop (−) r r n,
  ld r [r + i]

lookupTVar ∈ TVar → [Var] → Reg → [Instr]
lookupTVar y Γ r :=
  let i := index y (Γ|TVar) in
  ld r env,
  repeat i (ld r [r + 1]),
  ld r [r + 0]

lookupKVar ∈ KVar → [Var] → [Instr]
lookupKVar k Γ :=
  let i := 2 * index k (Γ|KVar) + 2 in
  bop (−) sp r i,
  ld ret ⟨sp + 0⟩,
  ld env ⟨sp + 1⟩

```

Figure 5.11: Code generated for variable lookups.

Labels are looked up in the static label environment that the loader places as part of the group header right before the data block in the heap (recall the memory layout from Figure 4.10 and the goal of position-independent code). In order for this to work, the various code generating functions keep track of the label environment's distance n relative to the code being produced.

$lookup_{\text{Lbl}}$ emits the code for label lookups. Each lookup consists of three instructions. It first gets its own code address (in the heap) and—using the distance n —calculates the absolute address of the label environment. Then it indexes the environment according to F 's position⁴ in Γ (restricted to labels) and loads the correct value into the given register r .

Term variables are looked up in a singly-linked list in the heap, which is immutable except that it can grow. Code emitted for expressions expects the `env` register to point to this list.

The lookup code, generated by $lookup_{\text{TVar}}$, is straightforward: it skips the first i elements in the list and then loads the next one into r , where i is the position of y in Γ .

Continuation variables are looked up on the stack. Each continuation takes up two slots: the first points to the continuation's code and the second to its term variable environment (both in the heap). Code emitted for expressions expects that the topmost stack cells correspond to the continuation variables in Γ . (Pilsner code does not use the stack for anything else, but imported functions that were not compiled by Pilsner may of course store arbitrary data there. For instance, Zwickel compiled code uses it to store intermediate values.)

$lookup_{\text{KVar}}$ calculates k 's position in the stack and then loads its code address into register `ret` and its environment address into register `env` (we will see later why it makes sense to always use these registers). Subsequently, it pops the continuation and all above it (more recently defined ones) from the stack by decrementing the stack pointer `sp`. This makes the lookup a destructive operation; it is safe because \mathcal{I} 's affinity property (cf. Section 4.2.4) ensures that using k now means that k won't be needed later and neither do any continuations that got defined after k (in the same lexical scope).

5.10.1.2 Compiling Expressions

Figures 5.12 and 5.13 show *acodegen* and *ecodegen*, which are responsible for translating basic expressions $a \in \text{BExp}$ and control expressions $e \in \text{Exp}$, respectively. Like above, the argument n is the size of the code emitted so far.

In the case of basic expressions only, the generated code writes the result value into the heap, at the address indicated by the `arg` register. This is always the last thing it does, so the execution will run into whatever instructions follow. On the

⁴The *index* function counts from 0.

other hand, code for control expressions will always jump to another function or continuation. Let us take a closer look at some representative examples.

Sums. The code generated for `inl x` first uses *lookup* to load the value of x into register `ret` (which is used as a temporary register here). It then allocates two heap cells, one for the tag of the sum (0 indicates "in left"), and one for the value of x . The code ends after writing the sum value itself, represented by the address of the first cell, into the result location (see above). Note that x may be either a label or a term variable. Here, and below, we write $lookup_{\text{Lbl} \cup \text{TVar}}$ for the obvious combination of $lookup_{\text{Lbl}}$ and $lookup_{\text{TVar}}$.

Let bindings. The code emitted for `let $y = a$ in e` (by *ecodegen*) is equally straightforward. It allocates a new list node (two heap cells) and sets its "next" pointer (the second cell) to the current environment. Then it executes c_1 , the code for a . Since it does so when `arg` points to the new node, a 's result will automatically end up in the right place, namely in the node's first cell. At this point the node represents the extension of the environment list with the value of a . The code sets `env` to this extended list and executes the code for e .

Note that we must be careful to provide the correct offsets when using *acodegen* and *ecodegen* to generate code for the subexpressions. For instance, we pass $3 + n$ to *acodegen* because c_1 will be preceded by three instructions.

Function calls. When compiling a function call, we want to produce the same code no matter what kind of function we are dealing with (imported from an unknown module vs. exported by the current module vs. defined locally in the current module). Hence when we compile a function, we always make it follow the regular calling convention (introduced in Section 4.2.5). This makes it very simple to define *ecodegen* for function calls. The code generated for $x_2 \ x_1 \ k$ first looks up the value of the argument x_1 and stores it into register `arg`. Next, it looks up the value of the function x_2 and stores it into register `clo`. Finally, it looks up the continuation variable k and performs the prescribed indirect jump to the function.

Recall that looking up a continuation has the effect of writing the continuation's code address into `ret` and its environment into `env`. This fits very well, because a function expects its return address in `ret`. Moreover, since `env` is a callee-save register, the continuation can be sure to find its environment intact when it eventually gets control.

Continuation calls. The code emitted by *ecodegen* for continuation calls $k \ x$ loads the argument value into register `arg`, loads the continuation and performs a direct jump to its code.

Functions. As far as the calling convention is concerned, a function value is just an address on which to perform an indirect jump, *i.e.*, the address of a heap cell that

$$acodegen \in \text{BExp} \rightarrow \text{Word} \rightarrow [\text{Var}] \rightarrow [\text{Instr}]$$

$$\begin{aligned} acodegen \text{ (inl } x) n \Gamma := & \\ & lookup_{\text{Lb} \cup \text{TV}_{\text{ar}}} x n \Gamma \text{ ret,} \\ & \text{ld aux}_1 2, \\ & \text{alloc clo aux}_1, \\ & \text{ld aux}_1 0, \\ & \text{sto [clo + 0] aux}_1, \\ & \text{sto [clo + 1] ret,} \\ & \text{sto [arg + 0] clo} \end{aligned}$$

$$\begin{aligned} acodegen \text{ (fix } f(y, k). e) n \Gamma := & \\ & \text{let } c := ecodegen e (22 + n) (\Gamma, f, y, k) \text{ in} \\ & \begin{array}{l} \text{ld aux}_1 2, \\ \text{alloc clo aux}_1, \\ \text{lpc aux}_1, \\ \text{bop (+) aux}_1 \text{ aux}_1 7, \\ \text{sto [clo + 0] aux}_1, \\ \text{sto [clo + 1] env,} \\ \text{sto [arg + 0] clo,} \\ \text{bop (+) clo aux}_1 (13 + \text{length } c), \\ \text{jmp clo,} \\ \text{sto [sp + 0] ret,} \\ \text{sto [sp + 1] env,} \\ \text{bop (+) sp sp 2,} \\ \text{ld aux}_1 2, \\ \text{alloc env aux}_1, \\ \text{sto [env + 0] clo,} \\ \text{ld aux}_1 [\text{clo} + 1], \\ \text{sto [env + 1] aux}_1, \\ \text{ld aux}_1 2, \\ \text{alloc clo aux}_1, \\ \text{sto [clo + 0] arg,} \\ \text{sto [clo + 1] env,} \\ \text{ld env clo,} \\ c \end{array} \end{aligned}$$

} active part

} inactive part

Figure 5.12: Code generation for basic expressions.

$$\begin{aligned}
& \text{ecodegen} \in \text{Exp} \rightarrow \text{Word} \rightarrow [\text{Var}] \rightarrow [\text{Instr}] \\
\\
& \text{ecodegen} (\text{let } y = a \text{ in } e) n \Gamma := \\
& \quad \text{let } c_1 := \text{acodegen } a (3 + n) \Gamma \text{ in} \\
& \quad \text{let } c_2 := \text{ecodegen } e (4 + \text{length } c_1 + n) (\Gamma, y) \text{ in} \\
& \quad \text{ld aux}_1 2, \\
& \quad \text{alloc arg aux}_1, \\
& \quad \text{sto } [\text{arg} + 1] \text{ env}, \\
& \quad c_1, \\
& \quad \text{ld env arg}, \\
& \quad c_2 \\
\\
& \text{ecodegen} (\text{let } k = \text{cont } y. e_2 \text{ in } e_1) n \Gamma := \\
& \quad \text{let } c_1 := \text{ecodegen } e_1 (5 + n) (\Gamma, k) \text{ in} \\
& \quad \text{let } c_2 := \text{ecodegen } e_2 (10 + \text{length } c_1 + n) (\Gamma, y) \text{ in} \\
& \quad \text{lpc aux}_1, \\
& \quad \text{bop } (+) \text{ aux}_1 \text{ aux}_1 (5 + \text{length } c_1), \\
& \quad \text{sto } \langle \text{sp} + 0 \rangle \text{ aux}_1, \\
& \quad \text{sto } \langle \text{sp} + 1 \rangle \text{ env}, \\
& \quad \text{bop } (+) \text{ sp sp } 2, \\
& \quad c_1, \\
& \quad \text{ld clo env}, \\
& \quad \text{ld aux}_1 2, \\
& \quad \text{alloc env aux}_1, \\
& \quad \text{sto } [\text{env} + 0] \text{ arg}, \\
& \quad \text{sto } [\text{env} + 1] \text{ clo}, \\
& \quad c_2 \\
\\
& \text{ecodegen} (x_2 x_1 k) n \Gamma := \\
& \quad \text{let } c_1 := \text{lookup}_{\text{Lbl} \cup \text{TVar}} x_1 n \Gamma \text{ arg in} \\
& \quad \text{let } c_2 := \text{lookup}_{\text{Lbl} \cup \text{TVar}} x_2 (\text{length } c_1 + n) \Gamma \text{ clo in} \\
& \quad \text{let } c_3 := \text{lookup}_{\text{KVar}} k \Gamma \text{ in} \\
& \quad c_1, c_2, c_3, \\
& \quad \text{jmp } [\text{clo} + 0] \\
\\
& \text{ecodegen} (k x) n \Gamma := \\
& \quad \text{let } c_1 := \text{lookup}_{\text{Lbl} \cup \text{TVar}} x n \Gamma \text{ arg in} \\
& \quad \text{let } c_2 := \text{lookup}_{\text{KVar}} k \Gamma \text{ in} \\
& \quad c_1, c_2, \\
& \quad \text{jmp ret}
\end{aligned}$$

Figure 5.13: Code generation for control expressions.

contains the address of the function's first instruction. For top-level functions, which are closed (modulo labels), this is all we need. In fact, the loader already creates these values as part of the group header, so besides producing the function code, we just need to fill these values with the correct code pointers.

For local functions, there's more to do: they need to be converted to proper closures, *i.e.*, pairs of code pointer and term variable environment, for which we must dynamically allocate heap storage. We represent closures as two consecutive cells, where the first one holds the code pointer and the second one the environment pointer. The function value itself then is the address of the first cell. This still matches the calling convention: an indirect jump on the function value will reach the first instruction of its code.

So let us look at the code emitted for local functions $\text{fix } f(y, k). e$ by *acodegen*. It consists of an active and an inactive part. The active part, executed immediately, creates and returns the function closure; the inactive part is the code that the closure references, *i.e.*, that gets executed when the function is called.

- The active part starts off by allocating two cells of heap storage for the closure. It then computes the absolute address of the beginning of the inactive part and writes it into the first cell. Next, it writes the current environment pointer into the second cell (recall that the values in this environment are read-only). Having created and populated the closure, it stores the function value, *i.e.*, the address of the first closure cell, at the designated result location in the heap (see above). Finally, it computes the address of the last instruction in the inactive part and jumps one behind it, *i.e.*, it skips over the inactive part and jumps to whatever code follows.
- The inactive part consists of a prologue and of c , the code for the function body e . The prologue first pushes the contents of registers `ret` and `env` on the stack. Since the stack acts as the continuation environment, pushing `ret` and `env` (in this order) means extending the continuation environment with an entry for the function's continuation argument k . If the function gets called from Pilsner-compiled code, then this makes perfect sense as `ret` will be the continuation's code address and `env` its term environment (see the compilation of function calls above). But even when the function gets called from arbitrary code, this makes perfect sense: the calling convention requires that the function returns to the address in `ret` and preserves the contents of `env`. Treating `ret` plus `env` as a Pilsner continuation will have exactly this effect.

