

Deliverable D4.1

Progress report for WP4: Reasoning about heterogenous systems

Project acronym	ADVENT
Project title	Architecture-driven verification of systems software
Funding scheme	FP7 FET Young Explorers
Scientific coordinator	Dr. Alexey Gotsman, IMDEA Software Institute, Alexey.Gotsman@imdea.org, +34 911 01 22 02

1 Summary

Work Package 4 in the ADVENT project is concerned with reasoning about heterogeneous systems, with a focus on managed languages and how compilation affects formal reasoning about such systems. Throughout the first year of the project, we have focused on developing the underlying theory for proving compositionally the correctness of compositional compilers, as well as for reasoning about the interfaces between compilers and runtime libraries.

Specifically, we have developed a theory for proving modular compilers correct, which we call *parametric bisimulations* (PBs). Our early work [5, 6] (prior to the official commencement of ADVENT) proposed the underlying theory (which we called *relation transition systems* at the time) and proved that our PB model is transitive. In the course of ADVENT, we have extended this model to account for programming languages with more advanced control flow primitives, such as exceptions and call/cc. In doing so, it was necessary to extend parametric bisimulations to account for *stuttering steps* in the execution. We presented our findings in Technical Report MPI-SWS-2014-003 [7]. We are currently working on applying PBs to verify a compositional multi-phase compiler from a core ML language to an idealised assembly language.

In addition, we have looked at runtime systems that support managed programming languages. Because of the increasing importance of executing large parallel computations efficiently and correctly, we have focused at what support the programming language and its runtime system can provide for detecting and recovering from execution failures in the context of a parallel computation. In a TASE 2013 publication [9] we specified the semantics of such a programming language and a runtime, and proved correspondence properties between a higher-level and a low-level design.

2 Parametric Bisimulations [5, 6, 7]

In the last several years, researchers have developed a number of effective methods, such as bisimulations and Kripke logical relations (KLRs), for reasoning about program equivalence in higher-order imperative languages like ML.

While these methods provide us with extremely powerful reasoning principles for proving program equivalence, they are not directly applicable to reason compositionally about equivalences between programs written in different languages, such as the source and target of a verified compiler.

Existing bisimulation methods for higher-order stateful languages rely crucially on “syntactic” devices (e.g., context closure) in order to deal properly with unknown higher-order values that may be passed in as function arguments. While these syntactic devices are appropriate for proving “contextual” properties (such as contextual equivalence), they bake in the assumption that the programs being related share a common syntactic notion of “context”, which is clearly not a valid assumption in the inter-language setting. In contrast, KLRs have been successfully generalized to the inter-language setting—specifically, to the goal of establishing “compositional compiler correctness” [4]—but they suffer from an orthogonal limitation, namely that in general they are not transitively composable.

Therefore, in a POPL’12 paper [5], we proposed a new technique, which we call *parametric bisimulations* (PBs), that overcomes the aforementioned limitations. PBs fruitfully synthesize the direct coinductive style of bisimulations with the flexible invariants on local state afforded by KLRs, thereby enabling clean and elegant proofs about local state and recursive features simultaneously. In addition, they support transitive composition of equivalence proofs.

However, the PB model of our previous work suffered from two limitations. First, it failed to validate the eta law for function values, which is important for our intended application of compiler certification. Second, it was not clear how to scale the method to reason about control effects.

In the technical report MPI-SWS-2014-003, we proposed *stuttering parametric bisimulations* (SPBs), an extension of PBs that addresses their aforementioned limitations. Interestingly, despite the fact that the η -law and control effects seem like unrelated issues, our solutions to both problems hinge on the same technical device, namely the use of a “logical” reduction semantics that permits finite but unbounded stuttering in between physical steps. This technique is closely related to the key idea in well-founded and stuttering bisimulations, adapted here for the first time to reasoning about open, higher-order programs. We present SPBs—along with meta-theoretic results and example applications—for a language with recursive types and first-class continuations. Following our previous account of PBs, we can easily extend SPBs to handle abstract types and general mutable references as well. All our results have been fully mechanized in the Coq proof assistant [2] and are available online.

We are currently working on extending the PBs to multi-language setting and on verifying a compositional multi-phase compiler from a core ML-like language to an idealised assembly language.

3 Fault-Tolerant Fork-Join Parallelism [9]

When running big parallel computations on thousands of processors, the probability that an individual processor will fail during the execution is actually quite high, and cannot be ignored by the language runtime. For example, if we assume that the *mean time between failures* (MTBF) for a single machine is one year, and we use one thousand machines for a single computation, then the MTBF for the whole computation becomes $1 \text{ year} \div 1000 \approx 9 \text{ hours}$.

As a result, language runtimes try to address this problem in one of two ways. A simple—albeit expensive—solution is to use *replication*. In theory, we can straightforwardly deal with a single fail-stop failure with 3-way replication [1], and with a single Byzantine failure with 4-way replication [8]. Replication, however, comes at a significant cost, not only in execution time (since fewer execution units are available), but also in the amount of energy required to compute the correct result.

The alternative approach to replication is to use *checkpointing*: that is, to run the computation optimistically with no replication, to detect any failures that occur, and to rerun the parts of the computation affected by those failures [3]. The benefit of checkpointing over the replication approach is that the effective replication rate is determined by the number of actual failures that occurred in an execution and how large a sub-computation was interrupted rather than the maximum number of failures that the system can tolerate. To implement checkpointing, one assumes that some part of the storage space is safe (non-failing) and uses that to store fields needed to recover from failures. This safe storage subsystem may internally be implemented using replication, but this kind of storage replication is much lighter-weight than replicating the entire computation.

As for proving the correctness of these two approaches, that of replication is relatively straightforward, because it uses correctly computed results from one of the replicas in the system. In the checkpointing approach, however, correctness is not so straightforward, because failed processors can be in inconsistent states and partially computed expressions are used in reexecutions.

In this work, we formalized checkpointing from a programming language perspective and proved its correctness. For simplicity, we performed this work in the context of a purely functional language with fork-join parallelism. For this language, we developed a high-level formal operational semantics capturing the essence of the checkpointing approach. In our semantics, the execution of a parallel computation may fail at any point; failures can then be detected and the appropriate parts of a failed computation can be restarted. This high-level semantics is easy to understand, and may thus be used as a basis for reasoning about fault-tolerant parallel programs.

To justify the completeness of our semantics with respect to actual implementations, we also developed a lower-level semantics, which models run-time failures and parallel task execution at the processor level. The low-level semantics also models recovery from failed executions, as may be performed by a runtime system, where a processor can observe that some other processor has failed, and restart its execution. We then proved theorems relating the two semantics and showing that our fault-aware semantics are sound: whenever a program evaluates to a value in the fault-aware semantics (perhaps by failing a few times and recovering from the failures), then it can also evaluate to the same value under the standard fault-free semantics. All lemmas and theorems were proved using the Coq proof assistant [2] and are available online.

References

- [1] A. Clement, M. Kapritsos, S. Lee, Y. Wang, L. Alvisi, M. Dahlin, and T. Riche. UpRight cluster services. In *Proceedings of the ACM SIGOPS 22nd SOSP*, pages 277–290. ACM, 2009.
- [2] Coq development team. The Coq proof assistant. <http://coq.inria.fr/>.
- [3] E. N. Elnozahy, L. Alvisi, Y.-M. Wang, and D. B. Johnson. A survey of rollback-recovery protocols in message-passing systems. *ACM Computing Surveys (CSUR)*, 34(3):375–408, 2002.
- [4] C.-K. Hur and D. Dreyer. A Kripke logical relation between ML and assembly. In *POPL*, pages 133–146. ACM, 2011.
- [5] C.-K. Hur, D. Dreyer, G. Neis, and V. Vafeiadis. The marriage of bisimulations and kripke logical relations. In *POPL*, pages 59–72, 2012.
- [6] C.-K. Hur, G. Neis, D. Dreyer, and V. Vafeiadis. The transitive composability of relation transition systems. Technical Report MPI-SWS-2012-002, MPI-SWS, 2012.
- [7] C.-K. Hur, G. Neis, D. Dreyer, and V. Vafeiadis. A logical step forward in parametric bisimulations. Technical Report MPI-SWS-2014-003, MPI-SWS, 2014.
- [8] L. Lamport, R. Shostak, and M. Pease. The Byzantine generals problem. *TOPLAS*, 4(3):382–401, 1982.
- [9] M. Zengin and V. Vafeiadis. A programming language approach to fault tolerance for fork-join parallelism. In *TASE*, pages 105–112, 2013.

List of Attached Papers

- [7] 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.
- [9] Mustafa Zengin and Viktor Vafeiadis. A programming language approach to fault tolerance for fork-join parallelism. In *TASE*, pages 105–112, 2013.

Technical Report MPI-SWS-2014-003
January 2014

A Logical Step Forward in Parametric Bisimulations

Chung-Kil Hur Georg Neis Derek Dreyer Viktor Vafeiadis
Seoul National University Max Planck Institute for Software Systems (MPI-SWS)
gil.hur@cse.snu.ac.kr {neis,dreyer,viktor}@mpi-sws.org

Abstract

In the last several years, a number of effective methods have been developed for reasoning about program equivalence in higher-order imperative languages like ML. Most recently, we proposed *parametric bisimulations (PBs)*, which fruitfully synthesize the direct coinductive style of bisimulations with the flexible invariants on local state afforded by Kripke logical relations, and which furthermore support transitive composition of equivalence proofs. However, the PB model of our previous work suffered from two limitations. First, it failed to validate the eta law for function values, which is important for our intended application of compiler certification. Second, it was not clear how to scale the method to reason about control effects.

In this paper, we propose *stuttering parametric bisimulations (SPBs)*, a variant of PBs that addresses their aforementioned limitations. Interestingly, despite the fact that the eta law and control effects seem like unrelated issues, our solutions to both problems hinge on the same technical device, namely the use of a “logical” reduction semantics that permits finite but unbounded stuttering in between physical steps. This technique is closely related to the key idea in *well-founded* and *stuttering* bisimulations, adapted here for the first time to reasoning about open, higher-order programs. We present SPBs—along with meta-theoretic results and example applications—for a language with recursive types and first-class continuations. Following our previous account of PBs, we can easily extend SPBs to handle abstract types and general mutable references as well (see the appendix for details). All our results have been fully mechanized in Coq.

CONTENTS

1	Introduction	3
2	The Languages λ^μ and λ_{cc}^μ	3
3	Parametric Bisimulations (PBs)	4
4	Stuttering Parametric Bisimulations (First Step: Validating Eta)	5
4.1	The Problem with Eta	5
4.2	Guardedness Revisited	7
4.3	Logical Reduction and the Stutter Budget	7
4.4	Stuttering Parametric Bisimulations (SPBs) for λ^μ	7
4.5	Eta Revisited	8
5	Stuttering Parametric Bisimulations (Second Step: Supporting Continuations)	8
5.1	Contextualizing the Model	9
5.2	Supporting First-Class Continuations	9
5.3	Example: callcc in a Loop	10
5.4	Using Parameterized Coinduction	10
6	Metatheory	10
6.1	Transitivity	11
7	Discussion and Related Work	11
	References	12
	Appendix A: SPBs for λ^μ and λ_{cc}^μ	13
A.1	Languages λ^μ and λ_{cc}^μ	13
A.1.1	Statics	13
A.1.2	Dynamics	14
A.2	Simple SPB Model for λ^μ	14
A.2.1	Definition	14
A.2.2	Properties	15
A.2.3	Example: Eta Equivalence (lambda version)	16
A.3	SPB Models for λ^μ and λ_{cc}^μ	16
A.3.1	Definition	16
A.3.2	Properties	17
A.3.3	Example: Callcc in a Loop	21
A.3.4	Example: Eta Equivalence (lambda version)	21
A.3.5	Example: Syntactic Minimal Invariance	22
A.3.6	Using Paco (parameterized coinduction)	22
	Appendix B: SPBs for $F^{\mu!}$ and $F_{cc}^{\mu!}$	24
B.1	Languages $F^{\mu!}$ and $F_{cc}^{\mu!}$	24
B.1.1	Statics	24
B.1.2	Dynamics	25
B.2	SPB Models for $F^{\mu!}$ and $F_{cc}^{\mu!}$	26
B.2.1	Definition	26
B.2.2	Example: Well-Bracketed State Change	28

1 – INTRODUCTION

A longstanding problem in semantics is to find effective methods for reasoning about program equivalence in ML-like languages supporting both functional and imperative features. In the last several years, considerable progress has been made on this problem, primarily by advancements to two different classes of proof methods—*bisimulations* [20, 17, 18] and *step-indexed Kripke logical relations (SKLRs)* [1, 6].

In recent work [8], we proposed a new method for proving contextual equivalences in ML-like languages, which at the time we called *relation transition systems*, but since then have been referring to as *parametric bisimulations (PBs)*. PBs marry together some of the best features of state-of-the-art bisimulations and SKLRs into a single method:

- Like bisimulations [17, 19, 18], PBs support reasoning about “recursive” language features (*e.g.*, recursive types, higher-order state) in a *direct, coinductive* style.
- Like SKLRs [1, 6], PBs provide a very flexible treatment of “local” state, in which one can define a per-module *state transition system* to express and enforce invariants on how the module’s local state may change over time.

Furthermore, PBs were designed to overcome some apparent limitations of their ancestors with regards to their potential to scale to *inter-language reasoning*—*i.e.*, reasoning compositionally about equivalences between programs written in different languages, such as the source and target of a certified compiler [2, 7]:

- Like SKLRs, PBs are truly *semantic*, in the sense that they do not rely fundamentally on the assumption that the programs they relate are written in the same language. (In contrast, bisimulations for higher-order languages rely crucially on this assumption.)
- Unlike SKLRs, the relation between programs induced by PBs is *transitive*. We believe this property will prove essential for our ultimate goal of using PBs for certifying *multi-phase* compilers, as it will enable correctness proofs for different phases to be linked transitively.

Unfortunately, the PB model we developed in our previous work [8] suffered from two key limitations:

- 1) It failed 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 [3]) which depend on it. This law is potentially quite important for our ultimate goal of compositional compiler certification, since eta-conversion is commonly used in compile-time optimization.

- 2) The language it modeled did not provide any form of control effects, such as *first-class continuations*, and it was not clear how to scale it to account for a language with such effects. In particular, the model baked in the assumption that the language in question had a “uniform” reduction semantics (*i.e.*, that the reduction relation was parametric in the evaluation context), an assumption that is of course not valid in the presence of *callcc*.

In this paper, we propose *stuttering parametric bisimulations (SPBs)*, which build closely on PBs and enjoy all their benefits, while also circumventing these two limitations. Interestingly, although the limitations of PBs appear at first glance to be completely unrelated, our solutions to both of them hinge on the same exact technical device, namely 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* [15] and *stuttering* [4] bisimulations, adapted here for the first time to reasoning about open, higher-order programs.

There is a superficial sense in which SPB proofs may seem similar to 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!

We present SPBs in the setting of λ_{cc}^μ , a CBV λ -calculus with recursive types and first-class continuations. λ_{cc}^μ provides a rich enough setting to explore the problems with PBs and the solutions offered by SPBs, while avoiding other orthogonal issues. Following our previous work on PBs, the SPB approach can be scaled straightforwardly to account for abstract types and general references, and the full SPB model supporting those features is given in the appendix. Moreover, our meta-theoretic results (soundness and transitivity—see Section 6) for the full PB model have been mechanized in Coq. This Coq development is available online at:

<http://www.mpi-sws.org/~neis/spbs>

We begin, in Section 2, by defining the language λ_{cc}^μ and its pure fragment λ^μ (lacking first-class continuations). In Section 3, we review our previous PB model for λ^μ , and the essential ideas behind it. We then present our SPB model in two stages. First, in Section 4, we explain the problem with the eta law, and how the logical reduction semantics fixes it, leading to the definition of a SPB model for λ^μ that closely follows the structure of the PB model. Second, in Section 5, we show how the same logical reduction semantics also aids in modeling control effects, and we use this insight to develop our generalized SPB model for λ_{cc}^μ . In Section 6, we summarize the key meta-theoretic results of the paper. Finally, in Section 7, we discuss related work and conclude.

2 – THE LANGUAGES λ^μ AND λ_{cc}^μ

Figure 1 gives the syntax of λ_{cc}^μ , along with excerpts of its static and dynamic semantics. The language is equipped with a standard type system, as well as a standard deterministic CBV dynamic semantics using evaluation contexts (aka

$$\begin{array}{l}
\text{Statics: } \boxed{\Gamma \vdash p : \tau} \\
\Gamma ::= \cdot \mid \Gamma, x:\tau \quad \text{where } \text{fv}(\tau) = \emptyset \\
\tau, \sigma ::= \alpha \mid \text{int} \mid \sigma_1 \times \sigma_2 \mid \sigma_1 + \sigma_2 \mid \sigma_1 \rightarrow \sigma_2 \mid \mu\alpha. \sigma \mid \text{cont } \sigma \\
p ::= x \mid n \mid p_1 \odot p_2 \mid \text{ifz } p \text{ then } p_1 \text{ else } p_2 \mid \langle p_1, p_2 \rangle \mid p.1 \mid \\
\quad p.2 \mid \text{inl}_\sigma p \mid \text{inr}_\sigma p \mid (\text{case } p \text{ of inl } x \Rightarrow p_1 \mid \text{inr } x \Rightarrow p_2) \mid \\
\quad \text{fix } f(x:\sigma_1):\sigma_2. p \mid p_1 p_2 \mid \text{roll}_\sigma p \mid \text{unroll } p \mid \\
\quad \text{callcc}_\sigma(x. p) \mid \text{throw}_\sigma p_1 \text{ to } p_2 \mid \text{isolate } p \\
\frac{\Gamma, f:(\tau_1 \rightarrow \tau_2), x:\tau_1 \vdash p : \tau_2}{\Gamma \vdash \text{fix } f(x:\tau_1):\tau_2. p : \tau_1 \rightarrow \tau_2} \quad \frac{\Gamma \vdash p_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma \vdash p_2 : \tau_1}{\Gamma \vdash p_1 p_2 : \tau_1} \\
\frac{\Gamma, x:\text{cont } \tau \vdash p : \tau}{\Gamma \vdash \text{callcc}_\tau(x. p) : \tau} \quad \frac{\Gamma \vdash p' : \tau' \quad \Gamma \vdash p : \text{cont } \tau'}{\Gamma \vdash \text{throw}_\tau p' \text{ to } p : \tau} \\
\frac{\Gamma \vdash p : \tau \rightarrow \text{int}}{\Gamma \vdash \text{isolate } p : \text{cont } \tau}
\end{array}$$

$$\begin{array}{l}
\text{Dynamics: } \boxed{e \hookrightarrow e} \\
v ::= x \mid n \mid \langle v_1, v_2 \rangle \mid \text{inl } v \mid \text{inr } v \mid \text{fix } f(x). e \mid \text{roll } v \mid \text{cont } K \\
e ::= v \mid e_1 \odot e_2 \mid \text{ifz } e_0 \text{ then } e_1 \text{ else } e_2 \mid \langle e_1, e_2 \rangle \mid e.1 \mid e.2 \mid \\
\text{inl } e \mid \text{inr } e \mid (\text{case } e \text{ of inl } x \Rightarrow e_1 \mid \text{inr } x \Rightarrow e_2) \mid e_1 e_2 \mid \\
\text{roll } e \mid \text{unroll } e \mid \text{callcc}(x. e) \mid \text{throw } e_1 \text{ to } e_2 \mid \text{isolate } e \\
K ::= \bullet \mid K \odot e \mid v \odot K \mid \text{ifz } K \text{ then } e_1 \text{ else } e_2 \mid \langle K, e \rangle \mid \\
\langle v, K \rangle \mid K.1 \mid K.2 \mid \text{inl } K \mid \text{inr } K \mid \\
(\text{case } K \text{ of inl } x \Rightarrow e_1 \mid \text{inr } x \Rightarrow e_2) \mid K e \mid v K \mid \text{roll } K \mid \\
\text{unroll } K \mid \text{throw } K \text{ to } e \mid \text{throw } v \text{ to } K \mid \text{isolate } K \\
\begin{array}{l}
K[(\text{fix } f(x). e) v] \quad \hookrightarrow \quad K[e[(\text{fix } f(x). e)/f, v/x]] \\
K[\text{callcc}(x. e)] \quad \hookrightarrow \quad K[e[\text{cont } K/x]] \\
K[\text{throw } v \text{ to cont } K'] \quad \hookrightarrow \quad K'[v] \\
K[\text{isolate } v] \quad \hookrightarrow \quad K[\text{cont}(v \bullet)]
\end{array}
\end{array}$$

Fig. 1. Statics and dynamics (excerpts) of λ^μ and λ_{cc}^μ .

continuations) K . While the static semantics talks about typed (annotated) programs p , the dynamic semantics is defined for erased programs $e = |p|$.

First-class continuations have type $\text{cont } \tau$, and all associated constructs and rules are **highlighted in red**. At runtime, a first-class continuation is just another value form ($\text{cont } K$). There are two ways a program can get hold of such a value: either by capturing its current continuation via callcc , or by turning a function into a continuation via isolate (thereby isolating it from the current continuation). It can then yield control to a value of type $\text{cont } \tau$ by throw -ing it a value of type τ .

The fragment λ^μ is obtained from λ_{cc}^μ by simply dropping the highlighted constructs and rules.

Contextual equivalence for either language is defined in the usual way with the help of contexts C (programs with a hole):

Definition 1 (Contextual Equivalence).

$$\begin{array}{l}
\boxed{\Gamma \vdash p_1 \sim_{\text{ctx}} p_2 : \tau} := \Gamma \vdash p_1 : \tau \wedge \Gamma \vdash p_2 : \tau \wedge \forall C. \\
\vdash C : (\Gamma; \tau) \rightsquigarrow (\cdot; \text{int}) \implies (|C[p_1]| \hookrightarrow^\omega \iff |C[p_2]| \hookrightarrow^\omega)
\end{array}$$

Here \hookrightarrow^ω stands for divergence (infinite reduction). The definition of contexts and their typing judgment straightforwardly follows that of programs p and is omitted for space reasons. Suffice it to say that $\vdash C : (\Gamma; \tau) \rightsquigarrow (\Gamma'; \tau')$ implies $\forall p. \Gamma \vdash p : \tau \implies \Gamma' \vdash C[p] : \tau'$.

In the remainder, we often write $\lambda x. e$ as shorthand for the non-recursive function $\text{fix } f(x). e$, where $f \notin \text{fv}(e)$.

3 – PARAMETRIC BISIMULATIONS (PBs)

In this section, we briefly review the main ideas behind our previous model of *parametric bisimulations (PBs)* [8] (originally named *relation transition systems*) as well as the formal definition of the PB model for λ^μ , which is given in the left column of Figure 2.

The top-level judgment of our PB model has the form $\Gamma \vdash e_1 \sim e_2 : \tau$, stipulating that e_1 and e_2 are equivalent terms at type τ . Typically, in coinductive proofs, to establish such an equivalence, one exhibits a *coinduction hypothesis* L , which relates e_1 and e_2 but also relates other auxiliary terms that one needs to prove equivalent in the course of proving e_1 and e_2 equivalent. The soundness of this hypothesis w.r.t. contextual equivalence is then established by proving it to be *consistent* in a certain sense (a.k.a. a “bisimulation”).

PB proofs work in a similar way, but with an unusual twist concerning the treatment of (higher-order) functions. If in the coinduction hypothesis L we claim that two functions $\lambda x. e_1$ and $\lambda x. e_2$ are equivalent at type $\sigma \rightarrow \tau$, the aforementioned consistency condition will require us to demonstrate that e_1 and e_2 do in fact behave equivalently (at type τ) when instantiated with “equivalent arguments” at type σ .¹ A natural question then arises: from what relation do we draw these “equivalent arguments”? This is a tough question—if we knew how to usefully characterize when values of arbitrary type σ were equivalent, we would have solved our original problem!

The essential novelty of PBs is to answer this question precisely by *not* answering it. To understand this cryptic statement, let us make a distinction between *global knowledge* and *local knowledge*. Local knowledge is another term for the coinduction hypothesis L : it specifies which values *we* “know” are equivalent *in our proof*, but it is local to our proof—it does not pretend to be a global characterization of which values are equivalent in general. In contrast, the global knowledge G represents the sum total of knowledge concerning equivalence of values in the “whole program” in which $\lambda x. e_i$ appear, and it is this G from which we draw “equivalent arguments”. The key insight of PB proofs is that, in order to verify the consistency of L , we do not need to know what G is. In fact, we do not even need to know that G is sound w.r.t. contextual equivalence. Rather, we simply take G as a *parameter* of our equivalence proof.