After pushing k on the stack, the code sets up the term variable environment for e by extending the function's own environment with entries for f and y . Concretely, it first creates a list node for f in which it writes the function value itself (stored in `clo` as per the calling convention) to it. It attaches this node to the function's existing term environment, which we know resides in $[\text{clo} + 1]$ (the second cell of the closure). Next, it creates a list node for y , writes the contents of `arg` to it, attaches it to the previous node, and writes its address

$$\begin{aligned}
& \text{codegen} \in \mathbf{Mod}_{\mathcal{I}} \rightarrow [\text{Lbl}] \rightarrow \mathbf{Mod}_{\mathcal{T}} \\
& \text{codegen } M \Gamma := (\text{mcodegen } M \text{ (length } \Gamma + 2 * \text{length } M) 0 \Gamma) \\
\\
& \text{mcodegen} \in \mathbf{Mod}_{\mathcal{I}} \rightarrow \text{Word} \rightarrow \text{Word} \rightarrow [\text{Lbl}] \rightarrow \mathbf{Mod}_{\mathcal{T}} \\
& \text{mcodegen } \epsilon m n \Gamma := (\epsilon, \epsilon) \\
& \text{mcodegen } (F=a, M) m n \Gamma := \\
& \quad \text{let } c := \text{tcodegen } a (m + n) \Gamma \text{ in} \\
& \quad \text{let } g := \text{mcodegen } m (n + \text{length } c) (\Gamma, F) \text{ in} \\
& \quad ((F, n), g.\text{cptrs}, (\text{map } \mathbb{E} c, g.\text{data})) \\
\\
& \text{tcodegen } (\text{fix } f(y, k). e) n \Gamma := \\
& \quad \text{let } c := \text{ecodegen } e (10 + n) (\Gamma, f, y, k) \text{ in} \\
& \quad \text{sto } \langle \text{sp} + 0 \rangle \text{ ret,} \\
& \quad \text{sto } \langle \text{sp} + 1 \rangle \text{ env,} \\
& \quad \text{bop } (+) \text{ sp sp } 2, \\
& \quad \text{ld aux}_1 2, \\
& \quad \text{alloc aux}_1 \text{ aux}_1, \\
& \quad \text{sto } [\text{aux}_1 + 0] \text{ clo,} \\
& \quad \text{ld env } 2, \\
& \quad \text{alloc env env,} \\
& \quad \text{sto } [\text{env} + 0] \text{ arg,} \\
& \quad \text{sto } [\text{env} + 1] \text{ aux}_1, \\
& \quad c
\end{aligned}$$

Figure 5.14: Code generation for modules.

into env. This concludes the function prologue and c is executed next.

Continuation bindings. As the final example, consider continuation bindings $\text{let } k = \text{cont } y. e_2 \text{ in } e_1$. The code emitted by *ecodegen* first gets its own address so that it can calculate the absolute address of the continuation code (starting with the first load instruction). It pushes this address onto the stack, together with the current environment pointer, thus extending the continuation environment with an entry for k . Next, the code for the binding's body e_1 is executed. If and when this eventually leads to the continuation being invoked (directly or indirectly), the continuation can be sure to find its term environment and input value in *env* and *arg*, respectively. It then extends its environment with the input value in order for e_2 's code, which it executes next, to find it there.

5.10.1.3 Compiling Modules

Code generation for modules is shown in Figure 5.14. The final transformation is *codegen*, which takes an \mathcal{I} module with a matching wellformedness context Γ (*i.e.*,

the list of imports), and produces a \mathcal{T} module. This module consists of a single group (recall the structure of machine modules from Section 4.2.5), which is generated by the helper function *mcodegen*. For each function definition $F=a$ in M , *mcodegen* adds an entry for F to the group's code pointer table and adds the code produced for a to the group's data block (after encoding the instructions as machine words). In order to create the correct code pointer table entries, *mcodegen* maintains the size n of the code generated so far (for the functions "to the left"). Initially this is 0, so the first function will be mapped to 0 as this is the offset in the data block where its code starts. When using *tcodegen* to compile the functions, it also passes the correct offset along, which is always the current n plus the fixed size m of the group header.

The helper *tcodegen* translates top-level functions, and top-level functions only (Figure 5.14 shows the fix case but omits the analogous Λ case). As discussed previously, top-level functions are somewhat simpler than local functions. The code generated by *tcodegen* corresponds to the inactive part that *acodegen* would generate, except that the prologue that sets up the term environment doesn't need to load and extend any existing environment.

5.10.2 Verification

Code generation is the most radical and low-level transformation in Pilsner, and so it comes as no surprise that its proof is also the longest and most tedious. Here we can only give an overview.

The main goal is to show the following theorem, which states the semantics preservation of *codegen* in terms of our canonical $\lesssim_{\mathcal{TI}}$ relation.

Theorem 21.

$$\frac{\Gamma \vdash M : \Gamma'}{\Gamma \vdash \text{codegen } M \Gamma \lesssim_{\mathcal{TI}} M : \Gamma'}$$

Analogous to earlier proofs, a key part of this is finding a relation $\preccurlyeq_{\mathcal{TI}}$ that formalizes when some machine code implements an open \mathcal{I} expressions. For this relation we must prove something along the lines of

$$\frac{\Gamma \vdash e}{\Gamma \vdash \text{ecodegen } e \preccurlyeq_{\mathcal{TI}} e}$$

(and similarly for *acodegen*).

In $\preccurlyeq_{\mathcal{TI}}$, we want to ultimately say that the two programs are similar according to **E**. But clearly the generated code makes many assumptions about its environment, *e.g.*, where term variables can be looked up, how continuations are laid out on the stack, where temporary results are placed, etc. In order to restrict the environments in which the code is placed, we must therefore express parts of the compiler-internal protocol that the code follows. Naturally, the choice of the local world in Theorem 21 plays an important role in this.

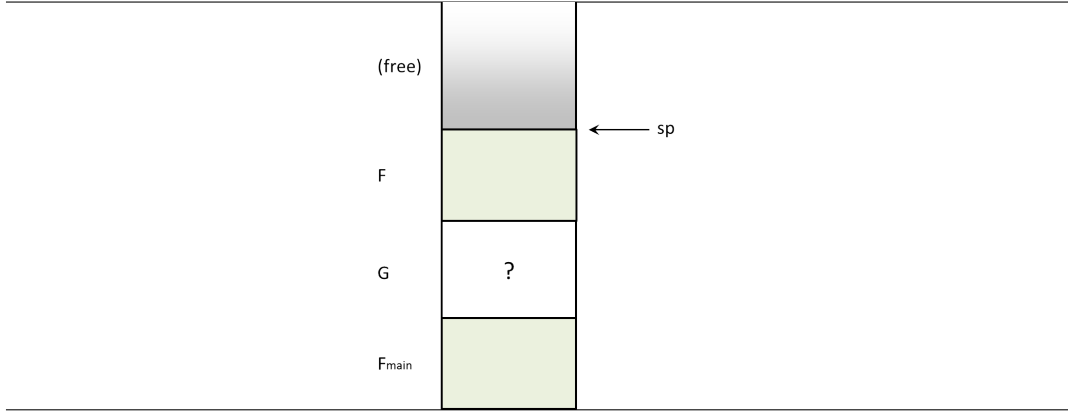


Figure 5.15: Example of Pilsner’s local stack.

5.10.2.1 The Local World

We define a local world w for the verification of code generation such that a state s consists of a stack (fragment) and a heap (fragment). They represent the memory used *internally* by the machine code generated for module M . The configuration relation $w.C$ ensures that \mathcal{T} ’s local heap is precisely the one given by the state, and similarly for the stack (so from now on we use the terms local and internal interchangeably).

Let us review the internal memory that Pilsner-generated code maintains.

Stack. The local stack $s.q$ consists of the current continuation environment. Note that this may not be one consecutive chunk due to calls to imported functions that haven’t returned yet.

To illustrate a bit better what the local stack typically looks like, let’s assume that the module for which we generate code provides the “main” function F_{main} . And that F_{main} calls an imported higher-order function G , passing as argument another function F from our module. After G calls F , the full machine stack looks as shown in Figure 5.15. The parts that constitute the local stack $s.q$ are shown in green.

The first chunk is part of our local stack and contains the continuations that F_{main} ’s code allocated before it called G . The second chunk is not part of our local stack, but part of the frame configuration that the \mathbf{E} relation quantifies over. It contains whatever G had pushed before it called F . We have no idea what that is (if anything at all) because G may not have been compiled by Pilsner. The third chunk is again part of our local stack. It contains the continuations that F ’s code just allocated. The final chunk of the stack is free and infinite. It is owned by the global world, as discussed in Section 4.4.2.3.

Heap. As is standard, the local heap $s.h$ contains the label environment. Of interest is what else it consists of: all the module’s code and the singly linked lists representing

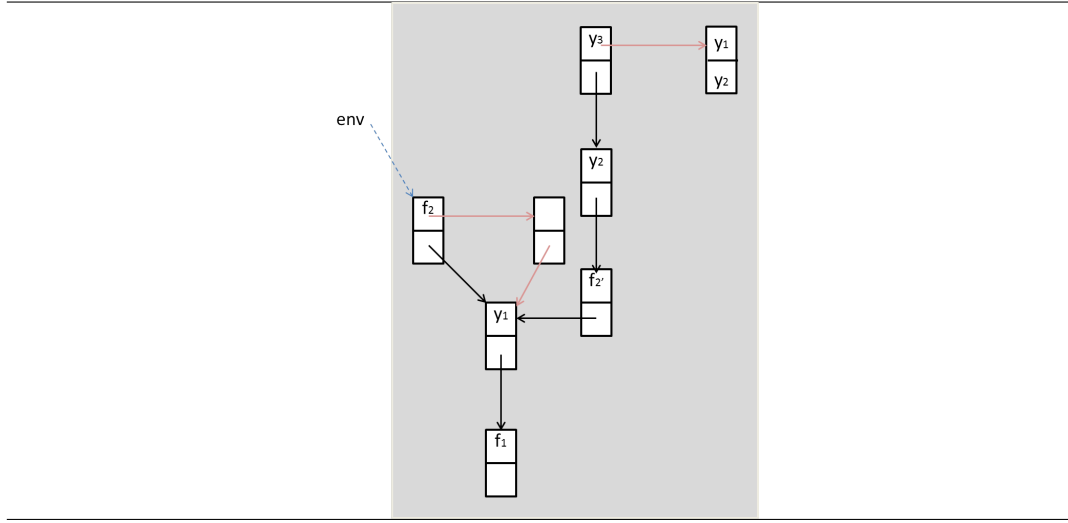


Figure 5.16: Example of Pilsner's local heap.

term variable environments that have been allocated so far (typically, many of them share nodes).

For instance, consider a module consisting of a single function F :

$$\begin{aligned}
 F &= \text{fix } f_1(y_1, k_1). \\
 &\quad \text{let } f_2 = (\text{fix } f_2'(y_2, k_2). \text{let } y_3 = \langle y_1, y_2 \rangle \text{ in } k_2 y_3) \text{ in} \\
 &\quad k_1 f_2
 \end{aligned}$$

Right before the execution of the main function, the local heap consists solely of the label environment and the code for F (which, of course, includes the code for the local function f_2). After executing the main function's prologue, the local heap additionally contains a term environment that provides the values for f_1 and y_1 (pointed to by register `env`). Executing the `let`-binding has two effects: (i) it allocates a new closure and (ii) it extends the environment with a new list node containing the closure (the value for f_2). Step (i) affects the local heap only partially: the second cell of the closure, which points to the initial term environment, is part of the local heap. However, the first cell, which contains the code pointer and represents the value of the function, is part of the global heap instead. This is enforced by our global world $\Omega_{\mathcal{TI}}$: if we want f_2 to be recognized as a proper function (*e.g.*, so that we can pass it to the unknown k_1) then we must register it as such in the global state's value registry. This registry, as we have seen in Section 4.4.2.3, is tied to the global heap.

The next time the local heap changes is when k_1 calls f_2 (assuming such a call is made, and assuming k_1 does not happen to call F itself before). The prologue of f_2 extends f_2 's environment with nodes for f_2' and y_2 . Then the `let`-binding (i) allocates a new pair and (ii) further extends the environment with a node for y_3 . Step (i) does not affect the local heap at all, because the representation of a pair as the address of the first of two heap cells is nothing internal but rather part the official "interface" of

values that every module must agree with. $\Omega_{\mathcal{TI}}$ enforces via the value registry that both cells of the pair are exposed.

As a summary, Figure 5.16 depicts the shape of the local heap when execution reaches the call to k_2 inside of f_2 . Besides the group environment and the code, it contains f_2 's closure term environment, the cell of f_2 's closure that points to this environment, and an extension of that environment matching the current execution of f_2 . The black arrows denote the "next" pointers of the singly-linked environments.