This idea of parameterizing the proof over the global knowledge G has important implications for how we define consistency for function values. If L relates $\lambda x. e_1$ and $\lambda x. e_2$ at $\sigma \rightarrow \tau$, then we must show that for all (v_1, v_2) related by G at σ , $e_1[v_1/x]$ and $e_2[v_2/x]$ “behave equivalently” at τ . The trouble is that, in an absolute sense, they don’t. Knowing that v_1 and v_2 are related by G tells us absolutely nothing, since G is a parameter that can relate any two values (at least at function type—see the discussion of “value closure” below).

¹The functions related by L may of course be recursive (e.g., $\text{fix } f(x). e$). We restrict attention to λ -terms here just to simplify the presentation.

For example, suppose that $\sigma = \text{int} \rightarrow \text{int}$, $\tau = \text{int}$, and $e_1 = e_2 = x(0)$. In this case, we should clearly be able to prove $\lambda x.e_1$ and $\lambda x.e_2$ equivalent—they are syntactically equal!—but it is possible that they are passed as arguments, say, $v_1 = \lambda x.x + 1$ and $v_2 = 5$, in which case $e_1[v_1/x] \hookrightarrow^* 1$, while $e_2[v_2/x] \hookrightarrow^* 5(0)$, a stuck term. Even if we were to restrict G to, at function type, only relate λ -expressions (instead of arbitrary junk like the integer 5), G might still relate v_1 here with, say, a divergent function, in which case $e_2[v_2/x] \hookrightarrow^\omega$.

While $e_1[v_1/x]$ and $e_2[v_2/x]$ in this example clearly do not have the same observable behavior, they *can* be understood to have the same *local* behavior. Intuitively, two terms are locally equivalent w.r.t. G if they behave equivalently *modulo* what happens during calls to functions related by G . In the above example, $e_1[v_1/x]$ and $e_2[v_2/x]$ apply values related by G (namely, v_1 and v_2) to the same integer argument (0). The fact that they behave differently is thus not their own fault, but G 's fault, so we can say that in fact they are locally equivalent.

This notion is formalized by the *local term equivalence* relation $\mathbf{E}(G)$. We say that two closed terms e_1 and e_2 are *locally equivalent* at a given type τ w.r.t. a global knowledge G —denoted $(e_1, e_2) \in \mathbf{E}(G)(\tau)$ —if one of these cases holds:

(**Case** \uparrow) they both diverge; or

(**Case** \downarrow) they both reduce to related values; or

(**Case** ζ) they both reduce to related “stuck” configurations ($\mathbf{S}(G, G)$), *i.e.*, function calls where related function values are applied to related argument values inside locally equivalent continuations ($\mathbf{K}(G)$). Locally equivalent continuations, in turn, are those which, when filled with related values, result (coinductively) in locally equivalent terms. In all cases, “related” values are drawn from the global knowledge G (or rather its value closure \overline{G} , as we explain below).

The definition of $\mathbf{E}(G)$ is highly reminiscent of *normal form* (or *open*) bisimulations [16, 12, 18], in which one establishes the consistency of equivalence of $\lambda x.e_1$ and $\lambda x.e_2$ by showing equivalence of e_1 and e_2 as open terms (assuming x is a “fresh” variable). Correspondingly, normal form bisimulations permit the proof to “get stuck” at a point where both terms apply the same function variable x (*e.g.*, in the example above, $x(0)$). This is no accident: PBs were inspired heavily by normal form bisimulations. The key difference is that PBs’ use of a global knowledge is more *semantic*: by drawing v_1 and v_2 from an unknown G instead of modeling them as the same variable x , we have the potential to scale to reasoning about equivalences between different languages (*e.g.*, involving assembly, which has no notion of variable binding [7]).

To conclude this section, we mention three important technical points in the formalization of PBs.

First, we restrict the local and global knowledge to only relate (closed) *values*, not arbitrary terms, and only at “flexible” types, CTyF , which for λ^μ means just function types. Value equivalence at flexible types, $R \in \text{VRelF}$, is then extended to all (closed) types by the inductively-constructed *value closure*, $\overline{R} \in \text{VRel}$. (We call the non-flexible types “rigid” since the value closure defines equivalence at those types in a fixed, canonical way.) Term equivalence is accounted for by $\mathbf{E}(R)$.

Second, when parameterizing over the global knowledge G , we impose the condition that G should at least contain the local knowledge L of our proof, since global subsumes local knowledge. This requirement is critical in enabling coinductive reasoning; for example, to show that two recursive functions f_1 and f_2 related by L are equivalent, we may wish to show that $f_1(v_1)$ and $f_2(v_2)$ evaluate to some expressions of the form $K_1[f_1(v'_1)]$ and $K_2[f_2(v'_2)]$ (with (v'_1, v'_2) related by \overline{G} and (K_1, K_2) by $\mathbf{K}(G)$). In this case, if we know that G contains L , then we also know that $(f_1, f_2) \in G$, and we can appeal to the “stuck” case (ζ) of $\mathbf{E}(G)$ to complete the proof.

Formally, the requirement that G subsumes L is slightly more complicated, because we allow L to be *itself* parameterized over G (so long as it is monotone in its G parameter—see the definition of LK). This additional parameterization of L over G is essential: it enables L to assert equivalences between open terms by quantifying over closing instantiations drawn from G . (We will see an example of this in Section 4.) As a result, the condition that “ G subsumes L ”, written $G \in \text{GK}(L)$, is defined to mean that $G \supseteq L(G)$.

Finally, in order to keep the kind of coinductive reasoning we just described sound, we must be quite careful in the definition of “consistent(L)”. Specifically, for each pair of functions (f_1, f_2) related by $L(G)$, and arguments (v_1, v_2) related by \overline{G} , we cannot simply require $(f_1(v_1), f_2(v_2))$ to be related by $\mathbf{E}(G)$ because, via the stuck case of $\mathbf{E}(G)$, this is a tautology! Instead, we demand that $f_i(v_i) \hookrightarrow e_i$, and that e_1 and e_2 are related by $\mathbf{E}(G)$. 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*. As we will see in the next section, however, the guardedness here is more restrictive than it ought to be.

4 – STUTTERING PARAMETRIC BISIMULATIONS (FIRST STEP: VALIDATING ETA)

4.1 The Problem with Eta

We begin by reviewing the inherent problem with eta in the PB model. 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 for the PB model presented above. Here we prove as much for the case of $\tau' = \tau = \text{int}$.

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 $(v_1, v_2) \in G(\text{int} \rightarrow \text{int})$ we have $(v_1, \lambda x.v_2 x) \in \mathbf{E}(G)(\text{int} \rightarrow \text{int})$. Being values, 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 $G^\zeta \in \text{GK}(L)$ 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^ζ be the least global knowledge that subsumes L and also relates $\lambda y. 0$ and $\lambda y. 1$ at $\text{int} \rightarrow \text{int}$.²

$$G^\zeta = L(G^\zeta) \cup \{(\text{int} \rightarrow \text{int}, \lambda y. 0, \lambda x. 1)\}$$

²This is a simple fixed point construction since L is a monotone function.

CTy(F), CVal, CExp, CCont = the sets of closed (flexible) types, closed values, closed expressions, and closed continuations, respectively

VRelF := $\mathbb{P}(\text{CTyF} \times \text{CVal} \times \text{CVal})$	VRelF := $\mathbb{P}(\text{CTyF} \times \mathbb{N} \times \text{CVal} \times \mathbb{N} \times \text{CVal})$
VRel := $\mathbb{P}(\text{CTy} \times \text{CVal} \times \text{CVal})$	VRel := $\mathbb{P}(\text{CTy} \times \mathbb{N} \times \text{CVal} \times \mathbb{N} \times \text{CVal})$
KRel := $\mathbb{P}(\text{CTy} \times \text{CTy} \times \text{CCont} \times \text{CCont})$	KRel := $\mathbb{P}(\text{CTy} \times \text{CTy} \times \mathbb{N} \times \text{CCont} \times \mathbb{N} \times \text{CCont})$
ERel := $\mathbb{P}(\text{CTy} \times \text{CExp} \times \text{CExp})$	ERel := $\mathbb{P}(\text{CTy} \times \mathbb{N} \times \text{CExp} \times \mathbb{N} \times \text{CExp})$
$\overline{R}(\tau \rightarrow \tau') := R(\tau \rightarrow \tau')$	$\overline{R}(\tau \rightarrow \tau') := R(\tau \rightarrow \tau')$
$\overline{R}(\text{int}) := \{(m, m)\}$	$\overline{R}(\text{int}) := \{(_, m, _, m)\}$
$\overline{R}(\tau \times \tau') := \{(\langle v_1, v_1' \rangle, \langle v_2, v_2' \rangle) \mid (v_1, v_2) \in \overline{R}(\tau) \cap \{(\langle v_1, v_1' \rangle, \langle v_2, v_2' \rangle) \mid (v_1', v_2') \in \overline{R}(\tau')\}\}$	$\overline{R}(\tau \times \tau') := \{(_, \langle v_1, v_1' \rangle, _, \langle v_2, v_2' \rangle) \mid (_, v_1, _, v_2) \in \overline{R}(\tau) \cap \{(_, \langle v_1, v_1' \rangle, _, \langle v_2, v_2' \rangle) \mid (_, v_1', _, v_2') \in \overline{R}(\tau')\}\}$
$\overline{R}(\tau + \tau') := \{(\text{inl } v_1, \text{inl } v_2) \mid (v_1, v_2) \in \overline{R}(\tau)\} \cup \{(\text{inr } v_1', \text{inr } v_2') \mid (v_1', v_2') \in \overline{R}(\tau')\}$	$\overline{R}(\tau + \tau') := \{(_, \text{inl } v_1, _, \text{inl } v_2) \mid (_, v_1, _, v_2) \in \overline{R}(\tau)\} \cup \{(_, \text{inr } v_1', _, \text{inr } v_2') \mid (_, v_1', _, v_2') \in \overline{R}(\tau')\}$
$\overline{R}(\mu\alpha. \sigma) := \{(\text{roll } v_1, \text{roll } v_2) \mid (v_1, v_2) \in \overline{R}(\sigma[\mu\alpha. \sigma/\alpha])\}$	$\overline{R}(\mu\alpha. \sigma) := \{(_, \text{roll } v_1, _, \text{roll } v_2) \mid (_, v_1, _, v_2) \in \overline{R}(\sigma[\mu\alpha. \sigma/\alpha])\}$
LK := $\{L \in \text{VRelF} \rightarrow \text{VRelF} \mid \forall R. \forall (\tau, v_1, v_2) \in L(R). (\forall R' \supseteq R. (\tau, v_1, v_2) \in L(R')) \wedge (\exists f_i, x_i, e_i. v_i = \text{fix } f_i(x_i). e_i)\}$	LK := $\{L \in \text{VRelF} \rightarrow \text{VRelF} \mid \forall R. \forall (\tau, n_1, v_1, n_2, v_2) \in L(R). \forall R' \supseteq R. n_1' \geq n_1, n_2' \geq n_2. (\tau, n_1', v_1, n_2', v_2) \in L(R')\}$
GK(L) := $\{G \in \text{VRelF} \mid G \supseteq L(G)\}$	GK(L) := $\{G \in \text{VRelF} \mid G \supseteq L(G) \wedge \forall (\tau, n_1', v_1, n_2, v_2) \in G. \forall n_1' \geq n_1, n_2' \geq n_2. (\tau, n_1', v_1, n_2', v_2) \in G\}$
$\mathbf{S} \in \text{VRelF} \times \text{VRelF} \rightarrow \text{ERel}$	$\mathbf{S} \in \text{VRelF} \times \text{VRelF} \rightarrow \text{ERel}$
$\mathbf{S}(R_c, R) := \{(\tau, v_1, v_1', v_2, v_2') \mid \exists \tau'. (\tau' \rightarrow \tau, v_1, v_2) \in R_c \wedge (\tau', v_1, v_2) \in \overline{R}\}$	$\mathbf{S}(R_c, R) := \{(\tau, n_1, v_1, v_1', n_2, v_2, v_2') \mid \exists \tau'. (\tau' \rightarrow \tau, n_1, v_1, n_2, v_2) \in R_c \wedge (\tau', _, v_1', _, v_2') \in \overline{R}\}$
$\mathbf{K} \in \text{VRelF} \rightarrow \text{KRel}$	$\mathbf{K} \in \text{VRelF} \rightarrow \text{KRel}$
$\mathbf{K}(R) := \{(\tau', \tau, K_1, K_2) \mid \forall (v_1, v_2) \in \overline{R}(\tau'). (\tau, K_1[v_1], K_2[v_2]) \in \mathbf{E}(R)\}$	$\mathbf{K}(R) := \{(\tau', \tau, k_1, K_1, k_2, K_2) \mid \forall (n_1, v_1, n_2, v_2) \in \overline{R}(\tau'). (\tau, k_1 + n_1, K_1[v_1], k_2 + n_2, K_2[v_2]) \in \mathbf{E}(R)\}$
$\mathbf{E} \in \text{VRelF} \rightarrow \text{ERel}$	$\mathbf{E} \in \text{VRelF} \rightarrow \text{ERel}$
$\mathbf{E}(R) := \{(\tau, e_1, e_2) \mid (e_1 \hookrightarrow^\omega \wedge e_2 \hookrightarrow^\omega) \vee (\exists (v_1, v_2) \in \overline{R}(\tau). e_1 \hookrightarrow^* v_1 \wedge e_2 \hookrightarrow^* v_2) \vee (\exists (\tau', e_1', e_2') \in \mathbf{S}(R, R). \exists (K_1, K_2) \in \mathbf{K}(R)(\tau', \tau). e_1 \hookrightarrow^* K_1[e_1'] \wedge e_2 \hookrightarrow^* K_2[e_2'])\}$	$\mathbf{E}(R) := \{(\tau, n_1, e_1, n_2, e_2) \mid (n_1, e_1 \hookrightarrow^\omega \wedge n_2, e_2 \hookrightarrow^\omega) \vee (\exists (n_1', v_1, n_2', v_2) \in \overline{R}(\tau). n_1, e_1 \hookrightarrow^* n_1', v_1 \wedge n_2, e_2 \hookrightarrow^* n_2', v_2) \vee (\exists (\tau', n_1', e_1', n_2', e_2') \in \mathbf{S}(R, R). \exists (_, K_1, _, K_2) \in \mathbf{K}(R)(\tau', \tau). n_1, e_1 \hookrightarrow^* n_1', K_1[e_1'] \wedge n_2, e_2 \hookrightarrow^* n_2', K_2[e_2'])\}$
$R^{\mathbf{S}} := \{(\tau, e_1, e_2) \mid \exists (e_1', e_2') \in R(\tau). e_1 \hookrightarrow e_1' \wedge e_2 \hookrightarrow e_2'\}$	$R^{\mathbf{S}} := \{(\tau, n_1, e_1, n_2, e_2) \mid \exists (n_1', e_1', n_2', e_2') \in R(\tau). n_1, e_1 \hookrightarrow n_1', e_1' \wedge n_2, e_2 \hookrightarrow n_2', e_2'\}$
consistent(L) := $\forall G \in \text{GK}(L). \mathbf{S}(L(G), G) \subseteq \mathbf{E}(G)^{\mathbf{S}}$	consistent(L) := $\forall G \in \text{GK}(L). \mathbf{S}(L(G), G) \subseteq \mathbf{E}(G)^{\mathbf{S}}$
$R(\cdot) := \{(\emptyset, \emptyset)\}$	$R(\cdot) := \{(\cdot, \emptyset, \cdot, \emptyset)\}$
$R(x:\tau, \Gamma) := \{(\gamma_1[x \mapsto v_1], \gamma_2[x \mapsto v_2]) \mid (v_1, v_2) \in R(\tau) \wedge (\gamma_1, \gamma_2) \in R(\Gamma)\}$	$R(x:\tau, \Gamma) := \{(n_1 :: N_1, \gamma_1[x \mapsto v_1], n_2 :: N_2, \gamma_2[x \mapsto v_2]) \mid (n_1, v_1, n_2, v_2) \in R(\tau) \wedge (N_1, \gamma_1, N_2, \gamma_2) \in R(\Gamma)\}$
$\Gamma \vdash e_1 \sim_L e_2 : \tau := \forall G \in \text{GK}(L). \forall (\gamma_1, \gamma_2) \in \overline{G}(\Gamma). (\tau, \gamma_1 e_1, \gamma_2 e_2) \in \mathbf{E}(G)$	$\Gamma \vdash e_1 \sim_{L, f_1, f_2} e_2 : \tau := \forall G \in \text{GK}(L). \forall (N_1, \gamma_1, N_2, \gamma_2) \in \overline{G}(\Gamma). (\tau, f_1(N_1), \gamma_1 e_1, f_2(N_2), \gamma_2 e_2) \in \mathbf{E}(G)$
$\Gamma \vdash e_1 \sim e_2 : \tau := \exists L. \text{consistent}(L) \wedge \Gamma \vdash e_1 \sim_L e_2 : \tau$	$\Gamma \vdash e_1 \sim e_2 : \tau := \exists L, f_1, f_2. \text{consistent}(L) \wedge \Gamma \vdash e_1 \sim_{L, f_1, f_2} e_2 : \tau$

Fig. 2. Models for λ^μ : Original PB model (left) and new SPB model with logical reduction (right).

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 consistent(L) and, say, $(\text{int}, 42, 42) \in \overline{G^{\sharp}}$ we know $(\text{int}, (\lambda y. 0) 42, (\lambda x. (\lambda y. 1) x) 42) \in \mathbf{E}(G^{\sharp})^{\mathbf{S}}$, i.e., $(\text{int}, 0, (\lambda y. 1) 42) \in \mathbf{E}(G^{\sharp})$. Since 0 is a value, this can only mean that $(\lambda y. 1) 42$ reduces to a related value, i.e., $(\text{int}, 0, 1) \in \overline{G^{\sharp}}$, which by definition of $\overline{G^{\sharp}}$ is false. ■

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 as follows:

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

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 $(\tau', v'_1, v'_2) \in \overline{G}$. For any $(v_1, v_2) \in G(\tau' \rightarrow \tau)$, we must show $(\tau, v_1 v'_1, (\lambda x. v_2 x) v'_2) \in \mathbf{E}(G)^\S$, i.e., $(\tau, e_1, e_2) \in \mathbf{E}(G)$, where $v_1 v'_1 \hookrightarrow e_1$ and $(\lambda x. v_2 x) v'_2 \hookrightarrow e_2$. The trouble is that, while we know that $e_2 = v_2 v'_2$, we have no way of knowing whether e_1 even exists. It is entirely possible, for instance, that v_1 is the integer 5, in which case $v_1 v'_1 \not\hookrightarrow$.

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) that it obtains from G , there is no way to know whether applications of such values reduce to well-behaved terms, or even if they are able to take a step of reduction at all.

4.2 Guardedness Revisited

As discussed at the end of Section 3, this requirement of “taking a step” in the definition of consistency is crucial in ensuring the soundness of PBs because it guarantees that the coinduction is suitably *guarded*. As the failed proof attempt 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 $(\tau, v_1 v'_1, v_2 v'_2) \in \mathbf{E}(G)$, which follows from $(\tau' \rightarrow \tau, v_1, v_2) \in G$ and $(\tau', v'_1, v'_2) \in \overline{G}$ (both given), and $(\tau, \tau, \bullet, \bullet) \in \mathbf{K}(G)$.

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.

4.3 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 $(\tau, n_1, v_1, n_2, v_2)$ rather than just (τ, v_1, v_2) . When proving consistency for such an item, i.e., when showing that the applications 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 in the introduction—to the subgoal of showing that the applications are in the \mathbf{E} relation. Using the “stuck” case, the \mathbf{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. (The exact changes to the model will be explained in a moment.) This logical reduction relation, operating on an expression and its budget, is defined as follows.

Definition 2 (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.

4.4 Stuttering Parametric Bisimulations (SPBs) for λ^μ

Our new stuttering parametric bisimulation (SPB) model for λ^μ is given in the right column of Figure 2, adjacent to the old PB model so that it is easy to see the (modest) changes: ignoring the stutter budget, there is no difference.

First, wherever the old model related two expressions, the new model additionally carries their budget. One may think of an expression and its part of the budget as a *logical expression*, in which case the change is that the new model relates logical expressions. Next, the value closure \overline{R} of a relation R does not care about the budget, except for function types, where it is passed on unmodified. For any other type, the closure relates values at any budget. (For the sake of readability we use an underscore, $_$, to stand for a fresh existentially quantified meta variable. When occurring in a binding position, such as the left side of a set comprehension, an underscore acts as a wildcard and matches anything.) The reason for this is that what really matters is the budget of functions, because we have to show consistency for them. If a pair or sum contains a function, then in order to access that function one already has to take a step (projection or case analysis, respectively), so progress is ensured regardless of the budget.

The definitions of LK and $\text{GK}(L)$ are slightly modified from those in the PB model, in order to stipulate the condition of *budget monotonicity*. This property says that two related values must stay related when their budget is increased (of

course, since it is always safe to give them more stuttering steps than they actually need). Unlike before, however, LK does not require the values to be proper syntactic function values. This is important because we want to be able to define local knowledges that, as in the proof sketch for the eta law, re-export values from the current global knowledge, which of course may be arbitrary junk. The reason for this restriction in the PB model had to do with an issue in the proof of transitivity, and we will see in Section 6.1 how we now instead make use of the stutter budget to resolve that issue.

Turning to the definition of the **S** relation, note that the only thing that counts is the budget of the *functions* being applied, not that of their arguments or continuations. This is important in the proof of transitivity. Intuitively, it is also sufficient because, before they can access their arguments or return to their continuations, the functions will have to take a physical beta-reduction step anyways. In the definition of the continuation relation, on the other hand, notice that the continuations' budgets are *added* to those of their input values. This corresponds to the idea that the continuations may stutter k_i times before passing the “hot potato” to their input values, which in turn demand stutter budgets of n_i .

Otherwise, the definition of the expression relation and of consistency are the same as before, just with physical reduction replaced by logical reduction.

Finally, in the equivalence judgment, we say that, besides a consistent local knowledge, there must exist a budget at which the programs are related. This budget may depend on the budgets of the values that are plugged in for free variables. It is not sufficient to just add up those budgets, because variables may occur multiple times and thus the dependency may not be linear. We therefore ask for a pair of budget functions f_1 and f_2 , which, given a list of the input budgets, calculate e_1 and e_2 's part of the final budget, respectively.

4.5 Eta Revisited

Using this model, we can now prove the eta law along the lines of the earlier attempt.

Theorem 1. $f : \tau' \rightarrow \tau \vdash f \sim (\lambda x. f x) : \tau' \rightarrow \tau$

Proof: We define the local knowledge L_η as follows:

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

We first show $f : \tau' \rightarrow \tau \vdash f \sim_{L, \text{hd}+1, \tilde{0}} (\lambda x. f x) : \tau' \rightarrow \tau$, where hd returns the first (here: only) element of a list, $\text{hd}+1$ is short for the meta function $\lambda l. \text{hd}(l) + 1$, and $\tilde{0}$ is the constant-0 function.

- Suppose $G \in \text{GK}(L_\eta)$ and $(n_1, v_1, n_2, v_2) \in G(\tau' \rightarrow \tau)$.
- We must show $(n_1+1, v_1, 0, \lambda x. v_2 x) \in \mathbf{E}(G)(\tau' \rightarrow \tau)$.
- This follows from $L_\eta(G) \subseteq G$ by construction of L_η .

It remains to show $\text{consistent}(L_\eta)$.