The formal definition of w is quite simple:

Definition 50 (Local world w for the verification of code generation).

$$\begin{aligned}
w &\in \text{LWorld}^{\Omega_{\mathcal{IS}}, \mathcal{T}} \\
w.\mathcal{T}.\mathcal{S} &:= \text{Stack} \times \text{Heap} \\
w.\mathcal{T}.\sqsubseteq &:= \{ ((q', h'), (q, h)) \mid h' \supseteq h \} \\
w.\mathcal{T}.\sqsubseteq_{\text{pub}} &:= \{ ((q', h'), (q, h)) \mid h' \supseteq h \wedge q' = q \} \\
w.\mathcal{C}(G)(s_g, (q, h)) &:= \{ (c_a, c_b) \mid c_a = (\emptyset, \emptyset, q, h) \wedge c_b = (\emptyset, \emptyset) \} \\
w.\mathcal{O} &:= \text{gwf}
\end{aligned}$$

Note that a state does not record any information about the \mathcal{I} program's memory. \mathcal{I} 's heap is only used for references, which we can all treat as global in this verification. Hence the \mathcal{I} configurations in $w.\mathcal{C}$ are empty.

Both transition relations allow the local heap to grow but not to be modified otherwise, which matches the use in Pilsner code. Regarding the local stack, the full transition relation allows arbitrary changes while the public one allows no changes at all. This reflects the calling convention: a function may modify the stack arbitrarily as long as it restores it again before returning.

We pick w 's initial state analogous to the example proof in Section 4.6. In particular, we choose its local heap to be exactly the local heap given by **cloud**. This heap consists of two continuous chunks: (i) the label environment of size $\text{length } \Gamma + \text{length } \Gamma'$, starting at whatever load address $\Psi_{\mathcal{T}}$ we are given, and (ii) the code, starting at address $\Psi_{\mathcal{T}} + \text{length } \Gamma + 2 * \text{length } \Gamma'$. The initial local stack is empty.

5.10.2.2 Relating Machine Code to Open Expressions

We now define the $\preceq_{\mathcal{TI}}$ relation, which is of the form $\Gamma \vdash f \preceq_{\mathcal{TI}} e$ and $\Gamma \vdash f \preceq_{\mathcal{TI}} a$. But what is f ? It doesn't make sense to directly relate machine code (instruction sequences) to control expressions or basic expressions, because the code that Pilsner generates depends on its distance n to the label environment. We model this by relating a meta-level function $f \in \text{Word} \rightarrow [\text{Word}]$ instead, which maps n to (encoded) instructions. In practice, f will always be $\lambda n. \text{map } \mathbb{E}(\text{ecodegen } e \ n \ \Gamma)$ and $\lambda n. \text{map } \mathbb{E}(\text{acodegen } a \ n \ \Gamma)$, respectively.

In the case of control expressions e , $\preceq_{\mathcal{TI}}$ is defined as follows.

Definition 51.

$$\begin{aligned}
\boxed{\Gamma \vdash f \preceq_{\mathcal{TI}} e} &:= \forall i. \exists i'. \forall G, s_0, s, K_a, K_b, \psi_a, \sigma_a, \sigma_b, m, cur, old. \\
&G \in \mathbf{GK}_{w\uparrow} \wedge \\
&s \sqsupseteq s_0 \wedge \\
&m \neq 0 \wedge \\
&(\exists n. s.h \supseteq [m \mapsto f\ n] \sqcup [m - n \mapsto map\ snd\ \psi_a]) \wedge \\
&(\forall F \in \Gamma. \exists v_a, v_b. \psi_a(F) = v_a \wedge \sigma_b(F) = v_b \wedge (v_a, v_b) \in \overline{G}(s)) \wedge \\
&\Gamma|_{\mathbf{Lbl}} = map\ fst\ \psi_a \wedge \\
&repr\ s.h\ (s.R(\mathbf{env}))\ (map\ snd\ \sigma_a) \wedge \\
&(\forall y \in \Gamma. \exists v_a, v_b. \sigma_a(y) = v_a \wedge \sigma_b(y) = v_b \wedge (v_a, v_b) \in \overline{G}(s)) \wedge \\
&\Gamma|_{\mathbf{TVar}} = map\ fst\ \sigma_a \wedge \\
&s.q = [s_0.R(\mathbf{sp}) \mapsto cur] \sqcup old \wedge \\
&s.R(\mathbf{sp}) = s_0.R(\mathbf{sp}) + length\ cur \wedge \\
&length\ cur = 2 * length\ \Gamma|_{\mathbf{KVar}} \wedge \\
&K\ G\ s_0\ s\ K_a\ K_b\ i\ \sigma_b\ \Gamma \\
&\implies ((m, \emptyset, \emptyset, \emptyset), (\emptyset, (\sigma_b, e))) \in \mathbf{E}_{w\uparrow}(i' + i)(K_a, K_b)(G)(s_0)(s)(\perp)
\end{aligned}$$

Definition 52.

$$\boxed{repr\ h\ v\ d} \quad \overline{repr\ h\ v\ e} \quad \frac{h(v) = w \quad h(v+1) = v' \quad repr\ h\ v'\ W}{repr\ h\ v\ (w, W)}$$

The conclusion of the implication states that a \mathcal{T} configuration consisting of a program counter m is \mathbf{E} -related in state s of $w\uparrow$ to an \mathcal{I} configuration consisting of e and some closing environment σ_b . Let us walk through the assumptions under which this must hold.

The first larger condition connects the program counter m to f : it says that the current local heap contains the sequence of words given by $f\ n$ (think: the code), starting at address m . The distance n is arbitrary but the condition also tells us that the local heap contains some data starting at address $m - n$.

The next condition gives meaning to this data: it is actually the label environment. For any label F in Γ , the data chunk contains a value that is related to whatever value F has in e 's environment σ_b . The subsequent condition says that the values in the data chunk are ordered according to Γ . In effect, these conditions tell us that looking up a label will yield the expected result (see also Lemma 78 below).

The next three conditions concern term variables. With the help of the predicate from Definition 52, the first of these states that the current value of register \mathbf{env} (exposed via the global state) points to a singly-linked list in the local heap containing certain elements (think: the term environment). The subsequent condition tells us something about these elements: for each term variable y in Γ , there is an element that is related to whatever value y has in e 's environment σ_b . And yet another condition says that these elements are ordered according to Γ . In effect, these three conditions tell us that looking up a term variable will yield the expected result (see also Lemma 78 below).

The remaining four conditions concern continuation variables. The first says that the local stack can be split into a "current" and an "old" part, and that the current part starts where the stack ended in state s_0 . Next, that the current part is actually the top part of the machine stack right now. The last two conditions intuitively say that the current part is the continuation environment matching Γ (recall that each continuation there takes up two slots) and that the stored continuations are related to the corresponding ones in e 's environment σ_b (see also Lemma 78 below).

This last condition is itself somewhat complex, as it must in turn describe the assumptions that code generated for continuations makes:

Definition 53.

$$\begin{aligned}
\boxed{\mathsf{K} \, G \, s_0 \, s \, K_a \, K_b \, i \, \sigma_b \, \Gamma} &:= \forall k \in \Gamma. \, \forall j, \mathbf{k}_a, w, \mathbf{k}_b. \, j = \text{index } k \, \Gamma|_{\mathsf{KVar}} \implies \\
&\quad s.q(s.R(\mathsf{sp}) - 2 * (1 + j)) = \mathbf{k}_a \wedge \\
&\quad s.q(s.R(\mathsf{sp}) - 2 * (1 + j) + 1) = w \wedge \\
&\quad \sigma_b(k) = \mathbf{k}_b \wedge \\
&\quad (\forall G', s', v_a, v_b. \\
&\quad \quad G' \supseteq G \wedge s' \sqsupseteq s \wedge s'.R(\mathsf{sp}) = s.R(\mathsf{sp}) - 2 * (1 + j) \wedge \\
&\quad \quad (\exists p \in [\mathsf{Word}]. \, \text{length } p = 2 * (1 + j) \wedge s.q = s'.q \sqcup [s'.R(\mathsf{sp}) \mapsto p]) \wedge \\
&\quad \quad s'.R(\mathsf{env}) = w \wedge (v_a, v_b) \in \overline{G'}(s')) \\
&\implies \\
&\quad v \uparrow . \text{cqh}_a(s')(\text{ret } v_a \, \mathbf{k}_a) \times w \uparrow . \text{cqh}_b(s')(\text{ret } v_b \, \mathbf{k}_b) \subseteq \\
&\quad \mathbf{E}_{w \uparrow}(i)(K_a, K_b)(G')(s_0)(s')(\perp)
\end{aligned}$$

For each continuation variable k in Γ , it expresses the following: The first three conditions state that, based on k 's position⁵ j in Γ , there are two adjacent slots associated with it in the local stack. The first contains \mathbf{k}_a (think: code pointer) and the second contains w (think: environment pointer). (k having position j means that there are j continuations above it on the stack). The next condition says that k has value \mathbf{k}_b in e 's environment σ_b .

The last condition then makes the connection between the \mathcal{T} continuation and the \mathcal{I} continuation for k . Roughly, when they are used in some future state s' with related input values v_a, v_b , then they must yield \mathbf{E} -related computations as long as the following preconditions are met: k itself and all continuations that were above it in the stack have been popped (p represents the popped part) and apart from that, the stack is unchanged; furthermore, k 's environment must have been loaded into register env . It is easy to see that the code generated by *ecodegen* for continuation calls establishes these conditions.

Let us now consider the definition of $\preceq_{\mathcal{TI}}$ for basic expressions $a \in \mathbf{BExp}$. Its purpose is to capture the contract that *acodegen*-generated code adheres to.

⁵A variable that occurs multiple times is only considered once, as *index* returns the index of the only relevant occurrence.

Definition 54.

$$\begin{aligned}
\boxed{\Gamma \vdash f \preceq_{\mathcal{TI}} a} &:= \forall i, G, s_0, s, K_a, K_b, \psi_a, \sigma_a, \sigma_b, m, n, v, i', x, e, c_a, c_b, \eta_a, \eta_b. \\
&G \in \mathbf{GK}_{w\uparrow} \wedge \\
&s \sqsupseteq s_0 \wedge \\
&m \neq 0 \wedge \\
&(\exists h. s.h = [m \mapsto f\ n] \sqcup [m - n \mapsto \text{map snd } \psi_a] \sqcup h) \wedge \\
&(\forall F \in \Gamma. \exists v_a, v_b. \psi_a(F) = v_a \wedge \sigma_b(F) = v_b \wedge (v_a, v_b) \in \overline{G}(s)) \wedge \\
&\Gamma|_{\text{Lbl}} = \text{map fst } \psi_a \wedge \\
&\text{repr } s.h (s.R(\text{env})) (\text{map snd } \sigma_a) \wedge \\
&(\forall y \in \Gamma. \exists v_a, v_b. \sigma_a(y) = v_a \wedge \sigma_b(y) = v_b \wedge (v_a, v_b) \in \overline{G}(s)) \wedge \\
&\Gamma|_{\text{TVar}} = \text{map fst } \sigma_a \wedge \\
&(\forall s' \sqsupseteq_{\text{pub}} s. \forall v_a, v_b. \\
&\quad s'.R(\text{arg}) = s.R(\text{arg}) \wedge \\
&\quad s'.h(s'.R(\text{arg})) = v_a \wedge \\
&\quad (v_a, v_b) \in \overline{G}(s') \\
&\quad \implies ((m + \text{length}(f\ n), \emptyset, \emptyset, \emptyset), \\
&\quad (\emptyset, (\sigma_b[x \mapsto v_b], e))) \in \mathbf{E}_{w\uparrow}(i')(K_a, K_b)(G)(s_0)(s')(\perp)) \wedge \\
&\quad (c_a, c_b) \in w\uparrow.C(G)(s) \\
&\implies ((m, \emptyset, \emptyset, [s.R(\text{arg}) \mapsto v]) \cdot c_a \cdot \eta_a, \\
&\quad (\emptyset, (\sigma_b, \text{let } x = a \text{ in } e)) \cdot c_b \cdot \eta_b) \in \mathbf{E}_{w\uparrow}(i)(K_a, K_b)(G)(s_0)(s)(\eta_a, \eta_b)
\end{aligned}$$

Several parts here look very much the same as for control expressions. Let us highlight the important differences.

1. Since we cannot directly relate a basic expression using \mathbf{E} (in the implication's conclusion), we wrap a with a let-binding. This does not restrict the applicability, because—excluding top-level functions—a let-binding is the only way a basic expression can be used. Top-level functions are not compiled using *acodegen*, so we do not need to consider them here.
2. Instead of using \mathbf{E} in the default mode, we use it in fixed mode (cf. Section 4.5.2) and then explicitly assume that parts of the configurations are world-related (namely c_a and c_b) and other parts are the frames (η_a and η_b). This would be equivalent to just using the default mode and hiding these configurations if it weren't for $[s.R(\text{arg}) \mapsto v]$. Here is the deal: We must assert ownership of the heap cell initially pointed to by register arg , because this is where the code is going to write its result (cf. Section 5.10.1.2). However, we cannot consider this cell part of the local heap before the result has actually been written, because w allows no mutation of the local heap. Concretely, if we were to state that $s.h$ contains $[s.R(\text{arg}) \mapsto v]$, where v is the arbitrary and irrelevant value currently stored in that cell, then we would get stuck when reasoning about the final write of the result: there would be no future state matching the final heap.

As an alternative to this definition's use of the fixed mode, we could probably have used a more complicated local world that allows for some cells of the

internal heap to initially be in an "uncommitted" state, where their contents can still change. The current approach seems simpler, though.

3. The last clause in the implication's premise connects the two big unknowns: the code following f 's code in memory (which the execution of f 's code will eventually run into), and e , the expression that we had to introduce as part of the let-binding.

It says that they must be **E**-related in any public future state s' in which the result cell is part of the local heap and contains a value v_a related to the one that x is mapped to in e 's environment. Note that now that the result cell is actually in the local heap, we can use **E** in its default mode again.

The reason for requiring s' to be a *public* extension is that the local stack cannot have changed (and it is important to know that). In fact, only registers and the value of the result cell can have changed, which is why there is no need to consider a future global knowledge. This is in contrast to control expressions, whose evaluation can change global references and call external functions.

The reason for requiring that the contents of register **arg** hasn't changed is merely that the code of let-bindings relies on this (this could easily be changed, though).

4. Nowhere do we say anything about continuation variables. This is fine, because basic expressions cannot refer to any continuations coming from an outer lexical scope (cf. Section 4.2.4.2).

5.10.2.3 Key Lemma

We are now able to formulate the desired key lemma.

Lemma 77.

$$\frac{\Gamma \vdash e}{\Gamma \vdash (\lambda n. \text{map } \mathbb{E}(\text{ecodegen } e \ n \ \Gamma)) \preceq_{\mathcal{TI}} e}$$

Proof. Reminiscent of compatibility lemmas (cf. Section 3.7.3), we prove one lemma for each form of control expression and basic expression. For instance, the statement for functions—as usual, the only one proven using coinduction—reads as follows:

$$\frac{\Gamma^\dagger, f, y, k \vdash (\lambda n. \text{map } \mathbb{E}(\text{ecodegen } e \ n \ (\Gamma^\dagger, f, y, k))) \preceq_{\mathcal{TI}} e}{\Gamma \vdash (\lambda n. \text{map } \mathbb{E}(\text{acodegen } (\text{fix } f(y, k). e) \ n \ \Gamma)) \preceq_{\mathcal{TI}} (\text{fix } f(y, k). e)}$$

As another example, here is the statement for applications:

$$\frac{x_1 \in \Gamma \quad x_2 \in \Gamma \quad k \in \Gamma}{\Gamma \vdash (\lambda n. \text{map } \mathbb{E}(\text{ecodegen } (x_1 \ x_2 \ k) \ n \ \Gamma)) \preceq_{\mathcal{TI}} (x_1 \ x_2 \ k)}$$

With an easy induction, these properties then yield the goal. \square

The various sub-proofs of Lemma 77 employ Lemmas 78 and 79. These provide Hoare-style reasoning principles for the variable lookup code generated by the *lookup* functions that we saw earlier. In particular the first one is tremendously useful, as code generation makes use of $lookup_{\text{Lbl} \cup \text{TVar}}$ in over twenty places.

Lemma 78 (Label or term variable lookup). If

1. $[pc - n \mapsto map\ snd\ \psi_a] \subseteq h \ \wedge \ \Gamma|_{\text{Lbl}} = map\ fst\ \psi_a$
2. $repr\ h\ (R(\text{env}))\ (map\ snd\ \sigma_a) \ \wedge \ \Gamma|_{\text{TVar}} = map\ fst\ \sigma_a$
3. $[pc \mapsto map\ \mathbb{E}\ c] \subseteq h \ \wedge \ c = lookup_{\text{Lbl} \cup \text{TVar}}\ x\ n\ \Gamma\ r \ \wedge \ pc \neq 0$
4. $\psi_a(x) = v \ \vee \ \sigma_a(x) = v$
5. $pc' = pc + length\ c \ \wedge \ R' = R[r \mapsto v] \ \wedge \ i' <^{length\ c} i$
6. $((pc', R', q, h), c_b) \in \mathbf{E}_{w\uparrow}(i')(K_a, K_b)(G)(s_0)(s)(\eta_a, \eta_b)$

then $((pc, R, q, h), c_b) \in \mathbf{E}_{w\uparrow}(i)(K_a, K_b)(G)(s_0)(s)(\eta_a, \eta_b)$.

The lemma describes the precise effects of executing the *lookup* code and can be read as follows: When reasoning about a \mathcal{T} machine whose program counter points to the first instruction of that code, it suffices to reason about the machine *after* executing that code, which is obtained by incrementing the program counter accordingly and by setting the r register to whatever the requested value is. (The requested value is determined by condition (4) depending on whether x is a label or a term variable.) Note how conditions (1) and (2) match the definition of $\preceq_{\mathcal{T}}$.

Lemma 79 is analogous:

Lemma 79 (Continuation variable lookup). If

1. $[pc \mapsto map\ \mathbb{E}\ c] \subseteq h \ \wedge \ c = lookup_{\text{KVar}}\ k\ \Gamma \ \wedge \ pc \neq 0$
2. $j = index\ k\ \Gamma|_{\text{KVar}} \ \wedge \ R(\text{sp}) = n + 2 * (1 + j) \ \wedge \ q(n) = v \ \wedge \ q(n + 1) = w$
3. $pc' = pc + length\ c \ \wedge \ R' = R[\text{sp} \mapsto n][\text{ret} \mapsto v][\text{env} \mapsto w] \ \wedge \ i' <^{length\ c} i$
4. $((pc', R', q, h), c_b) \in \mathbf{E}_{w\uparrow}(i')(K_a, K_b)(G)(s_0)(s)(\eta_a, \eta_b)$

then $((pc, R, q, h), c_b) \in \mathbf{E}_{w\uparrow}(i)(K_a, K_b)(G)(s_0)(s)(\eta_a, \eta_b)$.

5.11 The Full Pilsner Compiler

At this point, we have verified each of Pilsner's passes separately. Let us summarize the results:

- The first pass:

$$\frac{\Gamma \vdash M : \Gamma'}{\Gamma \vdash cps\ M \preceq_{\mathcal{IS}} M : \Gamma' \quad |\Gamma| \vdash cps\ M : |\Gamma'| \quad uniqmod(cps\ M)}$$

- The intermediate passes:

$$\frac{\Gamma \vdash M : \Gamma' \quad \text{uniqmod}(M)}{\Gamma \vdash f \ M \lesssim_{\mathcal{IL}}^* M : \Gamma' \quad \Gamma \vdash f \ M : \Gamma' \quad \text{uniqmod}(f \ M)}$$

for $f \in \{ \text{inline}, \text{contify}, \text{dce}, \text{hoist}, \text{commute}, \text{dedup} \}$

- The final pass:

$$\frac{\Gamma \vdash M : \Gamma' \quad \text{uniqmod}(M)}{\Gamma \vdash \text{codegen } M \ \Gamma \lesssim_{\mathcal{TS}} M : \Gamma' \quad \vdash \text{codegen } M \ \Gamma : \Gamma'}$$

Thanks to transitivity of PILS (Theorem 15), the composition of all these transformations is correct w.r.t. $\lesssim_{\mathcal{TS}}$.

Note that it doesn't actually matter which of the intermediate transformations are applied. So we can easily allow the user of Pilsner to selectively disable some optimizations⁶. Formally, we define Pilsner as follows. It takes a sequence of Boolean flags, each enabling or disabling one intermediate transformation.

Definition 55 (Pilsner with selection).

$$\begin{aligned} \text{pilsner} &\in \mathbb{B}^6 \rightarrow [\text{Lb}] \rightarrow \mathbf{Mod} \rightarrow \mathbf{Mod} \\ \text{pilsner flags } \Gamma &:= \\ &(\lambda M. \text{codegen } M \ \Gamma) \circ \\ &(\text{if flags.6 then dedup else id}) \circ \\ &(\text{if flags.5 then commute else id}) \circ \\ &(\text{if flags.4 then hoist else id}) \circ \\ &(\text{if flags.3 then dce else id}) \circ \\ &(\text{if flags.2 then contify else id}) \circ \\ &(\text{if flags.1 then inline else id}) \circ \\ &\text{cps} \end{aligned}$$

No matter which flags the user chooses, Pilsner is always correct:

Theorem 22 (Correctness of Pilsner).

$$\frac{\Gamma \vdash M : \Gamma'}{\Gamma \vdash \text{pilsner flags } |\Gamma| \ M \lesssim_{\mathcal{TS}} M : \Gamma' \quad \vdash \text{pilsner flags } |\Gamma| \ M : |\Gamma'|}$$

5.12 The Zwickel Compiler

The Zwickel compiler is a straightforward compiler from \mathcal{S} directly down to \mathcal{T} , not involving any intermediate transformations. As such, it is rather different from Pilsner, and being able to verify it using the PILS system developed in Chapter 4 provides some evidence of the flexibility of our approach.

⁶We could even let the user choose the order in which optimizations are performed, but that may not be very useful as optimizations are typically written with a certain order in mind.

```

lookupLbl ∈ Lbl → Word → [Var] → [Instr]
lookupLbl F n Γ :=
  let i := index F (Γ|Lbl) in
  lpc aux1,
  bop (−) aux1 aux1 n,
  ld aux1 [aux1 + i],
  bop (+) sp sp 1,
  sto ⟨sp − 1⟩ aux1

lookupTVar ∈ TVar → [Var] → [Instr]
lookupTVar y Γ :=
  let i := index y (Γ|TVar) in
  sto clo env
  repeat i (ld clo [clo + 1]),
  bop (+) sp sp 1,
  ld clo [clo + 0],
  sto ⟨sp − 1⟩ clo

```

Figure 5.17: Code generated for variable lookups.

5.12.1 Transformation

Despite the big difference to Pilsner as a whole, Zwickel is somewhat similar to Pilsner’s *code generation* pass (Section 5.10). We can therefore keep this section relatively short by focussing on the differences.

Recall that Pilsner’s codegen compiles \mathcal{I} programs, not \mathcal{S} programs. The code that it generates for a pure expression writes the result value into a new environment slot in the heap (because pure expressions can only occur in let-bindings, which extend the environment); control expressions, on the other hand, do not directly produce a value because they are in CPS. In \mathcal{S} , there is no distinction between pure and control expressions, and every expression produces a value (if it terminates). For Zwickel, we thus follow the convention that the code for an expression e pushes e ’s value on the stack. We could have used a fixed register instead, but using the stack makes Zwickel more different from Pilsner, where the stack is reserved for \mathcal{I} ’s continuations instead.

Figures 5.17–5.19 below roughly correspond to Figures 5.11–5.13. Like Pilsner’s codegen, the Zwickel code for looking up labels F —shown in Figure 5.17—reads from the loader environment whose start address is known relative to the current instruction. Similarly, term variables x are looked up in the linked-list environment pointed to by register `env`. Since labels and term variables in \mathcal{S} are syntactically expressions, however, the found value is in both cases then pushed on the stack. (This is why the *lookup* functions don’t take a register argument, as they did in Pilsner.) Also, recall that there are no continuation variables in \mathcal{S} .