- Suppose $G \in \text{GK}(L_\eta)$, $(n_1, v_1, -, v_2) \in G(\tau' \rightarrow \tau)$, $n'_1 > n_1$, n'_2 arbitrary, and $(-, v'_1, -, v'_2) \in \overline{G}(\tau')$.
- We must show $(n'_1, v_1 v'_1, n'_2, (\lambda x. v_2 x) v'_2) \in \mathbf{E}(G)^\S(\tau)$.

$$\begin{aligned} \text{DRel} &:= \mathbb{P}((\text{CTyF} \times \mathbb{N} \times \text{CVal} \times \mathbb{N} \times \text{CVal}) \\ &\quad \uplus (\text{CTy} \times \mathbb{N} \times \text{CCont} \times \mathbb{N} \times \text{CCont})) \\ \text{ERel} &:= \mathbb{P}(\mathbb{N} \times \text{CExp} \times \mathbb{N} \times \text{CExp}) \end{aligned}$$

$$\overline{R}(\text{cont } \tau) := \{(_, \text{cont } K_1, _, \text{cont } K_2) \mid (_, K_1, _, K_2) \in R(\tau)\}$$

$$\text{LK} := \{L \in \text{DRel} \rightarrow \text{DRel} \mid \forall R. \forall (\tau, n_1, d_1, n_2, d_2) \in L(R). \\ \forall R' \supseteq R, n'_1 \geq n_1, n'_2 \geq n_2. (\tau, n'_1, d_1, n'_2, d_2) \in L(R')\}$$

$$\text{GK}(L) := \{G \in \text{DRel} \mid G \supseteq L(G) \wedge (0, \bullet, 0, \bullet) \in G(\text{int}) \wedge \\ \forall (n_1, d_1, n_2, d_2) \in G. \forall n'_1 \geq n_1, n'_2 \geq n_2. \\ (n'_1, d_1, n'_2, d_2) \in G\}$$

$$\mathbf{S} \in \text{DRel} \times \text{DRel} \rightarrow \text{ERel}$$

$$\begin{aligned} \mathbf{S}(R_c, R) &:= \{(k_1 + n_1, K_1[v_1], k_2 + n_2, K_2[v_2]) \mid \exists \tau. \\ &\quad (\tau, k_1, K_1, k_2, K_2) \in R_c \wedge (\tau, n_1, v_1, n_2, v_2) \in \overline{R}\} \cup \\ &\quad \{(n_1, K_1[v_1 v'_1], n_2, K_2[v_2 v'_2]) \mid \exists \tau, \tau'. \\ &\quad (\tau' \rightarrow \tau, n_1, v_1, n_2, v_2) \in R_c \wedge \\ &\quad (\tau', -, v'_1, -, v'_2) \in \overline{R} \wedge (\tau, -, K_1, -, K_2) \in R\} \end{aligned}$$

$$\mathbf{E} \in \text{DRel} \rightarrow \text{ERel}$$

$$\begin{aligned} \mathbf{E}(R) &:= \{(n_1, e_1, n_2, e_2) \mid (e_1 \hookrightarrow^\omega \wedge e_2 \hookrightarrow^\omega) \vee \\ &\quad (\exists (n'_1, e'_1, n'_2, e'_2) \in \mathbf{S}(R, R). \\ &\quad n_1, e_1 \hookrightarrow^* n'_1, e'_1 \wedge n_2, e_2 \hookrightarrow^* n'_2, e'_2)\} \end{aligned}$$

$$R^\S := \{(n_1, e_1, n_2, e_2) \mid \exists (n'_1, e'_1, n'_2, e'_2) \in R. \\ n_1, e_1 \hookrightarrow n'_1, e'_1 \wedge n_2, e_2 \hookrightarrow n'_2, e'_2\}$$

$$\text{consistent}(L) := \forall G \in \text{GK}(L). \mathbf{S}(L(G), G) \subseteq \mathbf{E}(G)^\S$$

$$\Gamma \vdash e_1 \sim_{L, f_1, f_2} e_2 : \tau := \forall G \in \text{GK}(L).$$

$$\forall (N_1, \gamma_1, N_2, \gamma_2) \in \overline{G}(\Gamma). \forall (k_1, K_1, k_2, K_2) \in G(\tau). \\ (f_1(N_1) + k_1, K_1[\gamma_1 e_1], f_2(N_2) + k_2, K_2[\gamma_2 e_2]) \in \mathbf{E}(G)$$

$$\Gamma \vdash e_1 \sim e_2 : \tau :=$$

$$\exists L, f_1, f_2. \text{consistent}(L) \wedge \Gamma \vdash e_1 \sim_{L, f_1, f_2} e_2 : \tau$$

Fig. 3. Stuttering parametric bisimulations with continuation knowledge.

- After taking a stutter step on the left side and a physical step on the right, it thus suffices to show $(n_1, v_1 v'_1, -, v_2 v'_2) \in \mathbf{E}(G)(\tau)$.
- Since $(-, \bullet, -, \bullet) \in \mathbf{K}(G)(\tau, \tau)$ (trivial to prove), we are done if we can show $(n_1, v_1 v'_1, -, v_2 v'_2) \in \mathbf{S}(G, G)(\tau)$.
- Indeed, this follows from $(n_1, v_1, -, v_2) \in G(\tau' \rightarrow \tau)$ and $(-, v'_1, -, v'_2) \in \overline{G}(\tau')$. \square

5 – STUTTERING PARAMETRIC BISIMULATIONS (SECOND STEP: SUPPORTING CONTINUATIONS)

In the previous section, we have seen how to enrich the PB model such that it validates the eta law. In this section, we will see how the added machinery, the stutter budget, enables us to also add support for control effects, such as first-class continuations.

The crucial characteristic of a language with control effects is that its semantics is context-sensitive, by which we mean that the following property does not hold universally:

$$e \hookrightarrow e' \implies K[e] \hookrightarrow K[e']$$

It is easy to check (see Figure 1) that λ^μ satisfies this property, while, due to callcc , λ_{cc}^μ does not.

In the model from Section 4 (as well as in the PB model),

proving two programs equivalent involves executing them directly in the empty context (*i.e.*, without a continuation). However, without the above property, this is clearly unsound. We therefore have to “contextualize” the model such that programs are always executed in an (arbitrary) evaluation context. We wish do this in a way that (1) is independent of first-class continuations—in order to scale to other context-sensitive features—and (2) requires only a modest change to actually support first-class continuations, *i.e.*, to obtain a sound model for λ_{cc}^μ .

5.1 Contextualizing the Model

Since we already have a continuation relation (\mathbf{K}), one may be tempted to use that in order to contextualize the model. However, this approach does not scale, as we will discuss in Section 7. Instead, the key to achieving the above two goals lies in allowing our knowledges to relate not only values but also continuations (thus rendering the \mathbf{K} relation obsolete). The final model is presented in Figure 3. For now, ignore the budget-related aspects and the parts highlighted in red.

As can be seen, we define local and global knowledges the same as before except that now they can relate both values and continuations (see \mathbf{DRel}). In the program equivalence judgment (at the bottom of the figure), we no longer consider the programs in isolation but in the context of (essentially arbitrary) continuations K_1, K_2 drawn from the global knowledge. Of course, we need to adapt the expression relation to account for this. In the new definition of \mathbf{E} , the termination case has disappeared. Instead, the \mathbf{S} relation has been extended with a “return” case allowing the two expressions to reduce to related values v_i inside related continuations K_i . Since these K_i may be the initially given continuations, this subsumes the old termination case. Due to this generalization, the \mathbf{E} and \mathbf{S} relations need no longer be type-indexed.

Also observe that, in the “stuck function calls” case, the condition that the continuations of the calls must be related has shifted from \mathbf{E} into \mathbf{S} (explained below) and is now using the continuation knowledge. The earlier continuation relation \mathbf{K} , derived from \mathbf{E} , has thus disappeared. A remarkable consequence is that the definition of the \mathbf{E} relation is no longer (co)recursive. Coinductive reasoning is now done entirely through the local knowledge and consistency.

The latter is defined the same way as before in terms of \mathbf{S} and \mathbf{E} . Of course, the new “return” case in \mathbf{S} results in an additional burden in proving consistency: whenever one relates two continuations in one’s local knowledge, one will be forced to prove that they behave equivalently when given related values. Furthermore, in proving consistency for function values, we now apply the applications inside continuations, since the condition about continuations has been shifted from \mathbf{E} to \mathbf{S} .

Why the Stutter Budget Matters. We now focus on the budget-related aspects of this model. First of all, observe that the treatment of the budget here very much follows that in the previous model. The only difference is that, since programs are now being run inside continuations, we add the budget of

these continuations to that of the programs (in the program equivalence judgment, and, via \mathbf{S} , in consistency).

More interesting is why, as we claimed earlier, the stutter budget plays an essential role in facilitating the addition of continuation knowledge. The answer is simple: without the budget, the above contextualization would rule out many basic equivalences, rendering the model all but useless. Assume for a moment that we contextualized the PB model (Section 3) rather than the SPB model (Section 4), *i.e.*, imagine the model in Figure 3 had no budget and was using physical rather than logical reduction. Now consider the example equivalence

$$f : \text{int} \rightarrow \text{int} \vdash \text{inl}(f\ 0) \sim \text{inl}(f\ 0) : \text{int} + \tau,$$

which certainly ought to hold (for any τ), and let us attempt to prove it. We define the following local knowledge, whose value part is empty:

$$L(R) := \{(\text{int}, K_1[\text{inl}\ \bullet], K_2[\text{inl}\ \bullet]) \mid (K_1, K_2) \in R(\text{int} + \tau)\}$$

Using the function call case in \mathbf{S} , it is easy to show that the two open programs are equivalent relative to L . It then remains to establish L ’s consistency. So suppose $G \in \mathbf{GK}(L)$, $(K_1, K_2) \in G(\text{int} + \tau)$, and $(v_1, v_2) \in \overline{G}(\text{int})$. We must show that there is $(e_1, e_2) \in \mathbf{E}(G)$ with $K_i[\text{inl}\ v_i] \hookrightarrow e_i$. Since $\text{inl}\ v_i$ is a value and we know nothing about K_i , we are stuck. As in the trouble with the eta law, the problem once again is that the guardedness condition of requiring terms—in this case, $K_i[\text{inl}\ v_i]$ —to take a *physical* step to e_i is too strong.

Also as before, the solution is to use the stutter budget to support a logical notion of progress, leading to a valid proof of the equivalence via the SPB model from Figure 3.

Proof: Define L as follows.

$$L(R) := \{(\text{int}, k'_1, K_1[\text{inl}\ \bullet], k'_2, K_2[\text{inl}\ \bullet]) \mid \exists k_1 < k'_1, k_2 < k'_2. (k_1, K_1, k_2, K_2) \in R(\text{int} + \tau)\}$$

We show $f : \text{int} \rightarrow \text{int} \vdash \text{inl}(f\ 0) \sim_{L, \text{hd}, \text{hd}} \text{inl}(f\ 0) : \text{int} + \tau$.

- Suppose $G \in \mathbf{GK}(L)$, $(n_1, v_1, n_2, v_2) \in \overline{G}(\text{int} \rightarrow \text{int})$ and $(k_1, K_1, k_2, K_2) \in G(\text{int} + \tau)$. We must show:

$$(n_1 + k_1, K_1[\text{inl}(v_1\ 0)], n_2 + k_2, K_2[\text{inl}(v_2\ 0)]) \in \mathbf{E}(G)$$
- Immediately using the “function call” case in \mathbf{S} , it suffices to show $(_, K_1[\text{inl}\ \bullet], _, K_2[\text{inl}\ \bullet]) \in G(\text{int})$.
- This follows from $G \supseteq L(G)$ by construction of L .

It remains to prove consistent(L).

- So suppose $G \in \mathbf{GK}(L)$, $(k_1, K_1, k_2, K_2) \in G(\text{int} + \tau)$, $k'_1 > k_1$, $k'_2 > k_2$, and $(n_1, v_1, n_2, v_2) \in \overline{G}(\text{int})$.
- To show: $(k'_1 + n_1, K_1[\text{inl}\ v_1], k'_2 + n_2, K_2[\text{inl}\ v_2]) \in \mathbf{E}(G)^\S$
- Taking a stutter step on either side, it suffices to show $(k_1, K_1[\text{inl}\ v_1], k_2, K_2[\text{inl}\ v_2]) \in \mathbf{E}(G)$, which is obvious due to $\mathbf{S}(G, G) \subseteq \mathbf{E}(G)$. \square

5.2 Supporting First-Class Continuations

As mentioned before, the contextualization is independent of first-class continuations. Consequently, the model in Figure 3 *excluding* the two parts **highlighted in red** is still a (sound) model for λ^μ . In order to obtain one for λ_{cc}^μ , all one then needs to do is include those parts.

The first one extends the definition of the value closure to give meaning to continuation types: two continuation values are related at type $\text{cont } \tau$ iff the underlying continuations are related at τ by (the continuation part of) the global knowledge. Similarly to other cases of the value closure, the budget is ignored. The second part is an additional requirement that any global knowledge must relate the empty continuations at the system type int , for any budget. This is necessary in the presence of the isolate construct, which allows a programmer to make up continuations that escape the initial ones.

5.3 Example: callcc in a Loop

Of course, the eta law still holds in this generalized model. We could show the proof here but it is almost the same as the one in Section 4 (in fact, even the same local knowledge can be used). Instead, we show the proof of an interesting equivalence involving callcc. The example is adapted from Størring and Lassen [18]. Consider the following two programs:

$$\begin{aligned} \tau &:= \mu\alpha. \text{cont int} \rightarrow \alpha \\ F_v &:= \text{fix } f(x:\tau):\text{int}. f(\text{unroll } x \ v) \\ e_1 &:= |\lambda y:\tau. \text{callcc}_{\text{int}}(k. F_k \ y)| \\ e_2 &:= |\text{fix } f(x:\tau):\text{int}. \text{callcc}_{\text{int}}(k. f(\text{unroll } x \ k))| \end{aligned}$$

When called, both e_1 and e_2 loop. Both also capture the current continuation. The difference is that e_2 captures it once (namely at the very beginning of its execution), while e_1 does it in every loop iteration. However, since the continuation does not actually change, e_1 is always capturing the same one (namely K_1), and so the two programs behave equivalently.

Formally, in order to prove $\vdash e_1 \sim e_2 : \tau \rightarrow \text{int}$ (for λ_{cc}^μ), we define the following local knowledge:

$$\begin{aligned} L(R) &:= \{(\tau \rightarrow \text{int}, _, e_1, _, e_2)\} \\ &\cup \{(\tau, _, K_1[F_{(\text{cont } K_1)} \bullet], _, K_2[e_2 \bullet]) \mid \\ &\quad (_, K_1, _, K_2) \in G(\text{int})\} \end{aligned}$$

(As obvious from this definition, the budget does not play an interesting role here.) We first show $\vdash e_1 \sim_{L, \bar{0}, \bar{0}} e_2 : \tau \rightarrow \text{int}$.

- Suppose $G \in \text{GK}(L)$ and $(k_1, K_1, k_2, K_2) \in G(\tau \rightarrow \text{int})$.
- We must show $(k_1, K_1[e_1], k_2, K_2[e_2]) \in \mathbf{E}(G)$.
- It suffices to show $(0, e_1, 0, e_2) \in \overline{G}(\tau \rightarrow \text{int})$, which holds due to $G \supseteq L(G)$.

It remains to show $\text{consistent}(L)$, which has two parts.

- 1) • Suppose $G \in \text{GK}(L)$, n_1 and n_2 arbitrary, (*) $(_, v_1, _, v_2) \in \overline{G}(\tau)$, and $(_, K_1, _, K_2) \in G(\text{int})$.
 - To show: $(n_1, K_1[e_1 \ v_1], n_2, K_2[e_2 \ v_2]) \in \mathbf{E}(G)^\S$
 - From (*) we know that there is $(_, v'_1, _, v'_2) \in \overline{G}(\text{cont int} \rightarrow \tau)$ such that $v_i = \text{roll } v'_i$.
 - Taking several physical steps, it thus suffices to show:
$$\begin{aligned} (_, K_1[F_{(\text{cont } K_1)}(v'_1(\text{cont } K_1))], \\ _, K_2[e_2(v'_2(\text{cont } K_2))]) \in \mathbf{S}(G, G) \end{aligned}$$
 - Since $(_, \text{cont } K_1, _, \text{cont } K_2) \in \overline{G}(\text{cont int})$, this follows with $(_, K_1[F_{(\text{cont } K_1)} \bullet], _, K_2[e_2 \bullet]) \in G(\tau)$, which holds due to $G \supseteq L(G)$.
- 2) • Suppose $G \in \text{GK}(L)$, $(_, K_1, _, K_2) \in G(\text{int})$, n_1 and n_2 arbitrary, and (*) $(m_1, v_1, m_2, v_2) \in \overline{G}(\tau)$.

- To show: $(n_1 + m_1, K_1[F_{(\text{cont } K_1)} v_1], n_2 + m_2, K_2[e_2 v_2]) \in \mathbf{E}(G)^\S$
- By (*) there is $(_, v'_1, _, v'_2) \in \overline{G}(\text{cont int} \rightarrow \tau)$ such that $v_i = \text{roll } v'_i$.
- Taking several physical steps, it thus suffices to show
$$\begin{aligned} (_, K_1[F_{(\text{cont } K_1)}(v'_1(\text{cont } K_1))], \\ _, K_2[e_2(v'_2(\text{cont } K_2))]) \in \mathbf{S}(G, G), \end{aligned}$$
which we do the same as in part (1).

5.4 Using Parameterized Coinduction

Writing down a suitable local knowledge at the beginning of a proof can be quite tedious for complex equivalences, even more so in the new model where one generally also has to add continuations. 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* [10, 14] to avoid this issue completely and instead write proofs in an incremental style, where we basically start with a knowledge containing just the programs in question, and extend it 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.

Definition 3. We define the wanted function $\mathfrak{f} \in \text{LK} \xrightarrow{\text{mon}} \text{LK}$.
 $\mathfrak{f}(L)(R) := \{(\tau, n_1, d_1, n_2, d_2) \mid \forall G \in \text{GK}(L). \\ G \supseteq R \implies \mathbf{S}(\{(\tau, n_1, d_1, n_2, d_2)\}, G) \subseteq \mathbf{E}(G)^\S\}$

Lemma 2. $L \subseteq \mathfrak{f}(L) \iff \text{consistent}(L)$

6 – METATHEORY

We briefly state the main meta-theoretical results for SPBs. They apply to both models for λ^μ (the one without continuation knowledge, and the one with), and to the model for λ_{cc}^μ . Their proofs pretty much follow those for PBs [8] (but see below for details on transitivity). A number of the compatibility lemmas needed in showing reflexivity and congruence are proven in the appendix. For further proofs, we refer the reader to our Coq formalization of SPBs (in the setting of a richer language).

Theorem 3 (Reflexivity, Symmetry, Transitivity, Congruence).

$$\begin{aligned} \frac{\Gamma \vdash p : \tau}{\Gamma \vdash |p| \sim |p| : \tau} \quad \frac{\Gamma \vdash e_2 \sim e_1 : \tau}{\Gamma \vdash e_1 \sim e_2 : \tau} \\ \frac{\Gamma \vdash e_1 \sim e_2 : \tau \quad \Gamma \vdash e_2 \sim e_3 : \tau}{\Gamma \vdash e_1 \sim e_3 : \tau} \\ \frac{\Gamma' \vdash e_1 \sim e_2 : \tau' \quad \vdash C : (\Gamma'; \tau') \rightsquigarrow (\Gamma; \tau)}{\Gamma \vdash |C|[e_1] \sim |C|[e_2] : \tau} \end{aligned}$$

Lemma 4 (Adequacy). If $\vdash e_1 \sim e_2 : \text{int}$, then either both $e_1 \hookrightarrow^\omega$ and $e_2 \hookrightarrow^\omega$, or both $e_1 \hookrightarrow^* n$ and $e_2 \hookrightarrow^* n$ for some integer value n .

Theorem 5 (Soundness).

$$\frac{\Gamma \vdash p_1 : \tau \quad \Gamma \vdash |p_1| \sim |p_2| : \tau \quad \Gamma \vdash p_2 : \tau}{\Gamma \vdash p_1 \sim_{\text{ctx}} p_2 : \tau}$$

6.1 Transitivity

As mentioned in the introduction, transitivity is of particular interest to us. We describe the proof of transitivity for PBs in detail in a technical report [9]. Developing this very complex proof required a lot of effort. Fortunately, adapting it to SPBs does not require many changes. We now briefly discuss the more interesting ones.

The transitivity proof must establish that, given consistent local knowledges L_1, L_2 with $\Gamma \vdash e_1 \sim_{L_1} e_2 : \tau$ and $\Gamma \vdash e_2 \sim_{L_2} e_3 : \tau$, there exists a consistent local knowledge L with $\Gamma \vdash e_1 \sim_L e_3 : \tau$. A key part of this proof is the “correct” decomposition of a global knowledge G for L into $G_{(1)}$ for L_1 and $G_{(2)}$ for L_2 , using which L is defined. In the PB case, we have (where \circ is relational composition):

$$L(R) := L_1(R_{(1)}) \circ L_2(R_{(2)})$$

The proof of L ’s consistency is very tricky. Amongst other things, it relies on the property that certain “bad” values can be related by a global knowledge but not by a local one. This is enforced in the PB model by a restriction on LK: values in a local knowledge must be syntactic functions, in contrast to a global knowledge. However, as explained in Section 4.4, we had to drop this discrimination in SPBs. As a consequence, we can no longer prove the consistency of such defined L .

Fortunately, the stutter budget enables a trick that lets us overcome this problem. The idea is to define L in a slightly different manner:

$$L(R) := (R_{(1)} \circ R_{(2)})^{++}$$

(Here $(-)^{++}$ increments both parts of each tuple’s budget by 1.) At a first glance, this seems very odd. Previously, since L was defined as the composition of L_1 and L_2 , we could rely on L_1 and L_2 ’s consistency in proving L ’s. Now, however, L is defined in terms of global knowledges which may contain junk. Fortunately, thanks to increasing the budget in the definition of L , this is not an issue and we can pretty easily prove L consistent. Of course, there is no free lunch, and the price to pay for this parlor trick is that other parts of the transitivity proof become even more subtle than they already were (and the model had to be designed very carefully to make them all go through). We intend to report on the details of this new transitivity proof in a future, extended version of this paper.

7 – DISCUSSION AND RELATED WORK

Modeling Higher-Order State and Abstract Types. Following in the footsteps of PBs, it is fairly straightforward to scale both the intermediate SPB model (Section 4) and the final

contextualized one (Section 5) to a language with polymorphism, abstract types, and general references. We have proven all the meta-theoretic results of the previous section for this full model, and mechanized the proofs in Coq. Our appendix presents the definition of this full model, together with an example application of it (namely, Dreyer *et al.*’s challenging “well-bracketed state change” example [6]).

Concerning the extensions to abstract types and state, perhaps the only significant difference between SPBs and PBs is a (minor) generalization of local knowledges, necessitated by SPBs’ relating of continuations in the local/global knowledge. In the full PB model, both local and global knowledges were indexed by states of a transition system used to describe invariants on a term’s local mutable state. This holds true for the full SPB model as well. The difference is that in PBs, a local knowledge was only able to query the global knowledge at the “current” state of the transition system, whereas in SPBs, when defining a local knowledge, it is important to be able to query whether some other continuations were related by the global knowledge in *previous* states of the transition system.

Contextualization. In Section 5 we commented that using the \mathbf{K} relation in order to contextualize the model does not scale. In particular, one might expect to see something like (ignoring budget and types)

$$\forall (K_1, K_2) \in \mathbf{K}(G). (K_1[\gamma_1 e_1], K_2[\gamma_2 e_2]) \in \mathbf{E}(G)$$

in the definition of program equivalence (and similarly for consistency). That is, one would consider the behavior of programs when run inside continuations related by \mathbf{K} . To understand why this approach does not scale, it is important to understand that in the model for the full language with state, equivalence is relative not only to a local knowledge, but to a whole *world* [1, 6, 8] that, amongst other things, limits the part of the heap that the programs may access. In particular, both the local term equivalence relation, \mathbf{E} , and the continuation relation, \mathbf{K} , require programs to conform to the world and are therefore *indexed* by a world W .

One key property then is that programs stay equivalent when the world is extended, *i.e.*, when access to additional parts of the heap is granted (intuitively, because the programs do not care about these parts). This is where things go wrong when using the contextualization sketched above. Assume two programs are equivalent in a world W , meaning (according to the definition sketched above) that they are related inside any continuations from \mathbf{K}_W . We would now have to show that the programs are also related when run inside continuations from $\mathbf{K}_{W'}$, for any larger world W' . However, since $\mathbf{K}_{W'}$ generally contains more continuations than \mathbf{K}_W —after all, those in $\mathbf{K}_{W'}$ may access a larger part of the heap than those in \mathbf{K}_W —this is clearly impossible to prove.

Well-Founded Bisimulations. Our technique of logical reduction is inspired by Namjoshi’s *well-founded* [15, 13] bisimulations, which were developed as an alternative formulation of *stuttering* bisimulations [4] 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.

Relational Reasoning About ML-Like Languages. As explained in the introduction, many methods have been developed for reasoning about equivalence in ML-like languages, but the most powerful methods to date (at least practically speaking, in terms of providing effective proof techniques) are step-indexed Kripke logical relations methods [1, 6, 7] and various bisimulation methods [11, 17, 19]. We proposed PBs [8] as a way of synthesizing, as it were, the best of both worlds, but our previous formulation of them was lacking in (1) its invalidation of the eta law, and (2) its inability to model control effects. Thus, for all their virtues, PBs could not be claimed to subsume other methods, since indeed they could not prove as many equivalences as, say, Dreyer *et al.*’s SKLRs [6], which are compatible with both control effects and eta.

In this paper, we have shown how the idea of logical reduction semantics employed by stuttering parametric bisimulations (SPBs) rectifies the inadequacies of PBs, thus offering a proof method on par with the state of the art. In addition, like PBs, SPBs continue to offer distinct advantages over SKLR and bisimulation methods, due to their transitivity and semantic nature, respectively. We hope to exploit those advantages in our future work on compositional compiler certification.

ACKNOWLEDGMENTS

The research was partially supported by the EC FET project ADVENT.

REFERENCES

- [1] A. Ahmed, D. Dreyer, and A. Rossberg, “State-dependent representation independence,” in *POPL*, 2009.
- [2] N. Benton and C.-K. Hur, “Biorthogonality, step-indexing and compiler correctness,” in *ICFP*, 2009.
- [3] L. Birkedal and R. W. Harper, “Constructing interpretations of recursive types in an operational setting,” *Information and Computation*, vol. 155, pp. 3–63, 1999.
- [4] M. C. Browne, E. M. Clarke, and O. Grumberg, “Characterizing finite Kripke structures in propositional temporal logic,” *Theor. Comput. Sci.*, vol. 59, no. 1-2, pp. 115–131, Jul. 1988.
- [5] D. Dreyer, G. Neis, and L. Birkedal, “The impact of higher-order state and control effects on local relational reasoning,” in *ICFP*, 2010.
- [6] D. Dreyer, G. Neis, and L. Birkedal, “The impact of higher-order state and control effects on local relational reasoning,” *J. Funct. Program.*, vol. 22, no. 4-5, pp. 477–528, 2012.
- [7] C.-K. Hur and D. Dreyer, “A Kripke logical relation between ML and assembly,” in *POPL*, 2011.
- [8] C.-K. Hur, D. Dreyer, G. Neis, and V. Vafeiadis, “The marriage of bisimulations and Kripke logical relations,” in *POPL*, 2012.
- [9] C.-K. Hur, G. Neis, D. Dreyer, and V. Vafeiadis, “The transitive composability of relation transition systems,” MPI-SWS, Tech. Rep. MPI-SWS-2012-002, May 2012.
- [10] C.-K. Hur, G. Neis, D. Dreyer, and V. Vafeiadis, “The power of parameterization in coinductive proof,” in *POPL*, 2013.
- [11] V. Koutavas and M. Wand, “Small bisimulations for reasoning about higher-order imperative programs,” in *POPL*, 2006.
- [12] S. B. Lassen and P. B. Levy, “Typed normal form bisimulation for parametric polymorphism,” in *LICS*, 2008.
- [13] P. Manolios, “Mechanical verification of reactive systems,” Ph.D. Thesis, Department of Computer Science, The University of Texas at Austin, Aug. 2001.
- [14] L. S. Moss, “Parametric corecursion,” *Theor. Comput. Sci.*, vol. 260, no. 1-2, pp. 139–163, Jun. 2001.
- [15] K. S. Namjoshi, “A simple characterization of stuttering bisimulation,” in *FSTTCS*, 1997, pp. 284–296.
- [16] D. Sangiorgi, “The lazy lambda calculus in a concurrency scenario,” *Information and Computation*, vol. 111, no. 1, pp. 120–153, 1994.
- [17] D. Sangiorgi, N. Kobayashi, and E. Sumii, “Environmental bisimulations for higher-order languages,” *TOPLAS*, vol. 33, no. 1, pp. 1–69, Jan. 2011.
- [18] K. Støvring and S. Lassen, “A complete, co-inductive syntactic theory of sequential control and state,” in *POPL*, 2007.
- [19] E. Sumii, “A complete characterization of observational equivalence in polymorphic λ -calculus with general references,” in *CSL*, 2009.
- [20] E. Sumii and B. Pierce, “A bisimulation for type abstraction and recursion,” *JACM*, vol. 54, no. 5, pp. 1–43, Oct. 2007.

APPENDIX A
SPBS FOR λ^μ AND λ_{cc}^μ

A.1 Languages λ^μ and λ_{cc}^μ

The fragment λ^μ is obtained by removing first-class continuations (the parts **highlighted in red**).

A.1.1 Statics.

$\sigma, \tau \in \text{Ty} ::= \alpha \mid \text{int} \mid \sigma_1 \times \sigma_2 \mid \sigma_1 + \sigma_2 \mid \sigma_1 \rightarrow \sigma_2 \mid \mu\alpha. \sigma \mid \text{cont } \sigma$
 $p \in \text{Prog} ::= x \mid n \mid p_1 \odot p_2 \mid \text{ifz } p \text{ then } p_1 \text{ else } p_2 \mid \langle p_1, p_2 \rangle \mid p.1 \mid p.2 \mid \text{inl}_\sigma p \mid \text{inr}_\sigma p \mid$
 (case p of $\text{inl } x \Rightarrow p_1 \mid \text{inr } x \Rightarrow p_2$) $\mid \text{fix } f(x:\sigma_1):\sigma_2. p \mid p_1 p_2 \mid \text{roll}_\sigma p \mid \text{unroll } p \mid$
 $\text{callcc}_\sigma(x. p) \mid \text{throw}_\sigma p_1 \text{ to } p_2 \mid \text{isolate } p$

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

$\Gamma \vdash p : \tau$

$$\frac{}{\Gamma \vdash x : \tau} \quad \frac{}{\Gamma \vdash n : \text{int}}$$

$$\frac{\Gamma \vdash p_1 : \text{int} \quad \Gamma \vdash p_2 : \text{int}}{\Gamma \vdash p_1 \odot p_2 : \text{int}} \quad \frac{\Gamma \vdash p_0 : \text{int} \quad \Gamma \vdash p_1 : \tau \quad \Gamma \vdash p_2 : \tau}{\Gamma \vdash \text{ifz } p_0 \text{ then } p_1 \text{ else } p_2 : \tau}$$

$$\frac{\Gamma \vdash p_1 : \tau_1 \quad \Gamma \vdash p_2 : \tau_2}{\Gamma \vdash \langle p_1, p_2 \rangle : \tau_1 \times \tau_2} \quad \frac{\Gamma \vdash p : \tau_1 \times \tau_2}{\Gamma \vdash p.1 : \tau_1} \quad \frac{\Gamma \vdash p : \tau_1 \times \tau_2}{\Gamma \vdash p.2 : \tau_2}$$

$$\frac{\Gamma \vdash p : \tau_1}{\Gamma \vdash \text{inl}_{\tau_2} p : \tau_1 + \tau_2} \quad \frac{\Gamma \vdash p : \tau_2}{\Gamma \vdash \text{inr}_{\tau_1} p : \tau_1 + \tau_2}$$

$$\frac{\Gamma \vdash p : \tau_1 + \tau_2 \quad \Gamma, x:\tau_1 \vdash p_1 : \tau \quad \Gamma, x:\tau_2 \vdash p_2 : \tau}{\Gamma \vdash \text{case } p \text{ of } \text{inl } x \Rightarrow p_1 \mid \text{inr } x \Rightarrow p_2 : \tau}$$

$$\frac{\Gamma, f:(\tau_1 \rightarrow \tau_2), x:\tau_1 \vdash p : \tau_2}{\Gamma \vdash \text{fix } f(x:\tau_1):\tau_2. p : \tau_1 \rightarrow \tau_2} \quad \frac{\Gamma \vdash p_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma \vdash p_2 : \tau_1}{\Gamma \vdash p_1 p_2 : \tau_2}$$

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

$$\frac{\Gamma, x:\text{cont } \tau \vdash p : \tau}{\Gamma \vdash \text{callcc}_\tau(x. p) : \tau} \quad \frac{\Gamma \vdash p' : \tau' \quad \Gamma \vdash p : \text{cont } \tau'}{\Gamma \vdash \text{throw}_\tau p' \text{ to } p : \tau} \quad \frac{\Gamma \vdash p : \tau \rightarrow \text{int}}{\Gamma \vdash \text{isolate } p : \text{cont } \tau}$$

A.1.2 Dynamics.

$$\begin{aligned}
v \in \text{Val} & ::= x \mid n \mid \langle v_1, v_2 \rangle \mid \text{inl } v \mid \text{inr } v \mid \text{fix } f(x).e \mid \text{roll } v \mid \text{cont } K \\
e \in \text{Exp} & ::= v \mid e_1 \odot e_2 \mid \text{ifz } e_0 \text{ then } e_1 \text{ else } e_2 \mid \langle e_1, e_2 \rangle \mid e.1 \mid e.2 \mid \text{inl } e \mid \text{inr } e \mid \\
& \quad (\text{case } e \text{ of inl } x \Rightarrow e_1 \mid \text{inr } x \Rightarrow e_2) \mid e_1 e_2 \mid \text{roll } e \mid \text{unroll } e \mid \\
& \quad \text{callcc } (x. e) \mid \text{throw } e_1 \text{ to } e_2 \mid \text{isolate } e \\
K \in \text{Cont} & ::= \bullet \mid K \odot e \mid v \odot K \mid \text{ifz } K \text{ then } e_1 \text{ else } e_2 \mid \langle K, e \rangle \mid \langle v, K \rangle \mid K.1 \mid K.2 \mid \\
& \quad \text{inl } K \mid \text{inr } K \mid (\text{case } K \text{ of inl } x \Rightarrow e_1 \mid \text{inr } x \Rightarrow e_2) \mid K e \mid v K \mid \text{roll } K \mid \\
& \quad \text{unroll } K \mid \text{throw } K \text{ to } e \mid \text{throw } v \text{ to } K \mid \text{isolate } K
\end{aligned}$$

$$\boxed{e \hookrightarrow e}$$

$$\begin{aligned}
K[n_1 \odot n_2] & \hookrightarrow K[n] & (n = \llbracket n_1 \odot n_2 \rrbracket) \\
K[\text{ifz } 0 \text{ then } e_1 \text{ else } e_2] & \hookrightarrow K[e_1] \\
K[\text{ifz } n \text{ then } e_1 \text{ else } e_2] & \hookrightarrow K[e_2] & (n \neq 0) \\
K[\langle v_1, v_2 \rangle.1] & \hookrightarrow K[v_1] \\
K[\langle v_1, v_2 \rangle.2] & \hookrightarrow K[v_2] \\
K[\text{case inl } v \text{ of inl } x \Rightarrow e_1 \mid \text{inr } x \Rightarrow e_2] & \hookrightarrow K[e_1[v/x]] \\
K[\text{case inr } v \text{ of inl } x \Rightarrow e_1 \mid \text{inr } x \Rightarrow e_2] & \hookrightarrow K[e_2[v/x]] \\
K[(\text{fix } f(x).e) v] & \hookrightarrow K[e[(\text{fix } f(x).e)/f, v/x]] \\
K[\text{unroll } (\text{roll } v)] & \hookrightarrow K[v] \\
K[\text{callcc } (x. e)] & \hookrightarrow K[e[\text{cont } K/x]] \\
K[\text{throw } v \text{ to cont } K'] & \hookrightarrow K'[v] \\
K[\text{isolate } v] & \hookrightarrow K[\text{cont } (v \bullet)]
\end{aligned}$$

A.2 Simple SPB Model for λ^μ

A.2.1 Definition.

$$\boxed{n, e \hookrightarrow n', e'}$$

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

$$\begin{aligned}
\text{CVal} & := \{v \in \text{Val} \mid \text{fv}(v) = \emptyset\} \\
\text{CExp} & := \{e \in \text{Exp} \mid \text{fv}(e) = \emptyset\} \\
\text{CCont} & := \{K \in \text{Cont} \mid \text{fv}(K[\langle \rangle]) = \emptyset\} \\
\text{CTy} & := \{\tau \in \text{Ty} \mid \text{fv}(\tau) = \emptyset\} \\
\text{CTyF} & := \{\tau_1 \rightarrow \tau_2 \mid \tau_1, \tau_2 \in \text{CTy}\}
\end{aligned}$$

$$\begin{aligned}
\text{VRelF} & := \mathbb{P}(\text{CTyF} \times \mathbb{N} \times \text{CVal} \times \mathbb{N} \times \text{CVal}) \\
\text{KRel} & := \mathbb{P}(\text{CTy} \times \text{CTy} \times \mathbb{N} \times \text{CCont} \times \mathbb{N} \times \text{CCont}) \\
\text{VRel} & := \mathbb{P}(\text{CTy} \times \mathbb{N} \times \text{CVal} \times \mathbb{N} \times \text{CVal}) \\
\text{ERel} & := \mathbb{P}(\text{CTy} \times \mathbb{N} \times \text{CExp} \times \mathbb{N} \times \text{CExp})
\end{aligned}$$

$$\overline{(-)} \in \text{VRelF} \rightarrow \text{VRel}$$

$$\begin{aligned}
\overline{R}(\text{int}) & := \{(n_1, m, n_2, m)\} \\
\overline{R}(\tau \times \tau') & := \{(n_1, \langle v_1, v'_1 \rangle, n_2, \langle v_2, v'_2 \rangle) \mid \exists m_1, m_2. (m_1, v_1, m_2, v_2) \in \overline{R}(\tau)\} \\
& \quad \cap \{(n_1, \langle v_1, v'_1 \rangle, n_2, \langle v_2, v'_2 \rangle) \mid \exists m'_1, m'_2. (m'_1, v'_1, m'_2, v'_2) \in \overline{R}(\tau')\} \\
\overline{R}(\tau + \tau') & := \{(n_1, \text{inl } v_1, n_2, \text{inl } v_2) \mid \exists m_1, m_2. (m_1, v_1, m_2, v_2) \in \overline{R}(\tau)\} \\
& \quad \cup \{(n_1, \text{inr } v'_1, n_2, \text{inr } v'_2) \mid \exists m'_1, m'_2. (m'_1, v'_1, m'_2, v'_2) \in \overline{R}(\tau')\} \\
\overline{R}(\tau \rightarrow \tau') & := \{(n_1, v_1, n_2, v_2) \in R(\tau \rightarrow \tau')\} \\
\overline{R}(\mu\alpha.\sigma) & := \{(n_1, \text{roll } v_1, n_2, \text{roll } v_2) \mid \exists m_1, m_2. (m_1, v_1, m_2, v_2) \in \overline{R}(\sigma[\mu\alpha.\sigma/\alpha])\}
\end{aligned}$$

$$\begin{aligned}
\text{LK} & := \{L \in \text{VRelF} \rightarrow \text{VRelF} \mid \forall R. \forall (\tau, n_1, v_1, n_2, v_2) \in L(R). \\
& \quad \forall R' \supseteq R, n'_1 \geq n_1, n'_2 \geq n_2. (\tau, n'_1, v_1, n'_2, v_2) \in L(R')\}
\end{aligned}$$

$$\begin{aligned}
\text{GK}(L) & := \{G \in \text{VRelF} \mid L(G) \subseteq G \wedge \\
& \quad \forall (\tau, n_1, v_1, n_2, v_2) \in G. \forall n'_1 \geq n_1, n'_2 \geq n_2. (\tau, n'_1, v_1, n'_2, v_2) \in G\}
\end{aligned}$$

$$\begin{aligned} S &\in \text{VRelF} \times \text{VRelF} \rightarrow \text{ERel} \\ S(R, G)(\tau) &:= \{(n_1, v_1, v'_1, n_2, v_2, v'_2) \mid \\ &\quad \exists \tau', m_1, m_2. (n_1, v_1, n_2, v_2) \in R(\tau' \rightarrow \tau) \wedge (m_1, v'_1, m_2, v'_2) \in \overline{G}(\tau')\} \end{aligned}$$

$$\begin{aligned} E &\in \text{VRelF} \times \text{VRelF} \rightarrow \text{ERel} \\ E(R, G)(\tau) &:= \{(n_1, e_1, n_2, e_2) \mid (e_1 \hookrightarrow^\omega \wedge e_2 \hookrightarrow^\omega) \vee \\ &\quad (\exists (n'_1, v_1, n'_2, v_2) \in \overline{G}(\tau). n_1, e_1 \hookrightarrow^* n'_1, v_1 \wedge n_2, e_2 \hookrightarrow^* n'_2, v_2) \vee \\ &\quad (\exists (\tau', n'_1, e'_1, n'_2, e'_2) \in S(R, G). \exists (k_1, K_1, k_2, K_2) \in K(R, G)(\tau', \tau). \\ &\quad \quad n_1, e_1 \hookrightarrow^* n'_1, K_1[e'_1] \wedge n_2, e_2 \hookrightarrow^* n'_2, K_2[e'_2])\} \end{aligned}$$

$$\begin{aligned} K &\in \text{VRelF} \times \text{VRelF} \rightarrow \text{KRel} \\ K(R, G)(\tau', \tau) &:= \{(k_1, K_1, k_2, K_2) \mid \forall (n_1, v_1, n_2, v_2) \in \overline{G}(\tau'). \\ &\quad (k_1 + n_1, K_1[v_1], k_2 + n_2, K_2[v_2]) \in E(R, G)(\tau)\} \end{aligned}$$

$$R^{\$} := \{(\tau, n_1, e_1, n_2, e_2) \mid \exists (n'_1, e'_1, n'_2, e'_2) \in R(\tau). n_1, e_1 \hookrightarrow n'_1, e'_1 \wedge n_2, e_2 \hookrightarrow n'_2, e'_2\}$$

$$\text{consistent}(L) := \forall G \in \text{GK}(L). S(L(G), G) \subseteq E(G, G)^{\$}$$

$$\Gamma \vdash e_1 \sim e_2 : \tau := \exists L, f_1, f_2. \text{consistent}(L) \wedge \Gamma \vdash e_1 \sim_{L, f_1, f_2} e_2 : \tau$$

$$\begin{aligned} \Gamma \vdash e_1 \sim_{L, f_1, f_2} e_2 : \tau &:= \forall G \in \text{GK}(L). \forall (N_1, \gamma_1, N_2, \gamma_2) \in \overline{G}(\Gamma). \\ &\quad (f_1(N_1), \gamma_1 e_1, f_2(N_2), \gamma_2 e_2) \in E(G, G)(\tau) \end{aligned}$$

$$\begin{aligned} R(\cdot) &:= \{(\cdot, \text{id}, \cdot, \text{id})\} \\ R(x:\tau, \Gamma) &:= \{(n_1 :: N_1, \gamma_1[x \mapsto v_1], n_2 :: N_2, \gamma_2[x \mapsto v_2]) \mid \\ &\quad (n_1, v_1, n_2, v_2) \in R(\tau) \wedge (N_1, \gamma_1, N_2, \gamma_2) \in R(\Gamma)\} \end{aligned}$$

A.2.2 Properties.

Lemma 6. If

- 1) $(\tau, n'_1, e'_1, n'_2, e'_2) \in E(R, G)$
- 2) $n_1, e_1 \hookrightarrow^* n'_1, e'_1$
- 3) $n_2, e_2 \hookrightarrow^* n'_2, e'_2$

then $(\tau, n_1, e_1, n_2, e_2) \in E(R, G)$.

Lemma 7.

- $\text{GK}(L) \cap \text{GK}(L') = \text{GK}(L \cup L')$
- $S(R, G) \cup S(R', G) = S(R \cup R', G)$
- $\text{consistent}(L) \wedge \text{consistent}(L') \implies \text{consistent}(L \cup L')$
- $\Gamma \vdash e_1 \sim_{L, f_1, f_2} e_2 : \tau \implies \Gamma \vdash e_1 \sim_{L \cup L', f_1, f_2} e_2 : \tau$

Definition 4.

$$\text{close}(R) := \{(\tau, n_1, e_1, n_2, e_2) \mid \exists m_1 \leq n_1, m_2 \leq n_2. (\tau, m_1, e_1, m_2, e_2) \in R\}$$

Definition 5.

$$n(N) := n \quad (f + f')(N) := f(N) + f'(N)$$

Lemma 8.

$$\forall \tau, k_1, k_2, G. (\tau, \tau, k_1, \bullet, k_2, \bullet) \in K(G)$$

A.2.3 Example: Eta Equivalence (lambda version).

$$\begin{aligned} e_1 &:= |\lambda f:\tau \rightarrow \tau'. \lambda x:\tau. f x| \\ e_2 &:= |\lambda f:\tau \rightarrow \tau'. f| \end{aligned}$$

In order to prove $\vdash e_1 \sim e_2 : (\tau \rightarrow \tau') \rightarrow \tau \rightarrow \tau'$ for λ_{cc}^μ (and thus λ^μ), we define a suitable local knowledge L as follows:

$$\begin{aligned} L(G) &:= \text{close}\{((\tau \rightarrow \tau') \rightarrow \tau \rightarrow \tau', 0, e_1, 0, e_2)\} \\ &\cup \text{close}\{(\tau \rightarrow \tau', 0, \lambda x. v_1 x, n_2 + 1, v_2) \mid (_, v_1, n_2, v_2) \in G(\tau \rightarrow \tau')\} \end{aligned}$$

We first show $\vdash e_1 \sim_{L,0,0} e_2 : (\tau \rightarrow \tau') \rightarrow \tau \rightarrow \tau'$:

- Suppose $G \in \text{GK}(L)$.
- We must show $(0, e_1, 0, e_2) \in \text{E}(G, G)((\tau \rightarrow \tau') \rightarrow \tau \rightarrow \tau')$.
- It suffices to show $(0, e_1, 0, e_2) \in G((\tau \rightarrow \tau') \rightarrow \tau \rightarrow \tau')$, which holds due to $G \supseteq L(G)$.

It remains to show consistent(L):

- 1) • Suppose $G \in \text{GK}(L)$, $m_1 \geq 0$, $m_2 \geq 0$, and $(_, v_1, _, v_2) \in \overline{G}(\tau \rightarrow \tau')$.
 - We must show $(m_1, e_1 v_1, m_2, e_2 v_2) \in \text{E}(G, G)^\S(\tau \rightarrow \tau')$.
 - It suffices to show $(_, \lambda x. v_1 x, _, v_2) \in \overline{G}(\tau \rightarrow \tau')$, which holds due to $G \supseteq L(G)$.
- 2) • Suppose $G \in \text{GK}(L)$, $(_, v_1, n_2, v_2) \in G(\tau \rightarrow \tau')$, $m_1 \geq 0$, $m_2 \geq 0$, and $(_, v'_1, _, v'_2) \in \overline{G}(\tau)$.
 - We must show $(m_1, (\lambda x. v_1 x) v'_1, n_2 + 1 + m_2, v_2 v'_2) \in \text{E}(G, G)^\S(\tau')$.
 - It suffices to show $(_, v_1 v'_1, n_2 + m_2, v_2 v'_2) \in \text{E}(G, G)(\tau')$.
 - With Lemma 8, it suffices to show $(_, v_1 v'_1, n_2, v_2 v'_2) \in \text{S}(G, G)(\tau')$, which follows with $(_, v_1, n_2, v_2) \in G(\tau \rightarrow \tau')$.

A.3 SPB Models for λ^μ and λ_{cc}^μ

The model for λ^μ is obtained by excluding the parts **highlighted in red**.