$$ezwickel \in \text{Exp} \rightarrow \text{Word} \rightarrow [\text{Var}] \rightarrow [\text{Instr}]$$

$$ezwickel\ x\ n\ \Gamma := \\ \text{lookup}_{\text{TVar}}\ x\ n\ \Gamma$$

$$ezwickel\ F\ n\ \Gamma := \\ \text{lookup}_{\text{Lbl}}\ F\ n\ \Gamma$$

$$ezwickel\ (\text{inl}\ e)\ n\ \Gamma := \\ ezwickel\ e\ n\ \Gamma, \\ \text{ld}\ \text{clo}\ 2, \\ \text{alloc}\ \text{clo}\ \text{clo}, \\ \text{ld}\ \text{aux}_1\ 0, \\ \text{sto}\ [\text{clo} + 0]\ \text{aux}_1, \\ \text{ld}\ \text{aux}_1\ \langle \text{sp} - 1 \rangle, \\ \text{sto}\ [\text{clo} + 1]\ \text{aux}_1, \\ \text{sto}\ \langle \text{sp} - 1 \rangle\ \text{clo}$$

Figure 5.18: Code generation for expressions.

Figures 5.18 and 5.19 show excerpts of Zwickel’s expression compilation *ezwickel*, which takes the same arguments as *ecodegen* and *acodegen* did in Pilsner. Let us here just focus on compiling functions and function calls. Functions now have an epilogue (the last four instructions) that moves the body’s value, found on the top of the stack, into the function call’s result register **arg**, as required by the calling convention that all modules adhere to. The epilogue also restores the callee-save register **env** that the function stored on the stack as well. The code for an application does not simply pass through a given return address, but provides its own such that it can push the result value back onto the stack (an application is an expression). Of course it must then also restore the original return address in **ret**, which it saved on the stack.

The module-level compilation $zwickel \in \mathbf{Mod}_S \rightarrow [\text{Lbl}] \rightarrow \mathbf{Mod}_T$ is defined analogous to that in Pilsner and thus not shown here.

5.12.2 Verification

The verification of Zwickel follows roughly that of Pilsner’s codegen (Section 5.10). In fact, we use essentially the same local world w .

$ezwickel(\text{fix } f(x).e) n \Gamma :=$
 $\text{let } c := ezwickel e (22 + n) (\Gamma, f, x) \text{ in}$
 $\text{ld clo } 2,$
 $\text{alloc clo clo},$
 $\text{lpc arg},$
 $\text{bop } (+) \text{ arg arg } 8,$
 $\text{sto } [\text{clo} + 0] \text{ arg},$
 $\text{sto } [\text{clo} + 1] \text{ env},$
 $\text{bop } (+) \text{ sp sp } 1,$
 $\text{sto } \langle \text{sp} - 1 \rangle \text{ clo},$
 $\text{bop } (+) \text{ arg arg } (16 + \text{length } c),$
 $\text{jmp arg},$
 $\text{bop } (+) \text{ sp sp } 1,$
 $\text{sto } \langle \text{sp} - 1 \rangle \text{ env},$
 $\text{ld env } 2,$
 $\text{alloc env env},$
 $\text{sto } [\text{env} + 0] \text{ clo},$
 $\text{ld aux}_1 [\text{clo} + 1],$
 $\text{sto } [\text{env} + 1] \text{ aux}_1,$
 $\text{ld clo } 2,$
 $\text{alloc clo clo},$
 $\text{sto } [\text{clo} + 0] \text{ arg},$
 $\text{sto } [\text{clo} + 1] \text{ env},$
 $\text{ld env clo},$
 $c,$
 $\text{ld arg } \langle \text{sp} - 1 \rangle,$
 $\text{bop } (-) \text{ sp sp } 2,$
 $\text{ld env } \langle \text{sp} + 0 \rangle,$
 jmp ret

$\left. \begin{array}{l} \text{active part} \\ \text{inactive part} \end{array} \right\}$

$ezwickel(e_1 e_2) n \Gamma :=$
 $\text{let } c_1 := ezwickel e_1 n \Gamma \text{ in}$
 $\text{let } c_2 := ezwickel e_2 (\text{length } c_1 + n) \Gamma \text{ in}$
 $c_1, c_2,$
 $\text{bop } (-) \text{ sp sp } 1,$
 $\text{ld clo } \langle \text{sp} - 1 \rangle,$
 $\text{sto } \langle \text{sp} - 1 \rangle \text{ ret},$
 $\text{ld arg } \langle \text{sp} + 0 \rangle,$
 $\text{lpc ret},$
 $\text{bop } (+) \text{ ret ret } 3,$
 $\text{jmp } [\text{clo} + 0],$
 $\text{ld ret } \langle \text{sp} - 1 \rangle,$
 $\text{sto } \langle \text{sp} - 1 \rangle \text{ arg}$

Figure 5.19: Code generation for expressions (continued).

Definition 56 (Local world w for the verification of Zwickel).

$$\begin{aligned}
w &\in \text{LWorld}^{\Omega_{\mathcal{T}\mathcal{S}}.\mathsf{T}} \\
w.\mathsf{T.S} &:= \text{Stack} \times \text{Heap} \\
w.\mathsf{T}.\sqsupseteq &:= \{ ((q', h'), (q, h)) \mid h' \supseteq h \} \\
w.\mathsf{T}.\sqsupseteq_{\text{pub}} &:= \{ ((q', h'), (q, h)) \mid h' \supseteq h \wedge q' = q \} \\
w.\mathsf{C}(G)(s_g, (q, h)) &:= \{ (c_a, c_b) \mid c_a = (\emptyset, \emptyset, q, h) \wedge c_b = (\emptyset, \emptyset, \emptyset) \} \\
w.\mathsf{N.NS} &:= \emptyset \\
w.\mathsf{N.NR} &:= \emptyset \\
w.\mathsf{O} &:= \text{gwf}
\end{aligned}$$

The only differences are due to this being a local world for $\lesssim_{\mathcal{T}\mathcal{S}}$ rather than for $\lesssim_{\mathcal{T}\mathcal{I}}$. In particular, since we are now working in the typed version of PILS, we must provide a set of type names $w.\mathsf{N.NS}$ and a relational interpretation $w.\mathsf{N.NR}$ for them—we choose the empty set. Otherwise, the structure of w is the same as before, although the contents of the stack $s.q$ will look rather different in our Zwickel proofs than they did for Pilsner.

Analogous to Definition 51 in Section 5.10, we define an auxiliary relation $\preceq_{\mathcal{T}\mathcal{S}}$ between machine code parameterized over its load address (f) and its source expression (e), relative to a type and typing context. Similar to before, m is the address at which we load the code $f\ m$, and the only f in which we are interested is $\lambda n. \text{map } \mathbb{E}(\text{ezwickel } e\ n \mid \Gamma)$.

Definition 57.

$$\begin{aligned}
\boxed{\Gamma \vdash f \preceq_{\mathcal{TS}} e : \tau} &:= \exists i. \forall G, s_0, s, \delta, \psi_a, \sigma_a, \sigma_b, m, n, j, \tau'. \\
&G \in \mathbf{GK}_{w\uparrow} \wedge \\
&s \supseteq s_0 \wedge \\
&m \neq 0 \wedge \\
&(\exists h. s.h = [m \mapsto f\ n] \sqcup [m - n \mapsto \text{map snd } \psi_a] \sqcup h) \wedge \\
&(\forall F: \tau' \in \Gamma. \exists v_a, v_b. \psi_a(F) = v_a \wedge \sigma_b(F) = v_b \wedge (v_a, v_b) \in \overline{G}(s)(\delta\tau') \wedge \\
&\text{map fst } \Gamma|_{\text{LbI}} = \text{map fst } \psi_a \wedge \\
&\text{repr } s.h (s.R(\text{env})) (\text{map snd } \sigma_a) \wedge \\
&(\forall y: \tau' \in \Gamma. \exists v_a, v_b. \sigma_a(y) = v_a \wedge \sigma_b(y) = v_b \wedge (v_a, v_b) \in \overline{G}(s)(\delta\tau')) \wedge \\
&\text{map fst } \Gamma|_{\text{TVar}} = \text{map fst } \sigma_a \wedge \\
&\mathbf{K} G s_0 s K_a K_b j (\delta\tau) \tau' (\text{length } (f\ n) + m) K \\
\implies &((m, \emptyset, \emptyset, \emptyset), (\emptyset, \emptyset, K[\sigma_b e])) \in \mathbf{E}_{w\uparrow}(i + j)(K_a, K_b)(G)(s_0)(s)(\perp)(\tau')
\end{aligned}$$

$$\begin{aligned}
\boxed{\mathbf{K} G s_0 s K_a K_b j \tau \tau' m K} &= \forall G', s', v_a, v_b. \\
&G' \supseteq G \wedge \\
&G' \in \mathbf{GK}_{w\uparrow} \wedge \\
&s' \supseteq s \wedge \\
&s'.R(\text{env}) = s.R(\text{env}) \wedge \\
&s'.R(\text{ret}) = s.R(\text{ret}) \wedge \\
&s'.R(\text{sp}) = s.R(\text{sp}) + 1 \wedge \\
&s'.q = [s.R(\text{sp}) \mapsto v_a] \sqcup s.q \wedge \\
&(v_a, v_b) \in \overline{G'}(s')(\tau) \\
\implies &((m, \emptyset, \emptyset, \emptyset), (\emptyset, \emptyset, K[v_b])) \in \mathbf{E}_{w\uparrow}(j)(K_a, K_b)(G')(s_0)(s')(\perp)(\tau')
\end{aligned}$$

The key difference to $\preceq_{\mathcal{TI}}$ concerns the last condition, which, earlier, described how the top-most stack cells correspond to whatever σ_b provides for e 's continuation variables. Here, however, \mathbf{K} relates the evaluation context K (in which a closed instance of e is being executed) to the machine code that gets executed after the code for $f\ n$. This code starts at address $m + \text{length } (f\ n)$, *i.e.*, it is simply whatever comes next in memory (cf. the definition of *ezwickel*). \mathbf{K} 's definition states that, under some assumptions, this code is \mathbf{E} -related to the filling of K with some value v_b . It quantifies over an extended global knowledge and a future state s' , whose register file agrees with s on `ret` and `env`. This is important, for instance, in the case of functions, whose code relies on `ret` having retained its original value when the function epilogue jumps to it (cf. Figure 5.19). Moreover, \mathbf{K} assumes that the `sp` register has been increased by one and, correspondingly, the stack extended with a new value v_a . Finally, this value must be related to v_b at s' .

Lemma 80.

$$\frac{\Gamma \vdash e : \tau}{\Gamma \vdash (\lambda n. \text{map } \mathbb{E}(\text{ezwickel } e\ n\ |\Gamma|)) \preceq_{\mathcal{TS}} e : \tau}$$

Theorem 23 (Correctness of Zwickel).

$$\frac{\Gamma \vdash M : \Gamma'}{\Gamma \vdash \text{zwickel } M \mid \Gamma \mid \lesssim_{\mathcal{TS}} M : \Gamma' \quad \vdash \text{zwickel } M \mid \Gamma \mid : |\Gamma'|}$$

5.13 The Self-Modifying Awkward Example

As we have now seen, the PILS system from Chapter 4 can be used to verify multiple different compilers with the same source and target languages (namely at least Pilsner and Zwickel). In this section we further demonstrate PILS’ flexibility by reporting on the proof of the challenging refinement from Hur and Dreyer [32] mentioned in Section 1.3, which relies on tricky manipulations of local state, far more involved than those of any traditional compiler.

This example is based on Pitts and Stark’s “awkward” example, which have visited in Section 2.2 and repeat here for convenience:

$$\begin{aligned} \tau &= (\text{unit} \rightarrow \text{unit}) \rightarrow \text{int} \\ e_1 &= \lambda f. (f \langle \rangle; 1) \\ e_3 &= \text{let } x = \text{ref } 0 \text{ in } \lambda f. (x := 1; f \langle \rangle; !x) \end{aligned}$$

Recall that both expressions evaluate to higher-order functions that, when applied, call the argument “callback” function f and then return a number. In e_1 this number is simply 1. In e_3 it is the result of dereferencing a local (private) reference x , which is initialized to 0. Notice, though, that when e_3 is called for the first time, it immediately writes 1 to x . Since there are no other writes, the value of x returned at the end will always be 1 as well.

Hur and Dreyer [32] adapt this example by substituting for e_1 a tricky self-modifying machine program that implements the same behavior, but in a rather baroque way. Figure 5.20 shows what our version⁷ of this program looks like in memory. It is parameterized by the load address n , and $\mathbb{E}(-)$ denotes the encoding of an instruction as a machine word. Notice that part of the code has been “encrypted” by adding 666 to its encoding. We briefly explain how the code works.