A.3.1 Definition.

$$\boxed{n, e \hookrightarrow n', e'}$$

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

$$\begin{aligned} \text{CVal} &:= \{v \in \text{Val} \mid \text{fv}(v) = \emptyset\} \\ \text{CExp} &:= \{e \in \text{Exp} \mid \text{fv}(e) = \emptyset\} \\ \text{CCont} &:= \{K \in \text{Cont} \mid \text{fv}(K[\langle \rangle]) = \emptyset\} \\ \text{CTy} &:= \{\tau \in \text{Ty} \mid \text{fv}(\tau) = \emptyset\} \\ \text{CTyF} &:= \{\tau_1 \rightarrow \tau_2 \mid \tau_1, \tau_2 \in \text{CTy}\} \end{aligned}$$

$$\begin{aligned} \text{DRel} &:= \mathbb{P}((\text{CTyF} \times \mathbb{N} \times \text{CVal} \times \mathbb{N} \times \text{CVal}) \uplus (\text{CTy} \times \mathbb{N} \times \text{CCont} \times \mathbb{N} \times \text{CCont})) \\ \text{VRel} &:= \mathbb{P}(\text{CTy} \times \mathbb{N} \times \text{CVal} \times \mathbb{N} \times \text{CVal}) \\ \text{ERel} &:= \mathbb{P}(\mathbb{N} \times \text{CExp} \times \mathbb{N} \times \text{CExp}) \end{aligned}$$

$$\overline{(_)} \in \text{DRel} \rightarrow \text{VRel}$$

$$\begin{aligned} \overline{R}(\text{int}) &:= \{(n_1, m, n_2, m)\} \\ \overline{R}(\tau \times \tau') &:= \{(n_1, \langle v_1, v'_1 \rangle, n_2, \langle v_2, v'_2 \rangle) \mid \exists m_1, m_2. (m_1, v_1, m_2, v_2) \in \overline{R}(\tau)\} \\ &\quad \cap \{(n_1, \langle v_1, v'_1 \rangle, n_2, \langle v_2, v'_2 \rangle) \mid \exists m'_1, m'_2. (m'_1, v'_1, m'_2, v'_2) \in \overline{R}(\tau')\} \\ \overline{R}(\tau + \tau') &:= \{(n_1, \text{inl } v_1, n_2, \text{inl } v_2) \mid \exists m_1, m_2. (m_1, v_1, m_2, v_2) \in \overline{R}(\tau)\} \\ &\quad \cup \{(n_1, \text{inr } v'_1, n_2, \text{inr } v'_2) \mid \exists m'_1, m'_2. (m'_1, v'_1, m'_2, v'_2) \in \overline{R}(\tau')\} \\ \overline{R}(\tau \rightarrow \tau') &:= \{(n_1, v_1, n_2, v_2) \in \overline{R}(\tau \rightarrow \tau')\} \\ \overline{R}(\mu\alpha. \sigma) &:= \{(n_1, \text{roll } v_1, n_2, \text{roll } v_2) \mid \exists m_1, m_2. (m_1, v_1, m_2, v_2) \in \overline{R}(\sigma[\mu\alpha. \sigma/\alpha])\} \\ \overline{R}(\text{cont } \tau) &:= \{(n_1, \text{cont } K_1, n_2, \text{cont } K_2) \mid \exists m_1, m_2. (m_1, K_1, m_2, K_2) \in \overline{R}(\tau)\} \end{aligned}$$

$$\begin{aligned} \text{LK} &:= \{L \in \text{DRel} \rightarrow \text{DRel} \mid \forall R. \forall (\tau, n_1, d_1, n_2, d_2) \in L(R). \\ &\quad \forall R' \supseteq R, n'_1 \geq n_1, n'_2 \geq n_2. (\tau, n'_1, d_1, n'_2, d_2) \in L(R')\} \end{aligned}$$

$$\begin{aligned} \text{GK}(L) &:= \{G \in \text{DRel} \mid L(G) \subseteq G \wedge (0, \bullet, 0, \bullet) \in G(\text{int}) \wedge \\ &\quad \forall (\tau, n_1, d_1, n_2, d_2) \in G. \forall n'_1 \geq n_1, n'_2 \geq n_2. (\tau, n'_1, d_1, n'_2, d_2) \in G\} \end{aligned}$$

$$\begin{aligned}
& S \in \text{DRel} \times \text{DRel} \rightarrow \text{ERel} \\
& S(R, G) := \{(k_1 + n_1, K_1[v_1], k_2 + n_2, K_2[v_2]) \mid \\
& \quad \exists \tau. (k_1, K_1, k_2, K_2) \in R(\tau) \wedge (n_1, v_1, n_2, v_2) \in \overline{G}(\tau)\} \cup \\
& \quad \{(n_1, K_1[v_1 v'_1], n_2, K_2[v_2 v'_2]) \mid \\
& \quad \exists \tau, \tau', m_1, m_2, k_1, k_2. (n_1, v_1, n_2, v_2) \in R(\tau' \rightarrow \tau) \wedge \\
& \quad (m_1, v'_1, m_2, v'_2) \in \overline{G}(\tau') \wedge (k_1, K_1, k_2, K_2) \in G(\tau)\}
\end{aligned}$$

$$\begin{aligned}
& E \in \text{DRel} \rightarrow \text{ERel} \\
& E(G) := \{(n_1, e_1, n_2, e_2) \mid (e_1 \hookrightarrow^\omega \wedge e_2 \hookrightarrow^\omega) \vee \\
& \quad \exists (n'_1, e'_1, n'_2, e'_2) \in S(G, G). n_1, e_1 \hookrightarrow^* n'_1, e'_1 \wedge n_2, e_2 \hookrightarrow^* n'_2, e'_2\}
\end{aligned}$$

$$R^S := \{(n_1, e_1, n_2, e_2) \mid n_1, n_2 > 0 \wedge (n_1 - 1, e_1, n_2 - 1, e_2) \in R\}$$

$$\text{consistent}(L) := \forall G \in \text{GK}(L). S(L(G), G) \subseteq E(G)^S$$

$$\Gamma \vdash e_1 \sim e_2 : \tau := \exists L, f_1, f_2. \text{consistent}(L) \wedge \Gamma \vdash e_1 \sim_{L, f_1, f_2} e_2 : \tau$$

$$\begin{aligned}
\Gamma \vdash e_1 \sim_{L, f_1, f_2} e_2 : \tau & := \forall G \in \text{GK}(L). \forall (N_1, \gamma_1, N_2, \gamma_2) \in \overline{G}(\Gamma). \\
& \forall (k_1, K_1, k_2, K_2) \in G(\tau). (f_1(N_1) + k_1, K_1[\gamma_1 e_1], f_2(N_2) + k_2, K_2[\gamma_2 e_2]) \in E(G)
\end{aligned}$$

$$R(\cdot) := \{(\cdot, \text{id}, \cdot, \text{id})\}$$

$$\begin{aligned}
R(x:\tau, \Gamma) & := \{(n_1 :: N_1, \gamma_1[x \mapsto v_1], n_2 :: N_2, \gamma_2[x \mapsto v_2]) \mid \\
& (n_1, v_1, n_2, v_2) \in R(\tau) \wedge (N_1, \gamma_1, N_2, \gamma_2) \in R(\Gamma)\}
\end{aligned}$$

A.3.2 Properties.

Lemma 9. If

- 1) $(n'_1, e'_1, n'_2, e'_2) \in E(G)$
- 2) $n_1, e_1 \hookrightarrow^* n'_1, e'_1$
- 3) $n_2, e_2 \hookrightarrow^* n'_2, e'_2$

then $(n_1, e_1, n_2, e_2) \in E(G)$.

Lemma 10.

- $\text{GK}(L) \cap \text{GK}(L') = \text{GK}(L \cup L')$
- $S(R, G) \cup S(R', G) = S(R \cup R', G)$
- $\text{consistent}(L) \wedge \text{consistent}(L') \implies \text{consistent}(L \cup L')$
- $\Gamma \vdash e_1 \sim_{L, f_1, f_2} e_2 : \tau \implies \Gamma \vdash e_1 \sim_{L \cup L', f_1, f_2} e_2 : \tau$

Definition 6.

$$\text{close}(R) := \{(\tau, n_1, e_1, n_2, e_2) \mid \exists m_1 \leq n_1, m_2 \leq n_2. (\tau, m_1, e_1, m_2, e_2) \in R\}$$

Definition 7.

$$n(N) := n \quad (f + f')(N) := f(N) + f'(N)$$

Lemma 11.

$$\frac{\vdash \Gamma \quad x:\tau \in \Gamma}{\Gamma \vdash x \sim x : \tau}$$

Proof:

- Let i be the index of $x:\tau$ in Γ .
- We show $\Gamma \vdash x \sim_{\emptyset, \Pi_i, \Pi_i} x : \tau$.
- So assume $G \in \text{GK}(\emptyset)$ and $(N_1, \gamma_1, N_2, \gamma_2) \in \overline{G}(\Gamma)$ and $(k_1, K_1, k_2, K_2) \in G(\tau)$.
- We must show $(\Pi_i(N_1) + k_1, K_1[\gamma_1 x], \Pi_i(N_2) + k_2, K_2[\gamma_2 x]) \in E(G)$.
- It suffices to show $(\Pi_i(N_1) + k_1, K_1[\gamma_1 x], \Pi_i(N_2) + k_2, K_2[\gamma_2 x]) \in S(G, G)$, which is obvious because $(\Pi_i(N_1), \gamma_1 x, \Pi_i(N_2), \gamma_2 x) \in \overline{G}(\tau)$.

■

Lemma 12.

$$\frac{\Gamma \vdash e_1 \sim e_2 : \tau \quad \Gamma \vdash e'_1 \sim e'_2 : \tau'}{\Gamma \vdash \langle e_1, e'_1 \rangle \sim \langle e_2, e'_2 \rangle : \tau \times \tau'}$$

Proof:

- We know (1) $\Gamma \vdash e_1 \sim_{L, f_1, f_2} e_2 : \tau$ and (2) $\Gamma \vdash e'_1 \sim_{L', f'_1, f'_2} e'_2 : \tau'$ with $\text{consistent}(L)$ and $\text{consistent}(L')$.
- We define L'' as follows:

$$\begin{aligned} L''(G) := & L(G) \cup L'(G) \\ & \cup \text{close}\{(\tau, f'_1(N_1) + 2 + k_1, K_1[\langle \bullet, \gamma_1 e'_1 \rangle], f'_2(N_2) + 2 + k_2, K_2[\langle \bullet, \gamma_2 e'_2 \rangle]) \mid \\ & (N_1, \gamma_1, N_2, \gamma_2) \in \overline{G}(\Gamma) \wedge (k_1, K_1, k_2, K_2) \in G(\tau \times \tau')\} \\ & \cup \text{close}\{(\tau', 1 + k_1, K_1[\langle v_1, \bullet \rangle], 1 + k_2, K_2[\langle v_2, \bullet \rangle]) \mid \\ & (_, v_1, _, v_2) \in \overline{G}(\tau) \wedge (k_1, K_1, k_2, K_2) \in G(\tau \times \tau')\} \end{aligned}$$

- We first show $\Gamma \vdash \langle e_1, e'_1 \rangle \sim_{L'', (f_1+f'_1+2), (f_2+f'_2+2)} \langle e_2, e'_2 \rangle : \tau \times \tau'$.
 - Assume $G \in \text{GK}(L'')$, $(N_1, \gamma_1, N_2, \gamma_2) \in \overline{G}(\Gamma)$, and $(k_1, K_1, k_2, K_2) \in G(\tau \times \tau')$.
 - Using (1) and Lemma 10, it suffices to show

$$(f'_1(N_1) + 2 + k_1, K_1[\langle \bullet, \gamma_1 e'_1 \rangle], f'_2(N_2) + 2 + k_2, K_2[\langle \bullet, \gamma_2 e'_2 \rangle]) \in G(\tau),$$

which holds due to $G \supseteq L''(G)$.

- It remains to prove $\text{consistent}(L'')$. With Lemma 10, this reduces to the following:

- 1) – Suppose $G \in \text{GK}(L'')$, $(N_1, \gamma_1, N_2, \gamma_2) \in \overline{G}(\Gamma)$, $(k_1, K_1, k_2, K_2) \in G(\tau \times \tau')$, $m_1 \geq 0$, $m_2 \geq 0$, and $(n_1, v_1, n_2, v_2) \in \overline{G}(\tau)$.
 - We must show $(f'_1(N_1) + 1 + k_1 + m_1 + n_1, K_1[\langle v_1, \gamma_1 e'_1 \rangle], f'_2(N_2) + 1 + k_2 + m_2 + n_2, K_2[\langle v_2, \gamma_2 e'_2 \rangle]) \in E(G)$.
 - Using (2) and Lemma 10, it suffices to show

$$(1 + k_1 + m_1 + n_1, K_1[\langle v_1, \bullet \rangle], 1 + k_2 + m_2 + n_2, K_2[\langle v_2, \bullet \rangle]) \in G(\tau'),$$

which holds due to $G \supseteq L''(G)$.

- 2) – Suppose $G \in \text{GK}(L'')$, $(_, v_1, _, v_2) \in \overline{G}(\tau)$, $(k_1, K_1, k_2, K_2) \in G(\tau \times \tau')$, $m_1 \geq 0$, $m_2 \geq 0$, and $(n'_1, v'_1, n'_2, v'_2) \in \overline{G}(\tau')$.
 - We must show $(k_1 + m_1 + n'_1, K_1[\langle v_1, v'_1 \rangle], k_2 + m_2 + n'_2, K_2[\langle v_2, v'_2 \rangle]) \in E(G)$.
 - This follows from $E(G) \supseteq S(G, G)$, $(k_1, K_1, k_2, K_2) \in G(\tau \times \tau')$, and $(m_1 + n'_1, \langle v_1, v'_1 \rangle, m_2 + n'_2, \langle v_2, v'_2 \rangle) \in \overline{G}(\tau \times \tau')$.

■

Lemma 13.

$$\frac{\Gamma \vdash e_1 \sim e_2 : \tau \times \tau'}{\Gamma \vdash e_1.1 \sim e_2.1 : \tau}$$

Proof:

- We know (*) $\Gamma \vdash e_1 \sim_{L, f_1, f_2} e_2 : \tau \times \tau'$ with $\text{consistent}(L)$.
- We define L' as follows:

$$\begin{aligned} L'(G) := & L(G) \\ & \cup \text{close}\{(\tau \times \tau', k_1 + 1, K_1[\langle \bullet.1 \rangle], k_2 + 1, K_2[\langle \bullet.1 \rangle]) \mid \\ & (k_1, K_1, k_2, K_2) \in G(\tau)\} \end{aligned}$$

- We first show $\Gamma \vdash e_1.1 \sim_{L', f_1+1, f_2+1} e_2.1 : \tau$.
 - Assume $G \in \text{GK}(L')$ and $(N_1, \gamma_1, N_2, \gamma_2) \in \overline{G}(\Gamma)$ and $(k_1, K_1, k_2, K_2) \in G(\tau)$.
 - Using (1) and Lemma 10, it suffices to show

$$(k_1 + 1, K_1[\langle \bullet.1 \rangle], k_2 + 1, K_2[\langle \bullet.1 \rangle]) \in G(\tau \times \tau'),$$

which holds due to $G \supseteq L'(G)$.

- It remains to show $\text{consistent}(L')$. With Lemma 10, this reduces to the following.

- Suppose $G \in \text{GK}(L')$, $m_1 \geq 0$, $m_2 \geq 0$, $(k_1, K_1, k_2, K_2) \in G(\tau)$, and (**) $(n_1, v_1, n_2, v_2) \in \overline{G}(\tau \times \tau')$.
- We must show $(k_1 + m_1 + n_1, K_1[v_1.1], k_2 + m_2 + n_2, K_2[v_2.1]) \in E(G)$.
- From (**) we know there are v'_1, v''_1, v'_2, v''_2 such that $v_i = \langle v'_i, v''_i \rangle$ and $(_, v'_1, _, v'_2) \in \overline{G}(\tau)$.
- It thus suffices to show $(_, K_1[v'_1], _, K_2[v'_2]) \in S(G, G)$, which is obvious.

■

Lemma 14.

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

Proof:

- We know (*) $\Gamma, f:(\tau' \rightarrow \tau), x:\tau' \vdash e_1 \sim_{L, f_1, f_2} e_2 : \tau$ with $\text{consistent}(L)$.
- We define L' as follows:

$$L'(G) := L(G) \cup \text{close}\{(\tau' \rightarrow \tau, 0, \text{fix } f(x). \gamma_1 e_1, 0, \text{fix } f(x). \gamma_2 e_2) \mid (_, \gamma_1, _, \gamma_2) \in \overline{G}(\Gamma)\}$$

- We first show $\Gamma \vdash \text{fix } f(x). e_1 \sim_{L', 0, 0} \text{fix } f(x). e_2 : \tau' \rightarrow \tau$.
 - Assume $G \in \text{GK}(L')$, $(N_1, \gamma_1, N_2, \gamma_2) \in \overline{G}(\Gamma)$, and $(k_1, K_1, k_2, K_2) \in G(\tau' \rightarrow \tau)$.
 - It suffices to show $(k_1, K_1[\text{fix } f(x). \gamma_1 e_1], k_2, K_2[\text{fix } f(x). \gamma_2 e_2]) \in S(G, G)$, which follows with $G \supseteq L'(G)$.
- It remains to show $\text{consistent}(L')$. With Lemma 10, this reduces to the following.
 - Suppose $G \in \text{GK}(L')$, $(_, \gamma_1, _, \gamma_2) \in \overline{G}(\Gamma)$, $m_1 \geq 0$, $m_2 \geq 0$, $(_, K_1, _, K_2) \in G(\tau)$, and $(_, v_1, _, v_2) \in \overline{G}(\tau')$.
 - We must show $(m_1, K_1[(\text{fix } f(x). \gamma_1 e_1) v_1], m_2, K_2[(\text{fix } f(x). \gamma_2 e_2) v_2]) \in E(G)$.
 - By Lemma 9 it suffices to show $(_, K_1[\gamma'_1 e_1], _, K_2[\gamma'_2 e_2]) \in E(G)$, where $\gamma'_i := \gamma_i, f \mapsto (\text{fix } f(x). \gamma_i e_i), x \mapsto v_i$.
 - Using Lemma 10, this follows from (*) if we can show

$$(_, \text{fix } f(x). \gamma_1 e_1, _, \text{fix } f(x). \gamma_2 e_2) \in \overline{G}(\tau' \rightarrow \tau)$$

and

$$(_, v_1, _, v_2) \in \overline{G}(\tau').$$

- The former follows again from $G \supseteq L'(G)$ and the latter is given. ■

Lemma 15.

$$\frac{\Gamma \vdash e_1 \sim e_2 : \tau' \rightarrow \tau \quad \Gamma \vdash e'_1 \sim e'_2 : \tau'}{\Gamma \vdash e_1 e'_1 \sim e_2 e'_2 : \tau}$$

Proof:

- We know (1) $\Gamma \vdash e_1 \sim_{L, f_1, f_2} e_2 : \tau' \rightarrow \tau$ and (2) $\Gamma \vdash e'_1 \sim_{L', f'_1, f'_2} e'_2 : \tau'$ with $\text{consistent}(L)$ and $\text{consistent}(L')$.
- We define L'' as follows:

$$L''(G) := L(G) \cup L'(G) \cup \text{close}\{(\tau' \rightarrow \tau, f'_1(N_1) + 2, K_1[\bullet \gamma_1 e'_1], f'_2(N_2) + 2, K_2[\bullet \gamma_2 e'_2]) \mid (N_1, \gamma_1, N_2, \gamma_2) \in \overline{G}(\Gamma) \wedge (_, K_1, _, K_2) \in G(\tau)\} \cup \text{close}\{(\tau', 1 + n_1, K_1[v_1 \bullet], 1 + n_2, K_2[v_2 \bullet]) \mid (_, K_1, _, K_2) \in G(\tau) \wedge (n_1, v_1, n_2, v_2) \in \overline{G}(\tau' \rightarrow \tau)\}$$

- We first show $\Gamma \vdash e_1 e'_1 \sim_{L'', (f_1 + f'_1 + 2), (f_2 + f'_2 + 2)} e_2 e'_2 : \tau$.
 - Assume $G \in \text{GK}(L'')$, $(N_1, \gamma_1, N_2, \gamma_2) \in \overline{G}(\Gamma)$, and $(k_1, K_1, k_2, K_2) \in G(\tau)$.
 - Using (1) and Lemma 10, it suffices to show $(f'_1(N_1) + 2 + k_1, K_1[\bullet \gamma_1 e'_1], f'_2(N_2) + 2 + k_2, K_2[\bullet \gamma_2 e'_2]) \in G(\tau' \rightarrow \tau)$, which holds due to $G \supseteq L''(G)$.
- It remains to show $\text{consistent}(L'')$. With Lemma 10, this reduces to the following.
 - 1) – Suppose $G \in \text{GK}(L'')$, $(N_1, \gamma_1, N_2, \gamma_2) \in \overline{G}(\Gamma)$, $(_, K_1, _, K_2) \in G(\tau)$, $m_1 \geq 0$, $m_2 \geq 0$, and $(n_1, v_1, n_2, v_2) \in \overline{G}(\tau' \rightarrow \tau)$.
 - We must show $(f'_1(N_1) + 1 + m_1 + n_1, K_1[v_1 \gamma_1 e'_1], f'_2(N_2) + 1 + m_2 + n_2, K_2[v_2 \gamma_2 e'_2]) \in E(G)$.
 - Using (2) and Lemma 10, it then suffices to show
$$(1 + m_1 + n_1, K_1[v_1 \bullet], 1 + m_2 + n_2, K_2[v_2 \bullet]) \in G(\tau'),$$
which holds due to $G \supseteq L''(G)$.
 - 2) – Suppose $G \in \text{GK}(L'')$, $(_, K_1, _, K_2) \in G(\tau)$, $(n_1, v_1, n_2, v_2) \in \overline{G}(\tau' \rightarrow \tau)$, $m_1 \geq 0$, $m_2 \geq 0$, and $(n'_1, v'_1, n'_2, v'_2) \in \overline{G}(\tau')$.
 - We must show $(n_1 + m_1 + n'_1, K_1[v_1 v'_1], n_2 + m_2 + n'_2, K_2[v_2 v'_2]) \in E(G)$.
 - It suffices to show $(n_1, K_1[v_1 v'_1], n_2, K_2[v_2 v'_2]) \in S(G, G)$, which is obvious. ■

Lemma 16.

$$\frac{\Gamma, x:\text{cont } \tau \vdash e_1 \sim e_2 : \tau}{\Gamma \vdash \text{callcc}(x. e_1) \sim \text{callcc}(x. e_2) : \tau}$$

Proof:

- We know (*) $\Gamma, x:\text{cont } \tau \vdash e_1 \sim_{L, f_1, f_2} e_2 : \tau$ with consistent(L).
- We show $\Gamma \vdash \text{callcc}(x. e_1) \sim_{L, 0, 0} \text{callcc}(x. e_2) : \tau$.
- So assume $G \in \text{GK}(L)$ and $(N_1, \gamma_1, N_2, \gamma_2) \in \overline{G}(\Gamma)$ and $(k_1, K_1, k_2, K_2) \in G(\tau)$.
- We must show $(k_1, K_1[\text{callcc}(x. \gamma_1 e_1)], k_2, K_2[\text{callcc}(x. \gamma_2 e_2)]) \in E(G)$.
- By Lemma 9 it suffices to show $(_, K_1[\gamma'_1 e_1], _, K_2[\gamma'_2 e_2]) \in E(G)$, where $\gamma'_i := \gamma_i, x \mapsto (\text{cont } K_i)$.
- This follows from (*) if we can show $(_, \text{cont } K_1, _, \text{cont } K_2) \in \overline{G}(\text{cont } \tau)$ and $(_, K_1, _, K_2) \in G(\tau)$.
- The former follows from the latter and the latter is given. ■

Lemma 17.

$$\frac{\Gamma \vdash e'_1 \sim e'_2 : \tau' \quad \Gamma \vdash e_1 \sim e_2 : \text{cont } \tau'}{\Gamma \vdash \text{throw } e'_1 \text{ to } e_1 \sim \text{throw } e'_2 \text{ to } e_2 : \tau}$$

Proof:

- We know (1) $\Gamma \vdash e'_1 \sim_{L', f'_1, f'_2} e'_2 : \tau'$ and (2) $\Gamma \vdash e_1 \sim_{L, f_1, f_2} e_2 : \text{cont } \tau'$ with consistent(L) and consistent(L').
- We define L'' as follows:

$$\begin{aligned} L''(G) := & L(G) \cup L'(G) \\ & \cup \text{close}\{(\tau', f_1(N_1) + 1, K_1[\text{throw } \bullet \text{ to } \gamma_1 e_1], f_2(N_2) + 1, K_2[\text{throw } \bullet \text{ to } \gamma_2 e_2]) \mid \\ & (N_1, \gamma_1, N_2, \gamma_2) \in \overline{G}(\Gamma) \wedge (_, K_1, _, K_2) \in G(\tau)\} \\ & \cup \text{close}\{(\text{cont } \tau', 0, K_1[\text{throw } v'_1 \text{ to } \bullet], 0, K_2[\text{throw } v'_2 \text{ to } \bullet]) \mid \\ & (_, K_1, _, K_2) \in G(\tau) \wedge (_, v'_1, _, v'_2) \in \overline{G}(\tau')\} \end{aligned}$$

- We first show $\Gamma \vdash \text{throw } e'_1 \text{ to } e_1 \sim_{L'', (f'_1+f_1+1), (f'_2+f_2+1)} \text{throw } e'_2 \text{ to } e_2 : \tau$.
 - Assume $G \in \text{GK}(L'')$ and $(N_1, \gamma_1, N_2, \gamma_2) \in \overline{G}(\Gamma)$ and $(k_1, K_1, k_2, K_2) \in G(\tau)$.
 - Using (1) and Lemma 10, it suffices to show $(f_1(N_1) + 1 + k_1, K_1[\text{throw } \bullet \text{ to } \gamma_1 e_1], f_2(N_2) + 1 + k_2, K_2[\text{throw } \bullet \text{ to } \gamma_2 e_2]) \in G(\tau')$, which holds due to $G \supseteq L''(G)$.
- It remains to show consistent(L''). With Lemma 10, this reduces to the following:
 - 1) – Suppose $G \in \text{GK}(L'')$, $(N_1, \gamma_1, N_2, \gamma_2) \in \overline{G}(\Gamma)$, $(_, K_1, _, K_2) \in G(\tau)$, $m_1 \geq 0$, $m_2 \geq 0$, and $(n'_1, v'_1, n'_2, v'_2) \in \overline{G}(\tau')$.
 - We must show $(f_1(N_1) + m_1 + n'_1, K_1[\text{throw } v'_1 \text{ to } \gamma_1 e_1], f_2(N_2) + m_2 + n'_2, K_2[\text{throw } v'_2 \text{ to } \gamma_2 e_2]) \in E(G)$.
 - Using (2) and Lemma 10, it then suffices to show
$$(m_1 + n'_1, K_1[\text{throw } v'_1 \text{ to } \bullet], m_2 + n'_2, K_2[\text{throw } v'_2 \text{ to } \bullet]) \in G(\text{cont } \tau'),$$
which holds due to $G \supseteq L''(G)$.
 - 2) – Suppose $G \in \text{GK}(L'')$, $(_, K_1, _, K_2) \in G(\tau)$, $(_, v'_1, _, v'_2) \in \overline{G}(\tau')$, $(n_1, v_1, n_2, v_2) \in \overline{G}(\text{cont } \tau')$, $m_1 + n_1 > 0$, and $m_2 + n_2 > 0$.
 - We must show $(m_1 + n_1 - 1, K_1[\text{throw } v'_1 \text{ to } v_1], m_2 + n_2 - 1, K_2[\text{throw } v'_2 \text{ to } v_2]) \in E(G)$.
 - We know $v_1 = \text{cont } K'_1$ and $v_2 = \text{cont } K'_2$ for some $(_, K'_1, _, K'_2) \in G(\tau')$.
 - It thus suffices to show $(_, K'_1[v'_1], _, K'_2[v'_2]) \in S(G, G)$, which is obvious. ■

Lemma 18.

$$\frac{\Gamma \vdash e_1 \sim e_2 : \tau \rightarrow \text{int}}{\Gamma \vdash \text{isolate } e_1 \sim \text{isolate } e_2 : \text{cont } \tau}$$

Proof:

- We know (*) $\Gamma \vdash e_1 \sim_{L, f_1, f_2} e_2 : \tau \rightarrow \text{int}$ with consistent(L).
- We define L' as follows:

$$\begin{aligned} L'(G) := & L(G) \\ & \cup \text{close}\{(\tau \rightarrow \text{int}, 0, K_1[\text{isolate } \bullet], 0, K_2[\text{isolate } \bullet]) \mid \\ & (_, K_1, _, K_2) \in G(\text{cont } \tau)\} \\ & \cup \text{close}\{(\tau, n_1 + 1, v_1 \bullet, n_2 + 1, v_2 \bullet) \mid \\ & (n_1, v_1, n_2, v_2) \in G(\tau \rightarrow \text{int})\} \end{aligned}$$

- We first show $\Gamma \vdash \text{isolate } e_1 \sim_{L', (f_1+1), (f_2+1)} \text{isolate } e_2 : \text{cont } \tau$.
 - Assume $G \in \text{GK}(L')$ and $(N_1, \gamma_1, N_2, \gamma_2) \in \overline{G}(\Gamma)$ and $(k_1, K_1, k_2, K_2) \in G(\text{cont } \tau)$.
 - We must show $(f_1(N_1) + k_1, K_1[\text{isolate } \gamma_1 e_1], f_2(N_2) + k_2, K_2[\text{isolate } \gamma_2 e_2]) \in E(G)$.

- Using (*) and Lemma 10, it suffices to show $(k_1, K_1[\text{isolate } \bullet], k_2, K_2[\text{isolate } \bullet]) \in G(\tau \rightarrow \text{int})$, which holds due to $G \supseteq L'(G)$.
- It remains to show consistent(L'). With Lemma 10, this reduces to the following:
 - 1) - Suppose $G \in \text{GK}(L')$, $(_, K_1, _, K_2) \in G(\text{cont } \tau)$, $m_1 \geq 0$, $m_2 \geq 0$, and $(n_1, v_1, n_2, v_2) \in \overline{G}(\tau \rightarrow \text{int})$.
 - We must show $(m_1 + n_1, K_1[\text{isolate } v_1], m_2 + n_2, K_2[\text{isolate } v_2]) \in E(G)$.
 - It suffices to show $(_, K_1[\text{cont } (v_1 \bullet)], _, K_2[\text{cont } (v_2 \bullet)]) \in S(G, G)$.
 - This follows from $(_, v_1 \bullet, _, v_2 \bullet) \in G(\tau)$, which holds due to $G \supseteq L'(G)$.
 - 2) - Suppose $G \in \text{GK}(L')$, $(n_1, v_1, n_2, v_2) \in G(\tau \rightarrow \text{int})$, $m_1 \geq 0$, $m_2 \geq 0$, and $(n'_1, v'_1, n'_2, v'_2) \in \overline{G}(\tau)$.
 - We must show $(n_1 + m_1 + n'_1, v_1 v'_1, n_2 + m_2 + n'_2, v_2 v'_2) \in E(G)$.
 - It suffices to show $(n_1, v_1 v'_1, n_2, v_2 v'_2) \in S(G, G)$, which follows from $(_, \bullet, _, \bullet) \in G(\text{int})$.

■

A.3.3 Example: Callcc in a Loop.

$$\begin{aligned} \tau &:= \mu\alpha. \text{cont int} \rightarrow \alpha \\ e_1 &:= |\lambda y:\tau. \text{callcc}_{\text{int}}(k. F_k y)| \quad \text{where } F_v := \text{fix } f(x). f(\text{unroll } x v) \\ e_2 &:= |\text{fix } f(x:\tau):\text{int}. \text{callcc}_{\text{int}}(k. f(\text{unroll } x k))| \end{aligned}$$

In order to prove $\vdash e_1 \sim e_2 : \tau \rightarrow \text{int}$ (in λ_{cc}^μ), we define a suitable local knowledge L as follows:

$$\begin{aligned} L(G) &:= \text{close}\{(\tau \rightarrow \text{int}, 0, e_1, 0, e_2)\} \\ &\quad \cup \text{close}\{(\tau, 0, K_1[F_{(\text{cont } K_1)} \bullet], 0, K_2[e_2 \bullet]) \mid (_, K_1, _, K_2) \in G(\text{int})\} \end{aligned}$$

We first show $\vdash e_1 \sim_{L,0,0} e_2 : \tau \rightarrow \text{int}$:

- Suppose $G \in \text{GK}(L)$ and $(k_1, K_1, k_2, K_2) \in G(\tau \rightarrow \text{int})$.
- We must show $(k_1, K_1[e_1], k_2, K_2[e_2]) \in E(G)$.
- It suffices to show $(0, e_1, 0, e_2) \in G(\tau \rightarrow \text{int})$, which holds due to $G \supseteq L(G)$.

It remains to show consistent(L):

- 1) • Suppose $G \in \text{GK}(L)$, $m_1 \geq 0$, $m_2 \geq 0$, (*) $(_, v_1, _, v_2) \in \overline{G}(\tau)$, and $(_, K_1, _, K_2) \in G(\text{int})$.
 - We must show $(m_1, K_1[e_1 v_1], m_2, K_2[e_2 v_2]) \in E(G)^\S$.
 - From (*) we know that there is $(_, v'_1, _, v'_2) \in \overline{G}(\text{cont int} \rightarrow \tau)$ such that $v_i = \text{roll } v'_i$.
 - It thus suffices to show $(_, K_1[F_{(\text{cont } K_1)}(v'_1(\text{cont } K_1))], _, K_2[e_2(v'_2(\text{cont } K_2))]) \in S(G, G)$.
 - Since $(_, \text{cont } K_1, _, \text{cont } K_2) \in \overline{G}(\text{cont int})$, this follows from $(_, K_1[F_{(\text{cont } K_1)} \bullet], _, K_2[e_2 \bullet]) \in G(\tau)$, which holds due to $G \supseteq L(G)$.
- 2) • Suppose $G \in \text{GK}(L)$, $(_, K_1, _, K_2) \in G(\text{int})$, $m_1 \geq 0$, $m_2 \geq 0$, and (*) $(n_1, v_1, n_2, v_2) \in \overline{G}(\tau)$.
 - We must show $(m_1 + n_1, K_1[F_{(\text{cont } K_1)} v_1], m_2 + n_2, K_2[e_2 v_2]) \in E(G)^\S$.
 - From (*) we know that there is $(_, v'_1, _, v'_2) \in \overline{G}(\text{cont int} \rightarrow \tau)$ such that $v_i = \text{roll } v'_i$.
 - It thus suffices to show $(_, K_1[F_{(\text{cont } K_1)}(v'_1(\text{cont } K_1))], _, K_2[e_2(v'_2(\text{cont } K_2))]) \in S(G, G)$.
 - Since $(_, \text{cont } K_1, _, \text{cont } K_2) \in \overline{G}(\text{cont int})$, this follows from $(_, K_1[F_{(\text{cont } K_1)} \bullet], _, K_2[e_2 \bullet]) \in G(\tau)$, which holds due to $G \supseteq L(G)$.

A.3.4 Example: Eta Equivalence (lambda version).

$$\begin{aligned} e_1 &:= |\lambda f:\tau \rightarrow \tau'. \lambda x:\tau. f x| \\ e_2 &:= |\lambda f:\tau \rightarrow \tau'. f| \end{aligned}$$

In order to prove $\vdash e_1 \sim e_2 : (\tau \rightarrow \tau') \rightarrow \tau \rightarrow \tau'$ for λ_{cc}^μ (and thus λ^μ), we define a suitable local knowledge L as follows:

$$\begin{aligned} L(G) &:= \text{close}\{((\tau \rightarrow \tau') \rightarrow \tau \rightarrow \tau', 0, e_1, 0, e_2)\} \\ &\quad \cup \text{close}\{(\tau \rightarrow \tau', 0, \lambda x. v_1 x, n_2 + 1, v_2) \mid (_, v_1, n_2, v_2) \in G(\tau \rightarrow \tau')\} \end{aligned}$$

We first show $\vdash e_1 \sim_{L,0,0} e_2 : (\tau \rightarrow \tau') \rightarrow \tau \rightarrow \tau'$:

- Suppose $G \in \text{GK}(L)$ and $(k_1, K_1, k_2, K_2) \in G((\tau \rightarrow \tau') \rightarrow \tau \rightarrow \tau')$.
- We must show $(k_1, K_1[e_1], k_2, K_2[e_2]) \in E(G)$.
- It suffices to show $(0, e_1, 0, e_2) \in G((\tau \rightarrow \tau') \rightarrow \tau \rightarrow \tau')$, which holds due to $G \supseteq L(G)$.

It remains to show consistent(L):

- 1) • Suppose $G \in \text{GK}(L)$, $m_1 \geq 0$, $m_2 \geq 0$, $(_, v_1, _, v_2) \in \overline{G}(\tau \rightarrow \tau')$, and $(_, K_1, _, K_2) \in G(\tau \rightarrow \tau')$.
 - We must show $(m_1, K_1[e_1 v_1], m_2, K_2[e_2 v_2]) \in E(G)^\S$.

- It suffices to show $(_, K_1[\lambda x. v_1 x], _, K_2[v_2]) \in S(G, G)$.
 - This follows from $(_, \lambda x. v_1 x, _, v_2) \in \overline{G}(\tau \rightarrow \tau')$, which holds due to $G \supseteq L(G)$.
- 2) Suppose $G \in \text{GK}(L)$, $(_, v_1, n_2, v_2) \in G(\tau \rightarrow \tau')$, $m_1 \geq 0$, $m_2 \geq 0$, $(_, v'_1, _, v'_2) \in \overline{G}(\tau)$ and $(_, K_1, _, K_2) \in G(\tau')$.
- We must show $(m_1, K_1[(\lambda x. v_1 x) v'_1], n_2 + 1 + m_2, K_2[v_2 v'_2]) \in E(G)^\S$.
 - It suffices to show $(_, K_1[v_1 v'_1], n_2, K_2[v_2 v'_2]) \in S(G, G)$, which follows with $(_, v_1, n_2, v_2) \in G(\tau \rightarrow \tau')$.

A.3.5 Example: Syntactic Minimal Invariance.

$$\begin{aligned} \tau &:= \mu\alpha. \text{int} + (\alpha \rightarrow \alpha) \\ e_1 &:= |\text{fix } f(x:\tau):\tau. \text{case unroll } x \text{ of inl } y \Rightarrow \text{roll } (\text{inl } y) \\ &\quad | \text{inr } g \Rightarrow \text{roll } (\text{inr } \lambda y:\tau. f(g(f y)))| \\ e_2 &:= |\lambda x:\tau. x| \end{aligned}$$

In order to prove $\vdash e_1 \sim e_2 : \tau \rightarrow \tau$ for λ_{cc}^μ (and thus λ^μ), we define a suitable local knowledge L as follows:

$$\begin{aligned} L(G) &:= \text{close}\{(\tau \rightarrow \tau, 0, e_1, 0, e_2)\} \\ &\quad \cup \text{close}\{(\tau \rightarrow \tau, 0, \lambda y. e_1(v_1(e_1 y)), n_2 + 1, v_2) \mid (_, v_1, n_2, v_2) \in G(\tau \rightarrow \tau)\} \\ &\quad \cup \text{close}\{(\tau, 0, K_1[e_1 \bullet], k_2 + 1, K_2) \mid (_, K_1, k_2, K_2) \in G(\tau)\} \end{aligned}$$

We first show $\vdash e_1 \sim_{L,0,0} e_2 : \tau \rightarrow \tau$:

- Suppose $G \in \text{GK}(L)$ and $(k_1, K_1, k_2, K_2) \in G(\tau \rightarrow \tau)$.
- We must show $(k_1, K_1[e_1], k_2, K_2[e_2]) \in E(G)$.
- It suffices to show $(0, e_1, 0, e_2) \in G(\tau \rightarrow \tau)$, which holds due to $G \supseteq L(G)$.

To simplify the proof of consistency, we now show the following property (\dagger):

$$\forall G \in \text{GK}(L), (_, v_1, n_2, v_2) \in \overline{G}(\tau). \forall K. \exists v'_1. K[e_1 v_1] \hookrightarrow^+ K[v'_1] \wedge (_, v'_1, n_2, v_2) \in \overline{G}(\tau)$$

- From $(_, v_1, n_2, v_2) \in \overline{G}(\tau)$ we know that there are v'_1, v'_2 such that $v_i = \text{roll } v'_i$ and $(_, v'_1, _, v'_2) \in \overline{G}(\text{int} + (\tau \rightarrow \tau))$.
- From $(_, v'_1, _, v'_2) \in \overline{G}(\text{int} + (\tau \rightarrow \tau))$ we know that there are v''_1, v''_2 such that either (a) $v'_i = \text{inl } v''_i$ and $(_, v''_1, _, v''_2) \in \overline{G}(\text{int})$ or (b) $v'_i = \text{inr } v''_i$ and $(_, v''_1, _, v''_2) \in \overline{G}(\tau \rightarrow \tau)$.
- In case (a) we know $K[e_1 v_1] \hookrightarrow^+ K[v_1]$, and we are done.
- Now consider case (b).
- Here we know $K[e_1 v_1] \hookrightarrow^+ K[\widehat{v}_1]$ for $\widehat{v}_1 := \text{roll } (\text{inr } \lambda y. e_1(v''_1(e_1 y)))$, so it remains to show $(_, \widehat{v}_1, n_2, \text{roll } (\text{inr } v''_2)) \in \overline{G}(\tau)$.
- This follows from $(_, \lambda y. e_1(v''_1(e_1 y)), _, v''_2) \in \overline{G}(\tau \rightarrow \tau)$, which holds due to $G \supseteq L(G)$.

It remains to show consistent(L):

- 1) • Suppose $G \in \text{GK}(L)$, $m_1 \geq 0$, $m_2 \geq 0$, $(_, v_1, _, v_2) \in \overline{G}(\tau)$, and $(_, K_1, _, K_2) \in G(\tau)$.
 - We must show $(m_1, K_1[e_1 v_1], m_2, K_2[e_2 v_2]) \in E(G)^\S$.
 - With the help of (\dagger), it suffices to show $(_, K_1[v'_1], _, K_2[v_2]) \in S(G, G)$ for any v'_1 with $(_, v'_1, _, v_2) \in \overline{G}(\tau)$, which is obvious.
- 2) • Suppose $G \in \text{GK}(L)$, $(_, v_1, n_2, v_2) \in G(\tau \rightarrow \tau)$, $m_1 \geq 0$, $m_2 \geq 0$, $(_, v'_1, _, v'_2) \in \overline{G}(\tau)$, and $(_, K_1, _, K_2) \in G(\tau)$.
 - We must show $(m_1, K_1[(\lambda y. e_1(v_1(e_1 y))) v'_1], n_2 + 1 + m_2, K_2[v_2 v'_2]) \in E(G)^\S$.
 - With the help of (\dagger), it suffices to show $(_, K_1[e_1(v_1 v'_1)], n_2, K_2[v_2 v'_2]) \in S(G, G)$ for any v'_1 with $(_, v'_1, _, v_2) \in \overline{G}(\tau)$.
 - This follows from $(_, K_1[e_1 \bullet], _, K_2) \in G(\tau)$, which holds due to $G \supseteq L(G)$.
- 3) • Suppose $G \in \text{GK}(L)$, $(_, K_1, k_2, K_2) \in G(\tau)$, $m_1 \geq 0$, $m_2 \geq 0$, and $(n_1, v_1, n_2, v_2) \in \overline{G}(\tau)$.
 - We must show $(m_1 + n_1, K_1[e_1 v_1], k_2 + 1 + m_2 + n_2, K_2[v_2]) \in E(G)^\S$.
 - With the help of (\dagger), it suffices to show $(_, K_1[v'_1], k_2 + n_2, K_2[v_2]) \in S(G, G)$ for any v'_1 with $(_, v'_1, n_2, v_2) \in \overline{G}(\tau)$, which is obvious.

A.3.6 Using Paco (parameterized coinduction).

Definition 8.

$$\begin{aligned} \mathfrak{f} &\in \text{LK} \xrightarrow{\text{mon}} \text{LK} \\ \mathfrak{f}(L)(G)(\tau) &:= \{(n_1, d_1, n_2, d_2) \mid \\ &\quad \forall G' \in \text{GK}(L). G' \supseteq G \implies S(\{(\tau, n_1, d_1, n_2, d_2)\}, G') \subseteq E(G')^\S\} \end{aligned}$$

Using Lemmas 9 and 10, it is easy to verify that \mathfrak{f} is well-defined, *i.e.*, that it returns valid local knowledges and is itself monotone.

Lemma 19.

$$L \subseteq \mathfrak{f}(L) \implies \text{consistent}(L)$$

Proof:

- Assume $L \subseteq \mathfrak{f}(L)$ and suppose $G \in \text{GK}(L)$ and $(n_1, e_1, n_2, e_2) \in \text{S}(L(G), G)$.
- We must show $(n_1, e_1, n_2, e_2) \in \text{E}(G)^\S$.
- From $(n_1, e_1, n_2, e_2) \in \text{S}(L(G), G)$ we know that there is $(\tau, m_1, d_1, m_2, d_2) \in L(G)$ such that $(n_1, e_1, n_2, e_2) \in \text{S}(\{(\tau, m_1, d_1, m_2, d_2)\}, G)$.
- Using the assumption, we have $(\tau, m_1, d_1, m_2, d_2) \in \mathfrak{f}(L)(G)$.
- Exploiting this, we are done because $G \supseteq G$.

■

Lemma 20.

$$\text{consistent}(L) \implies L \subseteq \mathfrak{f}(L)$$

Proof:

- Assume $\text{consistent}(L)$ and suppose $G \in \text{DRel}$ and $(\tau, n_1, d_1, n_2, d_2) \in L(G)$.
- We must show $(\tau, n_1, d_1, n_2, d_2) \in \mathfrak{f}(L)(G)$.
- So suppose $G' \in \text{GK}(L)$, $G' \supseteq G$, and $(m_1, e_1, m_2, e_2) \in \text{S}(\{(\tau, n_1, d_1, n_2, d_2)\}, G')$.
- Using monotonicity of $\text{S}(-, G')$ and of L , this implies $(m_1, e_1, m_2, e_2) \in \text{S}(L(G'), G')$.
- From $\text{consistent}(L)$ we then get $(m_1, e_1, m_2, e_2) \in \text{E}(G')^\S$.

■

Definition 9.

$$\begin{aligned} \mathfrak{L} &\in \text{LK} \rightarrow \text{LK} \\ \mathfrak{L}(L) &:= \nu X. \mathfrak{f}(X \cup L) \end{aligned}$$

Lemma 21 ($\mathfrak{L}(\emptyset)$ is the greatest consistent local knowledge).

- $\text{consistent}(\mathfrak{L}(\emptyset))$
- $\text{consistent}(L) \implies L \subseteq \mathfrak{L}(\emptyset)$

Proof: Since $\mathfrak{L}(\emptyset)$ is by definition the greatest postfix point of \mathfrak{f} , this follows from Lemmas 19 and 20.

■

Lemma 22.

$$\Gamma \vdash e_1 \sim e_2 : \tau \iff \exists f_1, f_2. \Gamma \vdash e_1 \sim_{\mathfrak{L}(\emptyset), f_1, f_2} e_2 : \tau$$

Proof: The \Leftarrow -direction is trivial with Lemma 21. Consider the \Rightarrow -direction:

- From $\Gamma \vdash e_1 \sim e_2 : \tau$ we know $\Gamma \vdash e_1 \sim_{L, f_1, f_2} e_2 : \tau$ with $\text{consistent}(L)$.
- The latter implies $L \subseteq \mathfrak{L}(\emptyset)$ by Lemma 21.
- We are done by Lemma 10.

■

Lemma 23 (ACCUMULATE).

$$L \subseteq \mathfrak{L}(L') \iff L \subseteq \mathfrak{L}(L \cup L')$$

Proof: See our Paco paper (The Power of Parameterization in Coinductive Proof).

■

All this can be done for the other models/languages too. See our Coq formalization.

B.1 Languages $F^{\mu!}$ and $F_{cc}^{\mu!}$

The fragment $F^{\mu!}$ is obtained by removing first-class continuations (the parts **highlighted in red**).

B.1.1 Statics.

$\sigma, \tau \in \text{Ty} ::= \alpha \mid \text{int} \mid \text{unit} \mid \sigma_1 \times \sigma_2 \mid \sigma_1 + \sigma_2 \mid \sigma_1 \rightarrow \sigma_2 \mid \mu\alpha. \sigma \mid \text{cont } \sigma \mid \forall\alpha. \sigma \mid \exists\alpha. \sigma \mid \text{ref } \sigma$
 $p \in \text{Prog} ::= x \mid n \mid \langle \rangle \mid p_1 \odot p_2 \mid \text{ifz } p \text{ then } p_1 \text{ else } p_2 \mid \langle p_1, p_2 \rangle \mid p.1 \mid p.2 \mid \text{inl}_\sigma p \mid \text{inr}_\sigma p \mid$
 $(\text{case } p \text{ of } \text{inl } x \Rightarrow p_1 \mid \text{inr } x \Rightarrow p_2) \mid \text{fix } f(x:\sigma_1):\sigma_2. p \mid p_1 p_2 \mid \text{roll}_\sigma p \mid \text{unroll } p \mid$
 $\text{callcc}_\sigma(x. p) \mid \text{throw}_\sigma p_1 \text{ to } p_2 \mid \text{isolate } p \mid \Lambda\alpha. p \mid p[\sigma] \mid \text{pack } \langle \sigma_1, p \rangle \text{ as } \sigma_2 \mid$
 $\text{unpack } p_1 \text{ as } \langle \alpha, x \rangle \text{ in } p_2 \mid \text{ref } p \mid p_1 := p_2 \mid p_1 == p_2 \mid !p$

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

$\Delta \vdash \sigma$

$$\frac{\text{fv}(\sigma) \subseteq \Delta \quad \text{names}(\sigma) = \emptyset}{\Delta \vdash \sigma}$$

$\Delta \vdash \Gamma$

$$\frac{\forall x:\sigma \in \Gamma. \Delta \vdash \sigma}{\Delta \vdash \Gamma}$$

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

$$\frac{\vdash \Delta; \Gamma \quad x:\tau \in \Delta; \Gamma}{\Delta; \Gamma \vdash x : \tau} \quad \frac{\vdash \Delta; \Gamma}{\Delta; \Gamma \vdash n : \text{int}}$$

$$\frac{\Delta; \Gamma \vdash p_1 : \text{int} \quad \Delta; \Gamma \vdash p_2 : \text{int}}{\Delta; \Gamma \vdash p_1 \odot p_2 : \text{int}} \quad \frac{\Delta; \Gamma \vdash p_0 : \text{int} \quad \Delta; \Gamma \vdash p_1 : \tau \quad \Delta; \Gamma \vdash p_2 : \tau}{\Delta; \Gamma \vdash \text{ifz } p_0 \text{ then } p_1 \text{ else } p_2 : \tau}$$

$$\frac{\Delta; \Gamma \vdash p_1 : \tau_1 \quad \Delta; \Gamma \vdash p_2 : \tau_2}{\Delta; \Gamma \vdash \langle p_1, p_2 \rangle : \tau_1 \times \tau_2} \quad \frac{\Delta; \Gamma \vdash p : \tau_1 \times \tau_2}{\Delta; \Gamma \vdash p.1 : \tau_1} \quad \frac{\Delta; \Gamma \vdash p : \tau_1 \times \tau_2}{\Delta; \Gamma \vdash p.2 : \tau_2}$$

$$\frac{\Delta; \Gamma \vdash p : \tau_1}{\Delta; \Gamma \vdash \text{inl}_{\tau_2} p : \tau_1 + \tau_2} \quad \frac{\Delta; \Gamma \vdash p : \tau_2}{\Delta; \Gamma \vdash \text{inr}_{\tau_1} p : \tau_1 + \tau_2}$$

$$\frac{\Delta; \Gamma \vdash p : \tau_1 + \tau_2 \quad \Delta; \Gamma, x:\tau_1 \vdash p_1 : \tau \quad \Delta; \Gamma, x:\tau_2 \vdash p_2 : \tau}{\Delta; \Gamma \vdash \text{case } p \text{ of } \text{inl } x \Rightarrow p_1 \mid \text{inr } x \Rightarrow p_2 : \tau}$$

$$\frac{\Delta; \Gamma, f:(\tau_1 \rightarrow \tau_2), x:\tau_1 \vdash p : \tau_2}{\Delta; \Gamma \vdash \text{fix } f(x:\tau_1):\tau_2. p : \tau_1 \rightarrow \tau_2} \quad \frac{\Delta; \Gamma \vdash p_1 : \tau_1 \rightarrow \tau_2 \quad \Delta; \Gamma \vdash p_2 : \tau_1}{\Delta; \Gamma \vdash p_1 p_2 : \tau_2}$$