The first few lines allocate a new function closure with empty environment and code pointer $n + 5$, and return it to the context. When this function gets called the first time, it starts out by decrypting the encrypted instructions (offsets 5–11), thus replacing the encrypted code in memory. Subsequently (offsets 12–14), it replaces its first instruction by a direct jump in order to skip over the decryption loop in future executions. The remaining code (offsets 15–23), which is also the target of that jump, simply performs the callback function call and then returns 1.

Hur and Dreyer showed that this contrived implementation refines the high-level program e_3 as a demonstration that their KLR approach is flexible enough to reason about semantically involved “transformations”, even ones whose correctness relies on low-level internal state changes that clearly have no high-level counterpart. By

⁷We have slightly modified the original one to account for a difference in calling convention.

<i>offset</i>	<i>word</i>	<i>comment</i>
$n + 0$	$\mathbb{E}(\text{ld arg } 1)$ $\mathbb{E}(\text{alloc arg arg})$ $\mathbb{E}(\text{ld aux}_1 \underline{n + 5})$ $\mathbb{E}(\text{sto } [\text{arg} + 0] \text{ aux}_1)$ $\mathbb{E}(\text{jmp ret})$	create and return closure
$n + 5$	$\mathbb{E}(\text{ld aux}_2 \underline{n + 11})$ $\mathbb{E}(\text{ld aux}_1 [\text{aux}_2 + 0])$ $\mathbb{E}(\text{bop} - \text{aux}_1 \text{ aux}_1 666)$ $\mathbb{E}(\text{sto } [\text{aux}_2 + 0] \text{ aux}_1)$ $\mathbb{E}(\text{bop} + \text{aux}_2 \text{ aux}_2 1)$	(closure begins here) decrypt in loop
$n + 10$	$\mathbb{E}(\text{bop} - \text{aux}_1 \underline{n + 24} \text{ aux}_2)$ $666 + \mathbb{E}(\text{jnz aux}_1 \underline{n + 6})$ $666 + \mathbb{E}(\text{ld aux}_1 \underline{\mathbb{E}(\text{jmp } n + 15)})$ $666 + \mathbb{E}(\text{ld aux}_2 \underline{n + 5})$ $666 + \mathbb{E}(\text{sto } [\text{aux}_2 + 0] \text{ aux}_1)$	update code
$n + 15$	$666 + \mathbb{E}(\text{sto } \langle \text{sp} + 0 \rangle \text{ ret})$ $666 + \mathbb{E}(\text{bop} + \text{sp sp } 1)$ $666 + \mathbb{E}(\text{ld ret } \underline{n + 20})$ $666 + \mathbb{E}(\text{sto clo arg})$ $666 + \mathbb{E}(\text{jmp } [\text{clo} + 0])$	save and set return address call function
$n + 20$	$666 + \mathbb{E}(\text{bop} - \text{sp sp } 1)$ $666 + \mathbb{E}(\text{ld ret } \langle \text{sp} + 0 \rangle)$ $666 + \mathbb{E}(\text{ld arg } 1)$ $666 + \mathbb{E}(\text{jmp ret})$	clean up and return 1

Figure 5.20: Self-modifying “awkward” example

verifying the same example (with respect to $\lesssim_{\mathcal{TS}}$), we aim to demonstrate that PILS are equally flexible.

Fortunately (and, as should be clear by now, not coincidentally), the high-level structure of our proof closely follows that of Hur and Dreyer’s proof. Let us here merely sketch the local world that we use. Full details can be found in the Coq development.

The \mathcal{S} module that we are interested in is a singleton module whose only function is a thunked e_3 , *i.e.*, $\lambda_. e_3$. The \mathcal{T} module under consideration is a singleton module whose function has the “code” shown in the figure. A state s of our local world consists of the following:

1. The load address of the machine module. Given how loading works, this address determines the function value (*i.e.*, the heap address at which the code pointer is stored), the code address itself, and the precise memory contents at that address (*i.e.*, the code). In order for the latter to be uniquely determined, the state also records...
2. ...whether or not the code has already been decrypted. This has an all-or-nothing semantics, meaning that we do not have states representing the intermediate stages of the memory where some instructions have been decrypted but others still remain encrypted. Such states are not necessary for the proof since no functions get called during decryption.
3. The part of the machine stack occupied by not-yet returned instances of the self-modifying function. It consists entirely of return addresses, one per invocation.
4. Concerning the \mathcal{S} module, the state records the allocated memory, *i.e.*, all references x (created by calling the exported function) and their current values (either 0 or 1)—in other words, the local heap.

Let’s look at the possible state transitions. The load address (of the machine module) remains the same in all future states. The encryption status can change, but only from encrypted from unencrypted, after which it will remain unencrypted in all future states. The local machine stack can change arbitrarily in private, but must remain unchanged in public. Finally, the local source heap can change in two dimensions: new references can be added (when the exported wrapper function gets called), and existing references can change their contents from 0 to 1 (when one of closures gets called for the first time).

5.14 Mechanization and Extraction

Even though having our model be somewhat language-generic helped us avoid a lot of duplicate work, the mechanization effort was huge and so is the actual Coq code, consisting of roughly 13,000 lines of definitions and 24,000 lines of proofs.

In general, proofs done directly in the model tend to be very tedious—in particular when the assembly language is involved, such as the example sketch in Section 4.6. There is certainly room for more sophisticated automation here, but ultimately we would like to develop better abstractions and reasoning principles on top of the model.

From the start we formalized our work on PILS in Coq. This was critical for gaining high confidence in the results. Given the amount of technical details involved, we consider it impractical to carefully write and check many of the proofs by hand. Of course it is also immensely helpful that, after having modified some definition, Coq points one to all places that need to be updated.

Besides obviously being a very time-demanding task, the formalization also posed several other challenges. The compilation (including proof checking) of the Coq development takes very long (over 1 hour for the whole project). This seems to be not only due to the large number of definitions and proofs but also due to the large size of some proof terms. In fact we had to break up some large proofs in unnatural ways just to avoid intolerable slowdowns.

To work more productively, we wrote a script that would first compile all files with (most) proofs erased, in order to more quickly produce the “.vo” files needed for interactive development. After that, a regular build would run. We welcome the recent work on parallelizing compilation in Coq, which implements a similar idea and other features aimed at speeding up compilation and workflow [86].

Our Coq development contains a script that extracts Pilsner, Zwickel, and the linker as OCaml code, and couples the compilers with code for parsing command-line arguments as well as a lexer and parser for the source language. In order to execute target machine code, we have implemented a single-step interpretation function in Coq and proved that it conforms to the operational semantics. This function is also extracted to OCaml and wrapped in a loop. Please see the “README.txt” file for further details and instructions on how to build and run the programs.

5.15 Putting It All Together

In Section 5.11 we have established the correctness of Pilsner w.r.t. $\lesssim_{\mathcal{TS}}$, and in Section 5.12 also that of Zwickel. As discussed in Section 5.13, we have moreover proven the correctness of a manual “translation” of the awkward example to a self-modifying \mathcal{T} machine program.

Thanks to the modularity of PILS, specifically that of $\lesssim_{\mathcal{TS}}$ (Theorem 14), this immediately means that we can link together the self-modifying machine program module with any Pilsner-produced and/or any Zwickel-produced code, and that doing so results in a \mathcal{T} program that refines the linking of the corresponding sources.

Our Coq development contains one simple example of this, involving a Zwickel-compiled main function that acts as client of the awkward function and of a Pilsner-compiled factorial function. Please see the “README.txt” file for details.

Chapter 6

Related Work

Contents

6.1	Logical Relations	241
6.2	Bisimulations	243
6.3	Compositional Compiler Correctness	244
6.4	Miscellaneous	246

6.1 Logical Relations

The body of work on applying and extending logical relations is vast. We can only give a few pointers here.

Early days The idea of logical relations is over 50 years old. Often quoted as one of the earliest use of the technique is Tait’s proof of strong normalization for the typed lambda calculus [84]. Other early work includes that of Plotkin and Statman, who investigated lambda-definability [66, 67, 76].

Parametricity Girard [26, 27, 25] and Reynolds [69] generalized logical relations to the polymorphic lambda calculus, a key insight being the quantification over relational interpretations of a program’s abstract types. Reynolds introduced *parametricity*, which led to further work on reasoning about abstract types [53, 54, 90, 36, 60].

State Although they were originally geared toward reasoning about pure λ -calculi, logical relations have been successfully generalized to reason about state. In Pitts and Stark’s seminal work on *Kripke logical relations (KLRs)* [65], logical relations are indexed by *possible worlds*, which characterize the runtime environment (*e.g.*, the assumptions about heaps) under which two programs are considered to be equivalent.

These early KLRs for reasoning about local state imposed serious restrictions on memory contents by allowing references only to integers [65] or, say, to references of

integers [8, 68]. The first to deal with full higher-order store were Birkedal *et al.* [16, 14].

In more recent work on KLRs, Dreyer *et al.* [4, 22] showed how to generalize Pitts and Stark’s technique to reason about (1) modules whose correctness proofs require fine-grained control over how local state evolves over time, and (2) ML-like languages with higher-order state. W.r.t. point (1), they model possible worlds as *state transition systems (STSs)*, as we have reviewed in Chapter 2. PBs (and PILS) adopt Dreyer *et al.*’s STS technique directly, and thus it is relatively straightforward to port all the $F^{\mu!}$ equivalence proofs given in their papers from using KLRs to using PBs.

W.r.t. point (2), the challenge of supporting higher-order state in Kripke logical relations is that a naive attempt to construct a model of general reference types leads to a circularity. Intuitively, ℓ_1 and ℓ_2 are related at $\text{ref } \tau$ under a possible world W iff W encodes the invariant that the heaps of the two programs map ℓ_1 and ℓ_2 to values v_1 and v_2 that are logically related at type τ . But how can the logical relation be indexed by a possible world W , which itself is defined in terms of the logical relation? If τ is restricted to base type (*e.g.*, int), there’s no issue because the logical relation at int is simply the identity relation, but at higher type we have a problem.

Dreyer *et al.* handle higher-order state by means of Appel, McAllester, and Ahmed’s technique of *step-indexed logical relations (SILRs)* [5, 2]. That is, they cut the aforementioned semantic circularity by indexing the model by a natural number (“step index”) k , which represents the number of steps left on “the clock” and which gets decremented every time around the cycle between logical relations and possible worlds. As discussed in Chapter 2, it seems fundamentally difficult to compose SILR proofs transitively.

PBs and PILS employ the idea of global knowledge in order to avoid the need for step-indexing in modeling higher-order state. Specifically, by parameterizing the heap relations in our worlds over the global knowledge G , we give heap invariants a way of referring to the global value equivalence, which is essentially what the step-indexed stratification of Kripke worlds is trying to achieve as well. Our method does not bake in any syntactic typing assumptions.

Dreyer *et al.* [22] also proposed the distinction between private and public transitions to prove equivalences that only hold in the absence of control features such as call/cc. Orthogonally, they showed how equivalences that only hold when restricting store to first order can be handled by allowing a kind of backtracking in the STSs.

Our notions of logical reduction steps and stuttering from Sections 3.10 and 4.5.1 are reminiscent of Svendsen *et al.*’s *transfinite step-indexing* [83], which relaxes the coupling of a logical relation’s step-index to the physical reduction steps of related programs. This enables proofs where one can decrease the step-index by an arbitrary finite amount even when there is only a single corresponding physical reduction step.

Logics Several logics based on logical relations models have been developed [23, 21, 87], with increasing power and complexity. They attempt to abstract away tech-

nical details of the models (*e.g.*, the step-indexing) and offer higher-level reasoning principles.

6.2 Bisimulations

Aside from their general coinductive flavor, PBs and PILS are closely related to two different bisimulation techniques.

From *normal form* (or *open*) bisimulations [71, 43, 78, 44, 45], we take the idea of treating unknown equivalent functions as black boxes. In particular, our expression equivalence relation \mathbf{E} , which deals explicitly with the possibility (in its third disjunct) that related terms may get stuck by calling unknown functions, is *highly* reminiscent of the formulation of normal form bisimulations. The main difference is that we express the notion of “stuckness” semantically, via the global knowledge parameter G , whereas normal form bisimulations express it syntactically by requiring related stuck terms to share a common head variable.

Normal form bisimulations draw much inspiration from game-semantics models [57], and our distinction between global and local knowledge has a seemingly gamey flavor as well. We leave a deeper study of the connection to game semantics to future work.