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

$$\frac{\Delta; \Gamma, x:\text{cont } \tau \vdash p : \tau}{\Delta; \Gamma \vdash \text{callcc}_\tau(x. p) : \tau} \quad \frac{\Delta; \Gamma \vdash p' : \tau' \quad \Delta; \Gamma \vdash p : \text{cont } \tau'}{\Delta; \Gamma \vdash \text{throw}_\tau p' \text{ to } p : \tau} \quad \frac{\Delta; \Gamma \vdash p : \tau \rightarrow \text{int}}{\Delta; \Gamma \vdash \text{isolate } p : \text{cont } \tau}$$

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

$$\begin{array}{c}
\frac{\Delta \vdash \sigma_1 \quad \Delta; \Gamma \vdash p : \sigma_2[\sigma_1/\alpha]}{\Delta; \Gamma \vdash \text{pack } \langle \sigma_1, p \rangle \text{ as } \exists \alpha. \sigma_2 : \exists \alpha. \sigma_2} \quad \frac{\Delta; \Gamma \vdash p_1 : \exists \alpha. \sigma_1 \quad \Delta, \alpha; \Gamma, x:\sigma_1 \vdash p_2 : \sigma_2 \quad \Delta \vdash \sigma_2}{\Delta; \Gamma \vdash \text{unpack } p_1 \text{ as } \langle \alpha, x \rangle \text{ in } p_2 : \sigma_2} \\
\\
\frac{\Delta; \Gamma \vdash p : \sigma}{\Delta; \Gamma \vdash \text{ref } p : \text{ref } \sigma} \quad \frac{\Delta; \Gamma \vdash p_1 : \text{ref } \sigma \quad \Delta; \Gamma \vdash p_2 : \sigma}{\Delta; \Gamma \vdash p_1 := p_2 : \text{unit}} \\
\\
\frac{\Delta; \Gamma \vdash p : \text{ref } \sigma}{\Delta; \Gamma \vdash !p : \sigma} \quad \frac{\Delta; \Gamma \vdash p_1 : \text{ref } \sigma \quad \Delta; \Gamma \vdash p_2 : \text{ref } \sigma}{\Delta; \Gamma \vdash p_1 == p_2 : \text{int}}
\end{array}$$

B.1.2 Dynamics.

$$\begin{array}{l}
v \in \text{Val} \quad ::= \quad x \mid n \mid \langle \rangle \mid l \mid \langle v_1, v_2 \rangle \mid \text{inl } v \mid \text{inr } v \mid \text{fix } f(x).e \mid \text{roll } v \mid \text{cont } K \mid \Lambda.e \mid \text{pack } e \\
e \in \text{Exp} \quad ::= \quad v \mid e_1 \odot e_2 \mid \text{ifz } e_0 \text{ then } e_1 \text{ else } e_2 \mid \langle e_1, e_2 \rangle \mid e.1 \mid e.2 \mid \text{inl } e \mid \text{inr } e \mid \\
\quad (\text{case } e \text{ of inl } x \Rightarrow e_1 \mid \text{inr } x \Rightarrow e_2) \mid e_1 e_2 \mid \text{roll } e \mid \text{unroll } e \mid \\
\quad \text{callcc } (x.e) \mid \text{throw } e_1 \text{ to } e_2 \mid \text{isolate } e \mid e[] \mid \text{unpack } e_1 \text{ as } x \text{ in } e_2 \mid \\
\quad \text{ref } e \mid e_1 := e_2 \mid e_1 == e_2 \mid !e \\
K \in \text{Cont} \quad ::= \quad \bullet \mid K \odot e \mid v \odot K \mid \text{ifz } K \text{ then } e_1 \text{ else } e_2 \mid \langle K, e \rangle \mid \langle v, K \rangle \mid K.1 \mid K.2 \mid \\
\quad \text{inl } K \mid \text{inr } K \mid (\text{case } K \text{ of inl } x \Rightarrow e_1 \mid \text{inr } x \Rightarrow e_2) \mid K e \mid v K \mid \text{roll } K \mid \\
\quad \text{unroll } K \mid \text{throw } K \text{ to } e \mid \text{throw } v \text{ to } K \mid \text{isolate } K \mid K[] \mid \text{pack } K \mid \\
\quad \text{unpack } K \text{ as } x \text{ in } e \mid \text{ref } K \mid K := e \mid v := K \mid K == e \mid v == K \mid !K
\end{array}$$

$$h, e \hookrightarrow h', e'$$

$$\begin{array}{l}
h, K[n_1 \odot n_2] \quad \hookrightarrow \quad h, K[n] \quad (n = \llbracket n_1 \odot n_2 \rrbracket) \\
h, K[\text{ifz } 0 \text{ then } e_1 \text{ else } e_2] \quad \hookrightarrow \quad h, K[e_1] \\
h, K[\text{ifz } n \text{ then } e_1 \text{ else } e_2] \quad \hookrightarrow \quad h, K[e_2] \quad (n \neq 0) \\
h, K[\langle v_1, v_2 \rangle.1] \quad \hookrightarrow \quad h, K[v_1] \\
h, K[\langle v_1, v_2 \rangle.2] \quad \hookrightarrow \quad h, K[v_2] \\
h, K[\text{case inl } v \text{ of inl } x \Rightarrow e_1 \mid \text{inr } x \Rightarrow e_2] \quad \hookrightarrow \quad h, K[e_1[v/x]] \\
h, K[\text{case inr } v \text{ of inl } x \Rightarrow e_1 \mid \text{inr } x \Rightarrow e_2] \quad \hookrightarrow \quad h, K[e_2[v/x]] \\
h, K[(\text{fix } f(x).e) v] \quad \hookrightarrow \quad h, K[e[(\text{fix } f(x).e)/f, v/x]] \\
h, K[\text{unroll } (\text{roll } v)] \quad \hookrightarrow \quad h, K[v] \\
h, K[\text{callcc } (x.e)] \quad \hookrightarrow \quad h, K[e[\text{cont } K/x]] \\
h, K[\text{throw } v \text{ to cont } K'] \quad \hookrightarrow \quad h, K'[v] \\
h, K[\text{isolate } v] \quad \hookrightarrow \quad h, K[\text{cont } (v \bullet)] \\
h, K[(\Lambda.e)[]] \quad \hookrightarrow \quad h, K[e] \\
h, K[\text{unpack } (\text{pack } v) \text{ as } x \text{ in } e] \quad \hookrightarrow \quad h, K[e[v/x]] \\
h, K[\text{ref } v] \quad \hookrightarrow \quad h \uplus [\ell \rightarrow v], K[\ell] \\
h \uplus [\ell \rightarrow v], K[! \ell] \quad \hookrightarrow \quad h \uplus [\ell \rightarrow v], K[v] \\
h \uplus [\ell \rightarrow v], K[\ell := v'] \quad \hookrightarrow \quad h \uplus [\ell \rightarrow v'], K[\langle \rangle] \\
h, K[\ell == \ell] \quad \hookrightarrow \quad h, K[1] \\
h, K[\ell == \ell'] \quad \hookrightarrow \quad h, K[0] \quad (\ell \neq \ell')
\end{array}$$

B.2 SPB Models for $F^{\mu!}$ and $F_{cc}^{\mu!}$

The model for $F^{\mu!}$ is obtained by *excluding* the parts **highlighted in red** and *including* the parts **highlighted in blue**. Conversely, the model for $F_{cc}^{\mu!}$ is obtained by *excluding* the parts **highlighted in blue** and *including* the parts **highlighted in red**. Accordingly, the model for $F_{cc}^{\mu!}$ allows only one transition relation (\sqsubseteq_{pub}) in a world. This is because, in the presence of callcc, every transition is observable and thus public [5].

B.2.1 Definition.

$$\boxed{h, n, e \hookrightarrow h', n', e'}$$

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

For the construction of the model, we extend the syntax of types with type names \mathbf{n} . They are used in modelling existential and universal types.

$$\begin{aligned} \text{CVal} &:= \{v \in \text{Val} \mid \text{fv}(v) = \emptyset\} \\ \text{CExp} &:= \{e \in \text{Exp} \mid \text{fv}(e) = \emptyset\} \\ \text{CCont} &:= \{K \in \text{Cont} \mid \text{fv}(K[\langle \rangle]) = \emptyset\} \\ \text{CTy} &:= \{\tau \in \text{Ty} \mid \text{fv}(\tau) = \emptyset\} \\ \text{CTyF} &:= \{\tau_1 \rightarrow \tau_2 \mid \tau_1, \tau_2 \in \text{CTy}\} \end{aligned}$$

$$\begin{aligned} \text{DRel} &:= \mathbb{P}((\text{CTyF} \times \mathbb{N} \times \text{CVal} \times \mathbb{N} \times \text{CVal}) \uplus (\text{CTy} \times \mathbb{N} \times \text{CCont} \times \mathbb{N} \times \text{CCont})) \\ \text{VRel} &:= \mathbb{P}(\text{CTyF} \times \mathbb{N} \times \text{CVal} \times \mathbb{N} \times \text{CVal}) \\ \text{ERel} &:= \mathbb{P}(\mathbb{N} \times \text{CExp} \times \mathbb{N} \times \text{CExp}) \\ \text{HRel} &:= \mathbb{P}(\text{Heap} \times \text{Heap}) \end{aligned}$$

$$\begin{aligned} \overline{(-)} \in \text{DRel} &\rightarrow \text{VRel} \\ \overline{R(\mathbf{n})} &:= \{(n_1, v_1, n_2, v_2) \in R(\mathbf{n})\} \\ \overline{R(\text{int})} &:= \{(n_1, m, n_2, m)\} \\ \overline{R(\tau \times \tau')} &:= \{(n_1, \langle v_1, v'_1 \rangle, n_2, \langle v_2, v'_2 \rangle) \mid \exists m_1, m_2. (m_1, v_1, m_2, v_2) \in \overline{R(\tau)}\} \\ &\quad \cap \{(n_1, \langle v_1, v'_1 \rangle, n_2, \langle v_2, v'_2 \rangle) \mid \exists m'_1, m'_2. (m'_1, v'_1, m'_2, v'_2) \in \overline{R(\tau')}\} \\ \overline{R(\tau + \tau')} &:= \{(n_1, \text{inl } v_1, n_2, \text{inl } v_2) \mid \exists m_1, m_2. (m_1, v_1, m_2, v_2) \in \overline{R(\tau)}\} \\ &\quad \cup \{(n_1, \text{inr } v'_1, n_2, \text{inr } v'_2) \mid \exists m'_1, m'_2. (m'_1, v'_1, m'_2, v'_2) \in \overline{R(\tau')}\} \\ \overline{R(\tau \rightarrow \tau')} &:= \{(n_1, v_1, n_2, v_2) \in R(\tau \rightarrow \tau')\} \\ \overline{R(\exists \alpha. \sigma)} &:= \{(n_1, \text{pack } v_1, n_2, \text{pack } v_2) \mid \exists m_1, m_2, \tau. (m_1, v_1, m_2, v_2) \in \overline{R(\sigma[\tau/\alpha])}\} \\ \overline{R(\forall \alpha. \sigma)} &:= \{(n_1, v_1, n_2, v_2) \in R(\forall \alpha. \sigma)\} \\ \overline{R(\mu \alpha. \sigma)} &:= \{(n_1, \text{roll } v_1, n_2, \text{roll } v_2) \mid \exists m_1, m_2. (m_1, v_1, m_2, v_2) \in \overline{R(\sigma[\mu \alpha. \sigma/\alpha])}\} \\ \overline{R(\text{ref } \tau)} &:= \{(n_1, v_1, n_2, v_2) \in R(\text{ref } \tau)\} \\ \overline{R(\text{cont } \tau)} &:= \{(n_1, \text{cont } K_1, n_2, \text{cont } K_2) \mid \exists m_1, m_2. (m_1, K_1, m_2, K_2) \in R(\tau)\} \end{aligned}$$

$\text{DepWorld}(P) := \{(\mathbb{S} \in \text{Set},$
 $\quad \sqsubseteq \in \mathbb{P}(\mathbb{S} \times \mathbb{S}),$
 $\quad \sqsubseteq_{\text{pub}} \in \mathbb{P}(\mathbb{S} \times \mathbb{S}),$
 $\quad \mathbb{N} \in \mathbb{P}(\text{TyNam}),$
 $\quad \mathbb{L} \in (\mathbb{S}_P \rightarrow \mathbb{S} \rightarrow \text{DRel}) \rightarrow \mathbb{S}_P \rightarrow \mathbb{S} \rightarrow \text{DRel},$
 $\quad \mathbb{H} \in (\mathbb{S}_P \rightarrow \mathbb{S} \rightarrow \text{DRel}) \rightarrow \mathbb{S}_P \rightarrow \mathbb{S} \rightarrow \text{HRel}) \mid$
 $\text{countable}(\mathbb{S}) \wedge \text{preorder}(\mathbb{S}, \sqsubseteq) \wedge \text{preorder}(\mathbb{S}, \sqsubseteq_{\text{pub}}) \wedge$
 $(\forall R, s_P, s. \forall (\tau, n_1, d_1, n_2, d_2) \in \mathbb{L}(R)(s_P)(s).$
 $\quad \forall R' \supseteq R, s'_P \supseteq_P s_P, s' \supseteq_{\text{pub}} s, n'_1 \geq n_1, n'_2 \geq n_2.$
 $\quad (\tau, n'_1, d_1, n'_2, d_2) \in \mathbb{L}(R')(s'_P)(s')) \wedge$
 $(\forall R, s_P, s. \forall R' \supseteq R \cap (\mathbb{S}_P \rightarrow \mathbb{S} \rightarrow \text{VRel}), s' \supseteq s.$
 $\quad \mathbb{L}(R')(s_P)(s') \supseteq \mathbb{L}(R)(s_P)(s) \cap \text{VRel}) \wedge$
 $(\forall R. \forall R' \supseteq R \cap (\mathbb{S}_P \rightarrow \mathbb{S} \rightarrow \text{VRel}).$
 $\quad \mathbb{H}(R') \supseteq \mathbb{H}(R) \cap (\mathbb{S}_P \rightarrow \mathbb{S} \rightarrow \text{VRel})) \wedge$
 $(\forall R, s_P, s. \forall (\mathbf{n}, n_1, v_1, n_2, v_2) \in \mathbb{L}(R)(s_P)(s). \mathbf{n} \in \mathbb{N})\}$
 $\text{where } P = (\mathbb{S}_P \in \text{Set}, \sqsubseteq_P: \mathbb{P}(\mathbb{S}_P \times \mathbb{S}_P))$

$\text{World} := \{W \in \text{DepWorld}(\{*\}, \{(*, *)\})\}$

$W_{\text{ref}}.\mathbb{S} := \{s_{\text{ref}} \in \mathbb{P}_{\text{fin}}(\text{CTy} \times \text{Loc} \times \text{Loc}) \mid$
 $\quad \forall (\tau, l_1, l_2) \in s_{\text{ref}}. \forall (\tau', l'_1, l'_2) \in s_{\text{ref}}.$
 $\quad (l_1 = l'_1 \implies \tau = \tau' \wedge l_2 = l'_2) \wedge$
 $\quad (l_2 = l'_2 \implies \tau = \tau' \wedge l_1 = l'_1)\}$
 $W_{\text{ref}}.\sqsubseteq := \sqsubseteq$
 $W_{\text{ref}}.\sqsubseteq_{\text{pub}} := \sqsubseteq$
 $W_{\text{ref}}.\mathbb{N} := \emptyset$
 $W_{\text{ref}}.\mathbb{L}(R)(s_{\text{ref}}) := \{(\text{ref } \tau, l_1, l_2) \mid (\tau, l_1, l_2) \in s_{\text{ref}}\}$
 $W_{\text{ref}}.\mathbb{H}(R)(s_{\text{ref}}) := \{(h_1, h_2) \mid$
 $\quad \text{dom}(h_1) = \{l_1 \mid \exists \tau, l_2. (\tau, l_1, l_2) \in s_{\text{ref}}\} \wedge$
 $\quad \text{dom}(h_2) = \{l_2 \mid \exists \tau, l_1. (\tau, l_1, l_2) \in s_{\text{ref}}\} \wedge$
 $\quad \forall (\tau, l_1, l_2) \in s_{\text{ref}}. \exists n_1, n_2. (n_1, h_1(l_1), n_2, h_2(l_2)) \in \overline{R(s_{\text{ref}})(\tau)}\}$

$\text{LWorld} := \{w \in \text{DepWorld}(W_{\text{ref}}.\mathbb{S}, W_{\text{ref}}.\sqsubseteq_{\text{pub}}) \mid \forall R, s_{\text{ref}}, s, \tau. \mathbb{L}(R)(s_{\text{ref}})(s)(\text{ref } \tau) = \emptyset\}$

$$\begin{aligned}
R_{[1]}^{s_2}(s_{\text{ref}})(s_1) &:= \{(\tau, n_1, d_1, n_2, d_2) \mid \exists s'_1 \sqsubseteq_{\text{pub}} s_1, n'_1 \leq n_1, n'_2 \leq n_2. \\
&\quad (\tau, n'_1, d_1, n'_2, d_2) \in R(s_{\text{ref}})(s'_1, s_2)\} \\
&\cup \{(\tau, n_1, v_1, n_2, v_2) \mid \exists s'_1 \sqsubseteq s_1, n'_1 \leq n_1, n'_2 \leq n_2. \\
&\quad (\tau, n'_1, v_1, n'_2, v_2) \in R(s_{\text{ref}})(s'_1, s_2)\}
\end{aligned}$$

$$\begin{aligned}
R_{[2]}^{s_1}(s_{\text{ref}})(s_2) &:= \{(\tau, n_1, d_1, n_2, d_2) \mid \exists s'_2 \sqsubseteq_{\text{pub}} s_2, n'_1 \leq n_1, n'_2 \leq n_2. \\
&\quad (\tau, n'_1, d_1, n'_2, d_2) \in R(s_{\text{ref}})(s_1, s'_2)\} \\
&\cup \{(\tau, n_1, v_1, n_2, v_2) \mid \exists s'_2 \sqsubseteq s_2, n'_1 \leq n_1, n'_2 \leq n_2. \\
&\quad (\tau, n'_1, v_1, n'_2, v_2) \in R(s_{\text{ref}})(s_1, s'_2)\}
\end{aligned}$$

$$\begin{aligned}
(w_1 \otimes w_2).\mathbb{S} &:= w_1.\mathbb{S} \times w_2.\mathbb{S} \\
(w_1 \otimes w_2).\sqsubseteq &:= \{(p, p') \mid p.1 \sqsubseteq p'.1 \wedge p.2 \sqsubseteq p'.2\} \\
(w_1 \otimes w_2).\sqsubseteq_{\text{pub}} &:= \{(p, p') \mid p.1 \sqsubseteq_{\text{pub}} p'.1 \wedge p.2 \sqsubseteq_{\text{pub}} p'.2\} \\
(w_1 \otimes w_2).\mathbb{N} &:= w_1.\mathbb{N} \uplus w_2.\mathbb{N} \\
(w_1 \otimes w_2).\mathbb{L}(R)(s_{\text{ref}})(s_1, s_2) &:= w_1.\mathbb{L}(R_{[1]}^{s_2})(s_{\text{ref}})(s_1) \cup w_2.\mathbb{L}(R_{[2]}^{s_1})(s_{\text{ref}})(s_2) \\
(w_1 \otimes w_2).\mathbb{H}(R)(s_{\text{ref}})(s_1, s_2) &:= w_1.\mathbb{H}(R_{[1]}^{s_2})(s_{\text{ref}})(s_1) \otimes w_2.\mathbb{H}(R_{[2]}^{s_1})(s_{\text{ref}})(s_2)
\end{aligned}$$

$$\begin{aligned}
w \uparrow.\mathbb{S} &:= W_{\text{ref}}.\mathbb{S} \times w.\mathbb{S} \\
w \uparrow.\sqsubseteq &:= \{(p, p') \mid p.1 \sqsubseteq p'.1 \wedge p.2 \sqsubseteq p'.2\} \\
w \uparrow.\sqsubseteq_{\text{pub}} &:= \{(p, p') \mid p.1 \sqsubseteq_{\text{pub}} p'.1 \wedge p.2 \sqsubseteq_{\text{pub}} p'.2\} \\
w \uparrow.\mathbb{N} &:= w.\mathbb{N} \\
w \uparrow.\mathbb{L}(R)(s_{\text{ref}}, s) &:= W_{\text{ref}}.\mathbb{L}(\emptyset)(s_{\text{ref}}) \cup w.\mathbb{L}(R(-, -))(s_{\text{ref}})(s) \\
w \uparrow.\mathbb{H}(R)(s_{\text{ref}}, s) &:= W_{\text{ref}}.\mathbb{H}(R(-, s))(s_{\text{ref}}) \otimes w.\mathbb{H}(R(-, -))(s_{\text{ref}})(s)
\end{aligned}$$

$$\begin{aligned}
\text{GK}(W) &:= \{G \in W.S \rightarrow \text{DRel} \mid \\
&G \supseteq W.L(G) \wedge \\
&(\forall s, \tau. G(s)(\text{ref } \tau) \subseteq W.L(G)(s)(\text{ref } \tau)) \wedge \\
&(\forall s. \forall \mathbf{n} \in W.N. G(s)(\mathbf{n}) \subseteq W.L(G)(s)(\mathbf{n})) \wedge \\
&\color{red}(\forall s. (0, \bullet, 0, \bullet) \in G(s)(\text{int})) \wedge \\
&\color{blue}(\forall s. \forall s' \sqsupseteq s. G(s') \supseteq G(s) \cap \text{VRel}) \wedge \\
&(\forall s. \forall (\tau, n_1, d_1, n_2, d_2) \in G(s). \forall s' \sqsupseteq_{\text{pub}} s, n'_1 \geq n_1, n'_2 \geq n_2. \\
&\quad (\tau, n'_1, d_1, n'_2, d_2) \in G(s'))\}
\end{aligned}$$

$$\begin{aligned}
S &\in \text{DRel} \times \text{DRel} \rightarrow \text{ERel} \\
S(R, G) &:= \{(k_1 + n_1, K_1[v_1], k_2 + n_2, K_2[v_2]) \mid \\
&\quad \exists \tau. (k_1, K_1, k_2, K_2) \in R(\tau) \wedge (n_1, v_1, n_2, v_2) \in \overline{G}(\tau)\} \cup \\
&\quad \{(n_1, K_1[v_1 \ v'_1], n_2, K_2[v_2 \ v'_2]) \mid \\
&\quad \exists \tau, \tau', m_1, m_2, k_1, k_2. (n_1, v_1, n_2, v_2) \in R(\tau' \rightarrow \tau) \wedge \\
&\quad (m_1, v'_1, m_2, v'_2) \in \overline{G}(\tau') \wedge (k_1, K_1, k_2, K_2) \in G(\tau)\} \cup \\
&\quad \{(n_1, K_1[v_1 []], n_2, K_2[v_2 []]) \mid \\
&\quad \exists \sigma, \tau, k_1, k_2. (n_1, v_1, n_2, v_2) \in R(\forall \alpha. \sigma) \wedge \\
&\quad (k_1, K_1, k_2, K_2) \in G(\sigma[\tau/\alpha])\}
\end{aligned}$$

$$\begin{aligned}
E_W &\in \text{GK}(W) \rightarrow \text{DRel} \rightarrow W.S \rightarrow \text{ERel} \\
E_W(G)(s) &:= \{(n_1, e_1, n_2, e_2) \mid \forall (h_1, h_2) \in W.H(G)(s). \forall h_1^F \# h_1, h_2^F \# h_2. \\
&\quad (h_1 \uplus h_1^F, e_1 \hookrightarrow^\omega \wedge h_2 \uplus h_2^F, e_2 \hookrightarrow^\omega) \vee \\
&\quad (\exists s'. s' \sqsupseteq_{\text{pub}} s \wedge s' \sqsupseteq s \wedge \\
&\quad \exists (n'_1, e'_1, n'_2, e'_2) \in S(G(s'), G(s')). \exists (h'_1, h'_2) \in W.H(G)(s'). \\
&\quad h_1 \uplus h_1^F, n_1, e_1 \hookrightarrow^* h'_1 \uplus h_1^F, n'_1, e'_1 \wedge h_2 \uplus h_2^F, n_2, e_2 \hookrightarrow^* h'_2 \uplus h_2^F, n'_2, e'_2)\}
\end{aligned}$$

$$\begin{aligned}
R^\S &:= \{(n_1, e_1, n_2, e_2) \mid \exists (n'_1, e'_1, n'_2, e'_2) \in R. \\
&\quad (\forall h. h, n_1, e_1 \hookrightarrow h, n'_1, e'_1) \wedge (\forall h. h, n_2, e_2 \hookrightarrow h, n'_2, e'_2)\}
\end{aligned}$$

$$\text{consistent}(W) := \forall G \in \text{GK}(W), s \in W.S. S(W.L(G)(s), G(s)) \subseteq E_W(G)(s)^\S$$

$$\begin{aligned}
\text{stable}(w) &:= \forall G \in \text{GK}(w\uparrow). \forall s_{\text{ref}}, s. \forall (h_1, h_2) \in w.H(G(-, -))(s_{\text{ref}})(s). \\
&\quad \forall s'_{\text{ref}} \sqsupseteq_{\text{pub}} s_{\text{ref}}. \forall (h'_1, h'_2) \in w.H(G(-, s))(s'_{\text{ref}}). \\
&\quad h'_1 \# h_1 \wedge h'_2 \# h_2 \implies \exists s' \sqsupseteq_{\text{pub}} s. (h_1, h_2) \in w.H(G(-, -))(s'_{\text{ref}})(s')
\end{aligned}$$

$$\text{inhabited}(W) := \exists s. \forall G \in \text{GK}(W). (\emptyset, \emptyset) \in W.H(G)(s)$$

$$\begin{aligned}
R(\cdot) &:= \{(\cdot, \text{id}, \cdot, \text{id})\} \\
R(x:\tau, \Gamma) &:= \{(n_1 :: N_1, \gamma_1[x \mapsto v_1], n_2 :: N_2, \gamma_2[x \mapsto v_2]) \mid \\
&\quad (n_1, v_1, n_2, v_2) \in R(\tau) \wedge (N_1, \gamma_1, N_2, \gamma_2) \in R(\Gamma)\}
\end{aligned}$$

$$\begin{aligned}
\Delta; \Gamma \vdash e_1 \sim e_2 : \tau &:= \forall \mathcal{N} \in \mathbb{P}(\text{TyNam}). \mathcal{N} \text{ countably infinite} \implies \exists w, f_1, f_2. \\
w.N \subseteq \mathcal{N} \wedge \text{stable}(w) \wedge \text{consistent}(w\uparrow) \wedge \text{inhabited}(w\uparrow) \wedge \Delta; \Gamma \vdash e_1 \sim_{w\uparrow, f_1, f_2} e_2 : \tau
\end{aligned}$$

$$\begin{aligned}
\Gamma \vdash e_1 \sim_{W, f_1, f_2} e_2 : \tau &:= \forall G \in \text{GK}(W). \forall s. \forall \delta \in \Delta \rightarrow \text{CTy}. \forall (N_1, \gamma_1, N_2, \gamma_2) \in \overline{G(s)}(\delta\Gamma). \\
&\quad \forall (k_1, K_1, k_2, K_2) \in G(s)(\delta\tau). (f_1(N_1) + k_1, K_1[\gamma_1 e_1], f_2(N_2) + k_2, K_2[\gamma_2 e_2]) \in E_W(G)(s)
\end{aligned}$$

B.2.2 Example: Well-Bracketed State Change.