Sumii *et al.*’s *environmental* bisimulations (aka “relation-sets bisimulations”) are perhaps the most powerful form of bisimulation yet developed for ML-like languages [63, 81, 41, 74, 79]. As the latter name suggests, these bisimulations are not term relations, but *sets* \mathcal{X} whose elements are themselves term relations R (possibly paired with some additional environmental information, such as knowledge about the state of the heap). In essence, each $R \in \mathcal{X}$ defines some piece of “local knowledge” (following our terminology) about program equivalence. In order to show \mathcal{X} to be a bisimulation, one must check that for all $R \in \mathcal{X}$, uses of terms related by R will never result in observably different outcomes and will always produce values that are related by some $R' \in \mathcal{X}$ s.t. $R' \supseteq R$.

Viewed in terms of PBs, one can understand an environmental bisimulation \mathcal{X} as effectively defining an abstract state space, with each $R \in \mathcal{X}$ as a distinct state. However, the accessibility (transition) relation between these states is essentially baked in: roughly speaking, a term relation R' is (publicly) accessible from another term relation R if $R' \supseteq R$. Thus, environmental bisimulations provide less control over the structure of the transition system than PBs do, and they do not support anything directly analogous to the distinction between public and private transitions.

As a consequence, environmental bisimulations are most effective at proving equivalences that require transition systems with only public transitions (*e.g.*, the twin abstraction example), and their proofs for examples where private transitions are required (*e.g.*, the well-bracketed state change example) are comparatively “brute-force”. It is an open question whether environmental bisimulations can be generalized to support the full power of PBs and PILS with both public and private transitions.

Our approach to reasoning about parametricity of ADTs, by populating the local

knowledge of a world with relations at abstract *type names*, is inspired directly by Sumii and Pierce [81].

Well-Founded Bisimulations Our technique of logical reduction is inspired by Namjoshi’s *well-founded* [58, 48] bisimulations, which were developed as an alternative formulation of *stuttering* bisimulations [17] that can be checked by only reasoning about single transitions instead of infinite computations. In order to support finite but unbounded stuttering, well-founded bisimulations employ a “rank function” mapping states to some well-founded ordering, and insist (roughly) that, for states related in a bisimulation, either both make physical transitions to related states or else one side makes a transition while the rank of the pair of states decreases. In our model, we use the stutter budget to effectively bake a particular rank function into our bisimulations, which is sufficient for our purposes and convenient to work with. As far as we aware, this is the first time that the idea of well-founded bisimulations has been adapted for use in reasoning about open programs in a higher-order language setting.

6.3 Compositional Compiler Correctness

Research on compiler correctness has a long history. Here we focus only on compositionality, referring to Dave’s extensive bibliography [20] for the broader area.

Using Logical Relations Benton and Hur [6] proposed the idea of using logical relations to define compositional semantics preservation. In addition to being inherently modular, logical relations are highly flexible, having been used in the past as an effective technique for proving correctness of a wide variety of program transformations in a wide variety of languages [4, 22]. Moreover, unlike contextual refinement, logical relations can be used to relate different source and target languages.

Hur and Dreyer [32] developed this idea further by formalizing the compositional correctness of a simple, single-pass compiler from an ML-like source language to an idealized assembly language. They additionally demonstrated the flexibility of their inter-language logical relations by using them to verify a contrived but illustrative example, wherein a higher-order ML function was implemented in a rather baroque way by some tricky hand-written *self-modifying* assembly code. Thanks to the modularity of their logical relations method, this highly non-standard assembly code could nonetheless be safely linked with assembly modules produced by their verified compiler, with the resulting assembly program guaranteed to preserve the semantics of the corresponding linked source modules.

Unfortunately, it is not clear how to scale Hur *et al.*’s approach from single- to multi-pass compilers because, although logical relations are modular and flexible, they are not typically transitive.

Our language-generic style of defining PILS is very much inspired by Hur and Dreyer’s generic KLR setup [32]. Their point, however, was mainly to “simplify and

clarify the formal presentation”—they only ever instantiated their definitions with a single language pair and proved only a few very basic properties about the generic model; all the rest was proven about the concrete instantiation.

Multi-language semantics. Motivated by the goal of supporting compiler verification for programs that interoperate between different languages, Perconti and Ahmed [62] propose an approach based on *multi-language semantics* [49]. In particular, they define a “big-tent” language that comprises the source, target, and intermediate languages of a compiler, and provides “wrapping” operations for embedding terms of each language within the others. They then use logical relations to prove that every source module is contextually equivalent to a suitably wrapped version of the target module to which it is compiled. In this way, their method synthesizes the benefits of logical relations (modularity and different source and target languages) and contextual equivalence (transitivity).

One downside of their approach is that the intermediate languages (ILs) used in a compiler show up explicitly in the statement of compiler correctness. This leads to a loss of flexibility: the semantics of source-level linking is not preserved when linking the results of compilers that have different ILs. Another limitation with respect to flexibility is that their approach seems to be restricted to compilers that use *typed* intermediate and assembly languages, and has only so far been applied to a purely functional source language. On the other hand, Perconti and Ahmed are more flexible than we are with respect to multi-language interoperation. One of their explicit goals is to reason about the linking of ML code with *arbitrary* typed assembly code, whereas we only support verified linking with assembly modules that refine *some* source-level counterpart. As we observed in footnote 1, we do not believe this is a fundamental limitation of our approach: it should in principle be possible to develop PILS for a different source language in which high- and low-level modules may interoperate, in which case the “source”-level specification of a “target”-level module could be the target-level module itself.

In the above work, the source and intermediate languages are purely functional, and the target language is still very high-level. In recent follow-up work, Patterson *et al.* [61] develop a multi-language semantics that provides safe interoperability between the former source language and a truly low-level typed assembly language. They also construct a logical relation that supports compositional reasoning, but they do not define a compiler.

Compositional verified compilation for C. Motivated by the goal of compositional compiler verification, Beringer *et al.* [9, 77] propose an adaptation of the CompCert framework based on a novel “interaction” semantics that differentiates between internal (intra-module) and external (inter-module) function calls. They introduce a notion of “structured simulation” that assumes little about the memory transformations performed by external function calls.

Beringer *et al.*’s approach is transitive, and like Perconti and Ahmed’s (but unlike

ours), it supports verified compilation of multi-language programs—in this case, programs that link C and assembly modules. However, also like Perconti and Ahmed’s approach, Beringer *et al.*’s is somewhat lacking in flexibility. It depends on compiler passes only performing a restricted set of memory transformations—additional transformations could potentially break the transitivity property. In addition, their method appears to be geared specifically toward compilers à la CompCert, which employ a uniform memory model across source, intermediate, and target languages. It is not clear how to generalize their technique to support richer (*e.g.*, ML-like) source languages, or compilers whose source and target languages have different memory models.

In recent work, Kang *et al.* [37] adapt CompCert such that it supports separate compilation when all modules are compiled with the same compiler. This is clearly weaker than our horizontal compositionality. However, the restriction to a single compiler makes it possible to phrase the compositional verification in terms of the existing whole-program verification and thus reuse most of the proofs. While carrying out the work, Kang *et al.* discovered two bugs in CompCert: an incorrect axiom and an analysis that becomes invalid in the presence of linking.

Wang *et al.* [91] have also recently explored compositional compiler verification for a restricted C-like language called Cito. Their approach embeds the verification statement within a Hoare logic for partial correctness of assembly modules, thus enabling support for verified cross-language linking, but without guaranteeing preservation of termination behavior. Further work is needed to better understand the relationship between this approach and traditional refinement-based compiler verification.

6.4 Miscellaneous

Large vs. Small Worlds While PBs (and PILS) build very closely on the state transition systems in Dreyer *et al.*’s KLRs [4, 22], there is a big difference between them, which we like to think of in terms of *large* vs. *small* worlds.

Under Dreyer *et al.*’s approach, in order to demonstrate the equivalence of functions f_1 and f_2 under a “possible world” W , one proves that they behave the same when passed arguments that are related under any “future world” W' of W , which may contain arbitrary new invariants concerning the local state of other modules in the program. One can really think of the “future world” relation (*i.e.*, the Kripke structure) as defining its own transition system (or *large* world), with the possible worlds W as its states.

In contrast, our PBs rely only on *small* worlds. For us, worlds W are static entities that contain only the local invariants relevant to the module we are reasoning about, and nothing about any invariants for other parts of the program. In proving equivalence of functions f_1 and f_2 under W , we never quantify over any future worlds that extend W . Of course, in order to support compositional reasoning—*i.e.*, in order to show that consistency of worlds is preserved under separating conjunction—we must show that f_1 and f_2 behave the same when applied to arguments drawn from

some larger relation than just W 's local knowledge; but for that purpose we quantify over the global knowledge G , which is not a world, but rather an arbitrary extension of W 's local knowledge.

These different accounts of worlds are strongly reminiscent of the different techniques that have been proposed for modeling resource invariants in logics of storable locks. Gotsman *et al.* [30] and Hobor *et al.* [31] presented, roughly contemporaneously, two different models of a concurrent separation logic for local reasoning about programs that dynamically allocate locks and store them in the heap. The central challenge in developing such a model is in dealing with the semantic circularity that arises when accounting for locks whose resource invariants are essentially recursive.

Gotsman *et al.* deal with this circularity syntactically, by assuming a static set of named “sorts” of resource invariants, which includes not all possible invariants, but only those needed for reasoning about a particular program. In contrast, Hobor *et al.* (and more recently, Buisse *et al.* [18]) deal with the circularity head-on, defining once and for all what recursive resource invariants *mean* using step-indexing. The latter is analogous to Dreyer *et al.*'s “large worlds” approach, which defines the space of all possible heap invariants, while the former is analogous to our “small world” approach of defining only the heap invariants needed within the module we are reasoning about.

Our small-world relations seem easier to compose transitively, precisely because we make no assumption whatsoever about the relatedness of functions defined outside of whatever module we are reasoning about. That is, the global knowledge G that we quantify over (*e.g.*, when proving world consistency) could include complete garbage, and our transitivity proofs rely in a fundamental way on the surgical insertion of contentful garbage into the global knowledge.

Bibliography

- [1] Samson Abramsky. The lazy lambda calculus. In D. A. Turner, editor, *Research Topics in Functional Programming*, pages 65–117. 1990.
- [2] Amal Ahmed. *Semantics of Types for Mutable State*. PhD thesis, Princeton University, 2004.
- [3] Amal Ahmed. Step-indexed syntactic logical relations for recursive and quantified types. In *ESOP*, 2006.
- [4] Amal Ahmed, Derek Dreyer, and Andreas Rossberg. State-dependent representation independence. In *POPL*, 2009.
- [5] Andrew Appel and David McAllester. An indexed model of recursive types for foundational proof-carrying code. *TOPLAS*, 23(5):657–683, 2001.
- [6] Nick Benton and Chung-Kil Hur. Biorthogonality, step-indexing and compiler correctness. In *ICFP*, 2009.
- [7] Nick Benton and Vasileios Koutavas. A mechanized bisimulation for the nu-calculus. *Journal of Higher Order and Symbolic Computation*, 2013.
- [8] Nick Benton and Benjamin Leperchey. Relational reasoning in a nominal semantics for storage. In *TLCA*, 2005.
- [9] Lennart Beringer, Gordon Stewart, Robert Dockins, and Andrew W. Appel. Verified compilation for shared-memory C. In *ESOP*, 2014.
- [10] L. Birkedal, N. Torp-Smith, and H. Yang. Semantics of separation-logic typing and higher-order frame rules for Algol-like languages. *LMCS*, 2(5:1), 2006.
- [11] Lars Birkedal and Aleš Bizjak. A note on the transitivity of step-indexed logical relations. Manuscript, November 2012.
- [12] Lars Birkedal and Robert W. Harper. Constructing interpretations of recursive types in an operational setting. *Information and Computation*, 155:3–63, 1999.
- [13] Lars Birkedal, Kristian Støvring, and Jacob Thamsborg. Realizability semantics of parametric polymorphism, general references, and recursive types. In *FOSSACS*, 2009.

- [14] Lars Birkedal, Kristian Støvring, and Jacob Thamsborg. Relational parametricity for references and recursive types. In *Proceedings of the 4th International Workshop on Types in Language Design and Implementation*, TLDI '09, pages 91–104, New York, NY, USA, 2009. ACM.
- [15] Aaron Bohannon, Jeffrey A. Vaughan, and Benjamin C. Pierce. Relational lenses: A language for updateable views. In *Principles of Database Systems (PODS)*, 2006. Extended version available as University of Pennsylvania technical report MS-CIS-05-27.
- [16] Nina Bohr. *Advances in Reasoning Principles for Contextual Equivalence and Termination*. PhD thesis, IT University of Copenhagen, 2007.
- [17] M. C. Browne, E. M. Clarke, and O. Grümberg. Characterizing finite Kripke structures in propositional temporal logic. *Theor. Comput. Sci.*, 59(1-2):115–131, July 1988.
- [18] Alexandre Buisse, Lars Birkedal, and Kristian Støvring. A step-indexed Kripke model of separation logic for storable locks. In *MFPS*, 2011.
- [19] The Coq Proof Assistant. <http://coq.inria.fr/>.
- [20] Maulik A. Dave. Compiler verification: A bibliography. *SIGSOFT Softw. Eng. Notes*, 28(6):2–2, November 2003.
- [21] Derek Dreyer, Amal Ahmed, and Lars Birkedal. Logical step-indexed logical relations. *LMCS*, 7(2:16):1–37, June 2011.
- [22] Derek Dreyer, Georg Neis, and Lars Birkedal. The impact of higher-order state and control effects on local relational reasoning. *JFP*, 22(4-5), 2012.
- [23] Derek Dreyer, Georg Neis, Andreas Rossberg, and Lars Birkedal. A relational modal logic for higher-order stateful ADTs. In *POPL*, 2010.
- [24] Matthew Fluet and Stephen Weeks. Contification using dominators. In *ICFP*, 2001.
- [25] Jean H Gallier. On Girard’s “candidats de reductibilité”. Manuscript, 1989.
- [26] Jean-Yves Girard. *Interprétation fonctionnelle et élimination des coupures de l’arithmétique d’ordre supérieur*. PhD thesis, Université Paris VII, 1972.
- [27] Jean-Yves Girard, Yves Lafont, and Paul Taylor. *Proofs and Types*, volume 7 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 1989.
- [28] Andrew D. Gordon. Bisimilarity as a theory of functional programming. *Theor. Comput. Sci.*, 228(1-2):5–47, October 1999.

- [29] Andrew D. Gordon and Gareth D. Rees. Bisimilarity for a first-order calculus of objects with subtyping. In *Proceedings of the 23rd ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, POPL '96, pages 386–395, New York, NY, USA, 1996. ACM.
- [30] Alexey Gotsman, Josh Berdine, Byron Cook, Noam Rinetzkyl, and Mooly Sagiv. Local reasoning about storable locks and threads. In *APLAS*, 2007.
- [31] Aquinas Hobor, Andrew Appel, and Francesco Zappa Nardelli. Oracle semantics for concurrent separation logic. In *ESOP*, 2008.
- [32] Chung-Kil Hur and Derek Dreyer. A Kripke logical relation between ML and assembly. In *POPL*, 2011.
- [33] Chung-Kil Hur, Derek Dreyer, Georg Neis, and Viktor Vafeiadis. The marriage of bisimulations and Kripke logical relations. In *POPL*, 2012.
- [34] Chung-Kil Hur, Georg Neis, Derek Dreyer, and Viktor Vafeiadis. The power of parameterization in coinductive proof. In *POPL*, 2013.
- [35] Chung-Kil Hur, Georg Neis, Derek Dreyer, and Viktor Vafeiadis. A logical step forward in parametric bisimulations. Technical Report MPI-SWS-2014-003, MPI-SWS, 2014.
- [36] Patricia Johann and Janis Voigtländer. The impact of *seq* on free theorems-based program transformations. *Fundamenta Informaticae*, 69(1–2):63–102, 2006.
- [37] Jeehoon Kang, Yoonseung Kim, Chung-Kil Hur, Derek Dreyer, and Viktor Vafeiadis. Lightweight verification of separate compilation. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '16, pages 178–190, New York, NY, USA, 2016. ACM.
- [38] Andrew Kennedy. Relational parametricity and units of measure. In Peter Lee, Fritz Henglein, and Neil D. Jones, editors, *POPL*, pages 442–455. ACM Press, 1997.
- [39] Andrew Kennedy. Compiling with continuations, continued. In *ICFP*, 2007.
- [40] Vasileios Koutavas, Paul Blain Levy, and Eijiro Sumii. From applicative to environmental bisimulation. In *MFPS*, 2011.
- [41] Vasileios Koutavas and Mitchell Wand. Small bisimulations for reasoning about higher-order imperative programs. In *POPL*, 2006.
- [42] Ramana Kumar, Magnus Myreen, Michael Norrish, and Scott Owens. CakeML: A verified implementation of ML. In *POPL*, 2014.
- [43] Soren Lassen. Eager normal form bisimulation. In *LICS*, 2005.

- [44] Søren B. Lassen and Paul Blain Levy. Typed normal form bisimulation. In *CSL*, 2007.
- [45] Søren B. Lassen and Paul Blain Levy. Typed normal form bisimulation for parametric polymorphism. In *LICS*, 2008.
- [46] Vu Le, Mehrdad Afshari, and Zhengdong Su. Compiler validation via equivalence modulo inputs. In *PLDI*, 2014.
- [47] Xavier Leroy. A formally verified compiler back-end. *Journal of Automated Reasoning*, 43(4):363–446, 2009.
- [48] Panagiotis Manolios. *Mechanical Verification of Reactive Systems*. Ph.D. thesis, Department of Computer Science, The University of Texas at Austin, August 2001.
- [49] Jacob Matthews and Robert Bruce Findler. Operational semantics for multi-language programs. In *POPL*, 2007.
- [50] John McCarthy and James Painter. Correctness of a compiler for arithmetic expressions. pages 33–41. American Mathematical Society, 1967.
- [51] Paul-André Mellies and Jérôme Vouillon. Recursive polymorphic types and parametricity in an operational framework. In *LICS*, 2005.
- [52] Robin Milner. *Communicating and Mobile Systems: The π -calculus*. Cambridge University Press, 1999.
- [53] John C. Mitchell. Representation independence and data abstraction. In *Proceedings of the 13th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, POPL '86, pages 263–276, New York, NY, USA, 1986. ACM.
- [54] John C. Mitchell and Gordon D. Plotkin. Abstract types have existential type. *ACM Transactions on Programming Languages and Systems*, 10(3):470–502, July 1988.
- [55] J. Strother Moore. A mechanically verified language implementation. *J. Autom. Reason.*, 5(4):461–492, November 1989.
- [56] Lawrence S. Moss. Parametric corecursion. *Theor. Comput. Sci.*, 260(1-2):139–163, June 2001.
- [57] Andrzej S. Murawski and Nikos Tzevelekos. Game semantics for good general references. In *LICS*, 2011.
- [58] Kedar S. Namjoshi. A simple characterization of stuttering bisimulation. In *FSTTCS*, pages 284–296, 1997.

- [59] Aleksandar Nanevski, Viktor Vafeiadis, and Josh Berdine. Structuring the verification of heap-manipulating programs. In *POPL*, 2010.
- [60] Georg Neis, Derek Dreyer, and Andreas Rossberg. Non-parametric parametricity. *JFP*, 21(4&5):497–562, 2011.
- [61] Daniel Patterson, Jamie Perconti, Christos Dimoulas, and Amal Ahmed. FunTAL: Reasonably mixing a functional language with assembly. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI 2017, pages 495–509, New York, NY, USA, 2017. ACM.
- [62] James T. Perconti and Amal Ahmed. Verifying an open compiler using multi-language semantics. In *ESOP*, 2014.
- [63] Benjamin C. Pierce and Davide Sangiorgi. Behavioral equivalence in the polymorphic pi-calculus. *Journal of the ACM*, 47(3):531–586, 2000.
- [64] Andrew Pitts. Typed operational reasoning. In B. C. Pierce, editor, *Advanced Topics in Types and Programming Languages*, chapter 7. MIT Press, 2005.
- [65] Andrew Pitts and Ian Stark. Operational reasoning for functions with local state. In *HOOTS*, 1998.
- [66] Gordon Plotkin. Lambda-definability and logical relations. Technical Report SAI-RM-4, Univ. of Edinburgh, School of Artificial Intelligence, 1973.
- [67] Gordon Plotkin. Lambda-definability in the full type hierarchy. In J. P. Seldin and J. R. Hindley, editors, *Combinatory Logic, Lambda Calculus, and Formalism (Curry Festschrift)*, pages 363–373. Academic Press, Amsterdam, 1980.
- [68] Uday S. Reddy and Hongseok Yang. Correctness of data representations involving heap data structures. *Science of Computer Programming*, 50(1–3):129–160, March 2004.
- [69] John C. Reynolds. Types, abstraction, and parametric polymorphism. *Information Processing*, 1983.
- [70] Andreas Rossberg, Claudio Russo, and Derek Dreyer. F-ing modules. *Journal of Functional Programming*, 24(5):529607, 2014.
- [71] Davide Sangiorgi. The lazy lambda calculus in a concurrency scenario. *Information and Computation*, 111(1):120–153, 1994.
- [72] Davide Sangiorgi. *Introduction to Bisimulation and Coinduction*. Cambridge University Press, 2012.
- [73] Davide Sangiorgi. Origins of bisimulation and coinduction. In Davide Sangiorgi and Jan Rutten, editors, *Advanced Topics in Bisimulation and Coinduction*. Cambridge University Press, 2012.

- [74] Davide Sangiorgi, Naoki Kobayashi, and Eijiro Sumii. Environmental bisimulations for higher-order languages. *TOPLAS*, 33(1), 2011.
- [75] Jaroslav Sevcik, Viktor Vafeiadis, Francesco Zappa Nardelli, Suresh Jagannathan, and Peter Sewell. CompCertTSO: A Verified Compiler for Relaxed-Memory Concurrency. *Journal of the ACM (JACM)*, 60(3):22, 2013.
- [76] Richard Statman. Logical relations and the typed lambda-calculus. *Information and Control*, 65(2/3):85–97, 1985.
- [77] Gordon Stewart, Lennart Beringer, Santiago Cuellar, and Andrew W. Appel. Compositional CompCert. In *POPL*, 2015.
- [78] Kristian Støvring and Søren Lassen. A complete, co-inductive syntactic theory of sequential control and state. In *POPL*, 2007.
- [79] Eijiro Sumii. A complete characterization of observational equivalence in polymorphic λ -calculus with general references. In *CSL*, 2009.
- [80] Eijiro Sumii and Benjamin Pierce. A bisimulation for type abstraction and recursion. *JACM*, 54(5):1–43, 2007.
- [81] Eijiro Sumii and Benjamin Pierce. A bisimulation for type abstraction and recursion. *Journal of the ACM*, 54(5):1–43, 2007.
- [82] Eijiro Sumii and Benjamin C. Pierce. Logical relation for encryption. *J. Comput. Secur.*, 11(4):521–554, July 2003.
- [83] Kasper Svendsen, Filip Sieczkowski, and Lars Birkedal. Transfinite step-indexing: Decoupling concrete and logical steps. In *Proceedings of the 25th European Symposium on Programming Languages and Systems - Volume 9632*, pages 727–751, New York, NY, USA, 2016. Springer-Verlag New York, Inc.
- [84] W. W. Tait. Intensional interpretations of functionals of finite type I. *Journal of Symbolic Logic*, 32(2):198–212, 1967.
- [85] Yong Kiam Tan, Magnus O. Myreen, Ramana Kumar, Anthony Fox, Scott Owens, and Michael Norrish. A new verified compiler backend for CakeML. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming*, ICFP 2016, pages 60–73, New York, NY, USA, 2016. ACM.
- [86] Enrico Tassi. Coq reference manual: Asynchronous and parallel proof processing. <https://coq.inria.fr/refman/async-proofs.html>.
- [87] Aaron Turon, Derek Dreyer, and Lars Birkedal. Unifying refinement and Hoare-style reasoning in a logic for higher-order concurrency. In *Proceedings of the 18th ACM SIGPLAN International Conference on Functional Programming*, ICFP '13, pages 377–390, New York, NY, USA, 2013. ACM.

- [88] Viktor Vafeiadis. Concurrent separation logic and operational semantics. In *MFPS*, 2011.
- [89] Viktor Vafeiadis and Matthew Parkinson. A marriage of rely/guarantee and separation logic. In *CONCUR*, 2007.
- [90] Philip Wadler. Theorems for free! In *FPCA*, 1989.
- [91] Peng Wang, Santiago Cuellar, and Adam Chlipala. Compiler verification meets cross-language linking via data abstraction. In *OOPSLA*, 2014.
- [92] 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 '11, pages 283–294, New York, NY, USA, 2011. ACM.