For $F^{\mu!}$, we prove $\vdash v_1 \sim e_2 : \tau$, where:

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

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

Since the programs don't involve abstract types, we can define $w.N$ to be empty. A state $s \in w.S$ is to be understood as follows: for each running instance of e_2 , identified by the location l that that instance initially allocated, $s(l)$ says whether the instance is in the 0-state (l points to 0) or in the 1-state (l 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 $v_2[l/x]$ for any location l belonging to an instance.

$$\begin{aligned}
w.S &:= \text{Loc} \stackrel{\text{fin}}{\subseteq} \{0, 1\} \subseteq \text{Heap} \\
w.\sqsubseteq &:= \{(s, s') \mid \text{dom}(s) \subseteq \text{dom}(s')\} \\
w.\sqsubseteq_{\text{pub}} &:= \{(s, s') \in w.\sqsubseteq \mid \forall (l, 1) \in s. (l, 1) \in s'\} \\
w.N &:= \emptyset \\
w.L(R)(s_{\text{ref}})(s) &:= \{(\tau, _, v_1, _, v_2[l/x]) \mid l \in \text{dom}(s)\} \\
&\quad \cup \{(\text{unit}, _, K_1[\bullet; v'_1 \langle \rangle; 1], _, K_2[\bullet; l := 1; v'_2 \langle \rangle; !l]) \mid \exists s'. \\
&\quad \quad s \sqsubseteq_{\text{pub}} (s' \setminus l) \uplus [l \mapsto 0] \wedge l \in \text{dom}(s') \wedge \\
&\quad \quad (_, v'_1, _, v'_2) \in R(s_{\text{ref}})(s)(\text{unit} \rightarrow \text{unit}) \wedge (_, K_1, _, K_2) \in R(s_{\text{ref}})(s')(\text{int})\} \\
&\quad \cup \{(\text{unit}, _, K_1[\bullet; 1], _, K_2[\bullet; !l]) \mid s(l) = 1 \wedge (_, K_1, _, K_2) \in R(s_{\text{ref}})(s)(\text{int})\} \\
w.H(R)(s_{\text{ref}})(s) &:= \{(\emptyset, s)\}
\end{aligned}$$

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 $s_0 = \emptyset$. To show $\vdash v_1 \sim e_2 : \tau$, two parts remain.

Consistency. Establishing $\text{consistent}(w\uparrow)$ is the real meat of the proof.

- 1) Let $G \in \text{GK}(w\uparrow)$ and consider two functions related by $w\uparrow.L(G)$ at a state (s_{ref}, s) . Clearly, one is v_1 and the other is $v_2[l/x]$ for some $l \in \text{dom}(s)$. So suppose we are given related continuations $(_, K_1, _, K_2) \in G(s_{\text{ref}}, s)(\text{int})$ and arguments $(v'_1, v'_2) \in \overline{G}(s_{\text{ref}}, s)(\text{unit} \rightarrow \text{unit})$ and let n_1, n_2 be arbitrary. After performing a beta step on either side, we need to show:

$$(_, K_1[v'_1 \langle \rangle; v'_1 \langle \rangle; 1], _, K_2[l := 0; v'_2 \langle \rangle; l := 1; v'_2 \langle \rangle; !l]) \in \mathbf{E}_{w\uparrow}(G)(s_{\text{ref}}, s)$$

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

$$\begin{aligned}
h_2 \uplus h_2^F, K_2[l := 0; v'_2 \langle \rangle; l := 1; v'_2 \langle \rangle; !l] &\hookrightarrow^* \\
(s \setminus l) \uplus [l \mapsto 0] \uplus h_2^F, K_2[v'_2 \langle \rangle; l := 1; v'_2 \langle \rangle; !l] &
\end{aligned}$$

where $s \setminus l$ denotes the restriction of s to domain $\text{dom}(s) \setminus \{l\}$. It thus suffices, by the ‘‘stuck function call’’ case, to find $s' \sqsupseteq s$ such that:

- a) $(\emptyset, (s \setminus l) \uplus [l \mapsto 0]) \in w.H(G)(s_{\text{ref}})(s')$
- b) $(_, K_1[\bullet; v'_1 \langle \rangle; 1], _, K_2[\bullet; l := 1; v'_2 \langle \rangle; !l]) \in G(s_{\text{ref}}, s')(\text{unit})$

Naturally, we pick $s' = (s \setminus l) \uplus [l \mapsto 0] \sqsupseteq s$. Then both (a) and (b) hold by construction of w .

- 2) Now suppose $G \in \text{GK}(L)$, s_{ref} and s arbitrary, $s' \sqsupseteq_{\text{pub}} (s \setminus l) \uplus [l \mapsto 0]$, $l \in \text{dom}(s)$, $(_, v'_1, _, v'_2) \in G(s_{\text{ref}}, s')(\text{unit} \rightarrow \text{unit})$ and $(_, K_1, _, K_2) \in G(s_{\text{ref}}, s)(\text{int})$. We must show (after one step of reduction):

$$(_, K_1[v'_1 \langle \rangle; 1], _, K_2[l := 1; v'_2 \langle \rangle; !l]) \in \mathbf{E}(G)(s_{\text{ref}}, s')$$

After repeating the previous procedure one more time, we arrive at the goal of finding $s'' \sqsupseteq s'$ such that:

- a) $(\emptyset, (s' \setminus l) \uplus [l \mapsto 1]) \in w.H(G)(s_{\text{ref}})(s'')$
- b) $(_, K_1[\bullet; 1], _, K_2[\bullet; !l]) \in G(s_{\text{ref}}, s'')(\text{unit})$

Naturally, we pick $s'' = (s' \setminus l) \uplus [l \mapsto 1] \sqsupseteq s'$. Then both (a) and (b) hold by construction of w , where, for (b), we rely on $s'' \sqsupseteq_{\text{pub}} s$.

- 3) Finally, suppose $G \in \text{GK}(L)$, s_{ref} and s arbitrary, $s(l) = 1$ and $(_, K_1, _, K_2) \in G(s_{\text{ref}}, s)(\text{int})$. We must show (after one step of reduction):

$$(_, K_1[1], _, K_2[!l]) \in \mathbf{E}(G)(s_{\text{ref}}, s)$$

Since $s(l) = 1$, we know for any $(h_1, h_2) \in w.H(s_{\text{ref}})(s)(G)(s_{\text{ref}}, s)$ by construction that $h_2(l) = 1$. Consequently, for any frame heap h_2^F we have:

$$h_2 \uplus h_2^F, K_2[!l] \hookrightarrow h_2 \uplus h_2^F, K_2[1]$$

Since of course $(1, 1) \in \overline{G}(s_{\text{ref}}, s)(\text{int})$ by definition, we are done with $(_, K_1, _, K_2) \in G(s_{\text{ref}}, s)(\text{int})$.

Showing the Programs Related. Given how we constructed our world, this final goal is fairly easy to accomplish. Formally, we prove $\vdash v_1 \sim_{w\uparrow, 0, 0} e_2 : \tau$, i.e., we must show

$$(k_1, K_1[v_1], k_2, K_2[e_2]) \in \mathbf{E}_{w\uparrow}(G)(s_{\text{ref}}, s)$$

for any $G \in \text{GK}(w\uparrow)$, $(k_1, K_1, k_2, K_2) \in G(s_{\text{ref}}, s)(\tau)$, s_{ref}, s . Note that if $(h_1, h_2) \in w.\mathbf{H}(G(-, -))(s_{\text{ref}})(s)$, then for any frame heap h_2^F and some fresh l we have $h_2 \uplus h_2^F, K_2[e_2] \hookrightarrow h_2 \uplus [l \mapsto 0] \uplus h_2^F, K_2[v_2[l/x]]$. It therefore suffices to find $s' \sqsupseteq s$ such that the following hold:

- 4) $(k_1, K_1, _, K_2) \in G(s_{\text{ref}}, s')(\tau)$
- 5) $(0, v_1, _, v_2[l/x]) \in G(s_{\text{ref}}, s')(\tau)$
- 6) $(h_1, h_2 \uplus [l \mapsto 0]) \in w\uparrow.\mathbf{H}(G)(s_{\text{ref}}, s')$

We pick $s' = s \uplus [l \mapsto 0]$. Note that s' is well-defined because l is fresh (so $l \notin \text{dom}(s)$), and also that $s' \sqsupseteq_{\text{pub}} s$. The latter and monotonicity of the continuation knowledge imply (5). To show (6), it suffices by definition of GK to show $(0, v_1, _, v_2[l/x]) \in w.\mathbf{L}(G(-, -))(s_{\text{ref}})(s')(\tau)$. This holds by construction of w and s' , and so does (7).

A Programming Language Approach to Fault Tolerance for Fork-Join Parallelism

Mustafa Zengin Viktor Vafeiadis

Max Planck Institute for Software Systems (MPI-SWS)
{zengin,viktor}@mpi-sws.org

Abstract—When running big parallel computations on thousands of processors, the probability that an individual processor will fail during the execution cannot be ignored. Computations should be replicated, or else failures should be detected at run-time and failed subcomputations reexecuted. We follow the latter approach and propose a high-level operational semantics that detects computation failures, and allows failed computations to be restarted from the point of failure. We implement this high-level semantics with a lower-level operational semantics that provides a more accurate account of processor failures, and prove in Coq the correspondence between the high- and low-level semantics.

I. INTRODUCTION

As processors get smaller and more distributed, parallel computations run on increasingly larger number of processors. In the not-so-distant future, we may have large simulations running on a million cores for a couple of days, perhaps in the context of an advanced physics or biology experiment, or perhaps used to certify the safety of an engineering design.

The likelihood of a single processing unit failing during such long-running parallel computations is actually quite high, and can no longer be ignored. For example, if we assume that the *mean time between failures* (MTBF) for a single machine is one year, and we use one thousand machines for a single computation, then the MTBF for the whole computation becomes

$$1 \text{ year} \div 1000 \approx 9 \text{ hours.}$$

A simple—albeit expensive—solution is to use *replication*. In theory, we can straightforwardly deal with a single fail-stop failure with 3-way replication [1], and with a single Byzantine failure with 4-way replication [2]. Replication, however, comes at a significant cost, not only in execution time (since fewer execution units are available), but also in the amount of energy required to compute the correct result.

The alternative approach to replication is to use *checkpointing*: that is, to run the computation optimistically with no replication, to detect any failures that occur, and to rerun the parts of the computation affected by those failures [3]. The benefit of checkpointing over the replication approach is that the effective replication rate is determined by the number of actual failures that occurred in an execution and how large a sub-computation was interrupted rather than the maximum number of failures that the system can tolerate. To implement checkpointing, one assumes that some part of the storage space is safe (non-failing) and uses that to store fields needed

to recover from failures. This safe storage subsystem may internally be implemented using replication, but this kind of storage replication is much lighter-weight than replicating the entire computation.

As for proving the correctness of these two approaches, that of replication is relatively straightforward, because it uses correctly computed results from one of the replicas in the system. In the checkpointing approach, however, correctness is not so straightforward, because failed processors can be in inconsistent states and partially computed expressions are used in reexecutions.

In this paper, we formalize checkpointing from a programming language perspective and prove its correctness. For simplicity, we will work in the context of a purely functional programming language with fork-join parallelism (see §II). For this language, we develop a high-level formal operational semantics capturing the essence of the checkpointing approach (see §III). In our semantics, the execution of a parallel computation may fail at any point; failures can then be detected and the appropriate parts of a failed computation can be restarted. This high-level semantics is quite simple to understand, and can thus be used as a basis for reasoning about fault-tolerant parallel programs.

To justify the completeness of our semantics with respect to actual implementations, we also develop a lower-level semantics, which models run-time failures and parallel task execution at the processor level (see §IV). We then prove theorems relating the two semantics and showing that our fault-aware semantics are sound: whenever a program evaluates to a value in the fault-aware semantics (perhaps by failing a few times and recovering from the failures), then it can also evaluate to the same value under the standard fault-free semantics (see §V). All lemmas and theorems in this paper are proved using the Coq proof assistant [4] and are available at:

<http://plv.mpi-sws.org/ftpar>

II. PROGRAMMING LANGUAGE

For simplicity, we focus on a minimal purely functional language with built-in parallel tuple evaluation, allowing us to express directly interesting large-scale parallel computations by following the fork-join and map-reduce patterns. As we will discuss further in §VI, the lack of side-effects means that parallel tasks are independent, and so failure detection and recovery can be done locally, at the task level.

$$\begin{array}{c}
\frac{}{n_1 \oplus n_2 \rightsquigarrow n_1 \oplus n_2} \quad \frac{e_1 \rightsquigarrow e'_1}{\text{let } x = e_1 \text{ in } e_2 \rightsquigarrow \text{let } x = e'_1 \text{ in } e_2} \\
\frac{}{\text{let } x = v_1 \text{ in } e_2 \rightsquigarrow e_2[v_1/x]} \\
\frac{}{(\text{fun } f(x). e) v \rightsquigarrow e[v/x][\text{fun } f(x). e/f]} \quad \frac{}{\text{fst}(v_1, v_2) \rightsquigarrow v_1} \\
\frac{}{\text{snd}(v_1, v_2) \rightsquigarrow v_2} \quad \frac{n \neq 0}{\text{if } \underline{n} \text{ then } e_1 \text{ else } e_2 \rightsquigarrow e_1} \\
\frac{}{\text{if } \underline{0} \text{ then } e_1 \text{ else } e_2 \rightsquigarrow e_2} \quad \frac{e_1 \rightsquigarrow e'_1}{\langle e_1, e_2 \rangle \rightsquigarrow \langle e'_1, e_2 \rangle} \\
\frac{e_2 \rightsquigarrow e'_2}{\langle e_1, e_2 \rangle \rightsquigarrow \langle e_1, e'_2 \rangle} \quad \frac{}{\langle v_1, v_2 \rangle \rightsquigarrow \langle v_1, v_2 \rangle}
\end{array}$$

Fig. 1. Rules for small-step fault-free evaluation, $e \rightsquigarrow e'$.

In the following, let x range over program variables, \underline{n} over natural numbers, and f over function names. Values, v , and expressions, e , of our language are given by the following grammar:

$$\begin{array}{l}
v ::= x \mid \underline{n} \mid (v_1, v_2) \mid \text{fun } f(x). e \\
\oplus ::= + \mid - \mid \times \mid \div \mid = \mid \neq \mid < \mid \leq \mid \dots \\
e ::= v \mid v_1 \oplus v_2 \mid v_1 v_2 \mid \text{let } x = e_1 \text{ in } e_2 \mid \langle e_1, e_2 \rangle \mid \\
\quad \text{fst } v \mid \text{snd } v \mid \text{if } v \text{ then } e_1 \text{ else } e_2
\end{array}$$

In our language, values can be variables, natural numbers, value pairs or (recursive) function definitions. Expressions are either values, arithmetic and logical expressions, function applications, let bindings, parallel tuples, first or second projections of pairs, or conditionals. As in C, our if-then-else construct treats $\underline{0}$ as false and non-zero numbers as true. We present the grammar of expressions in *A Normal Form* (ANF) [5] just to make evaluation order explicit. The only place where we differ from standard ANF and have expressions rather than values is in the parallel tuple construct, $\langle e_1, e_2 \rangle$, because we want to model fork-join parallelism by possibly evaluating the two expressions simultaneously. The language can easily be extended with more constructs, types, etc., but such features are orthogonal to the problem at hand.

III. HIGH-LEVEL SEMANTICS

This section contains the standard fault-free semantics for our language both in big-step and small-step style, as well as high-level big-step fault-prone and recovery semantics. Fault prone semantics allows arbitrarily nested fail-stop failures, and recovery semantics re-executes parts of a failed computation. At the end of the section, we shall show correspondences among the high-level semantics.

A. Fault-Free Evaluation

We have two standard fault-free evaluation semantics: (i) a big-step fault-free evaluation, $e \Downarrow v$, which is totally standard and omitted for conciseness, and (ii) a small-step reduction relation, $e \rightsquigarrow e'$, which is defined as the least fixed point of the rules in Fig. 1. The rules are fairly standard. For example, the first rule says that arithmetic operations of the programming language simply perform the corresponding operation over

$$\begin{array}{c}
\frac{}{v \Downarrow^{\text{fp}} v} \quad \frac{}{n_1 \oplus n_2 \Downarrow^{\text{fp}} n_1 \oplus n_2} \quad \frac{e_1 \Downarrow^{\text{fp}} v_1 \quad e_2[v_1/x] \Downarrow^{\text{fp}} v_2}{\text{let } x = e_1 \text{ in } e_2 \Downarrow^{\text{fp}} v_2} \\
\frac{e[v_2/x][\text{fun } f(x). e/f] \Downarrow^{\text{fp}} r}{(\text{fun } f(x). e) v_2 \Downarrow^{\text{fp}} r} \quad \frac{}{\text{fst}(v_1, v_2) \Downarrow^{\text{fp}} v_1} \\
\frac{}{\text{snd}(v_1, v_2) \Downarrow^{\text{fp}} v_2} \quad \frac{n \neq 0 \quad e_1 \Downarrow^{\text{fp}} r_1}{\text{if } \underline{n} \text{ then } e_1 \text{ else } e_2 \Downarrow^{\text{fp}} r_1} \\
\frac{e_2 \Downarrow^{\text{fp}} r_2}{\text{if } \underline{0} \text{ then } e_1 \text{ else } e_2 \Downarrow^{\text{fp}} r_2} \quad \frac{e_1 \Downarrow^{\text{fp}} r_1 \quad e_2 \Downarrow^{\text{fp}} r_2}{\langle e_1, e_2 \rangle \Downarrow^{\text{fp}} r_1 \# r_2} \\
\frac{}{e \Downarrow^{\text{fp}} \text{BOT}(e)} \quad \frac{e_1 \Downarrow^{\text{fp}} r_1 \quad \text{not_value}(r_1)}{\text{let } x = e_1 \text{ in } e_2 \Downarrow^{\text{fp}} \text{LETL}(r_1, x, e_2)} \\
\frac{e_1 \Downarrow^{\text{fp}} v_1 \quad e_2[v_1/x] \Downarrow^{\text{fp}} r \quad \text{not_value}(r)}{\text{let } x = e_1 \text{ in } e_2 \Downarrow^{\text{fp}} \text{LETR}(r)}
\end{array}$$

Fig. 2. Rules for big-step fault-prone evaluation, $e \Downarrow^{\text{fp}} r$.

natural numbers. Perhaps, the only interesting rules are the last three concerning parallel tuples. The two expressions can be evaluated independently; when both have become values, we get the result as a value pair.

By standard proofs, we can show that big-step evaluation is deterministic and that small-step evaluation is sound and complete with respect to big-step evaluation.

Theorem 1. If $e \Downarrow v_1$ and $e \Downarrow v_2$, then $v_1 = v_2$.

Theorem 2. $e \rightsquigarrow^* v$ if and only if $e \Downarrow v$.

B. Fault-Prone Evaluation

Fault-prone evaluation, $e \Downarrow^{\text{fp}} r$ (see Fig. 2), reduces an expression, e , to a result r , which may be one of the following:

$$r ::= v \mid \text{BOT}(e) \mid \text{LETL}(r_1, x, e) \mid \text{LETR}(r) \mid \text{PTUPLE}(r_1, r_2)$$

If we get a value, v , the evaluation is successful, thus not requiring any reexecution (we can easily show that $e \Downarrow^{\text{fp}} v \iff e \Downarrow v$). If, however, e reduces to a non-value result r , then we can inspect r to find where the failure occurred and what parts of the computation had already been successfully executed.

The primary cause of failure in the semantics is that any expression can evaluate to bottom, represented as $\text{BOT}(e)$. The failed expression e is recorded so that it may be used in the recovery. Evaluation failures are propagated either by passing the results directly (e.g., as in the case of function application), when the premises of these rules are enough to recover the final result of the expression, or by creating a special data structure, such as $\text{PTUPLE}(r_1, r_2)$, $\text{LETL}(r_1, x, e)$ and $\text{LETR}(r)$. We use these structures to store and avoid re-execution of the successful sub-evaluations that are part of a failed computation. For example, in the rule for parallel tuples, there are four possible outcomes, namely, both branches are successful, both

$$\frac{\frac{1 + 1 \Downarrow^{\text{fp}} 2}{(1 + 1, 3)} \Downarrow^{\text{fp}} \text{PTUPLE}(2, \text{BOT}(3))}{\text{PTUPLE}(2, \text{BOT}(3)) \Downarrow^{\text{fp}} (2, 3)}$$

Fig. 3. Fault-prone evaluation example

$$\frac{\frac{v \Downarrow^{\text{recover}} v}{\text{BOT}(e) \Downarrow^{\text{recover}} r} \quad \frac{r \Downarrow^{\text{recover}} v}{e[v/x] \Downarrow^{\text{fp}} v_2} \quad \frac{r \Downarrow^{\text{recover}} v}{\text{LETL}(r, x, e) \Downarrow^{\text{recover}} v_2}}{\frac{r \Downarrow^{\text{recover}} r_1 \quad \text{not_value}(r_1)}{\text{LETL}(r, x, e) \Downarrow^{\text{recover}} \text{LETL}(r_1, x, e)} \quad \frac{r \Downarrow^{\text{recover}} v}{\text{LETR}(r) \Downarrow^{\text{recover}} v}}{\frac{r \Downarrow^{\text{recover}} v \quad \text{not_value}(r_2) \quad e[v/x] \Downarrow^{\text{fp}} r_2}{\text{LETL}(r, x, e) \Downarrow^{\text{recover}} \text{LETR}(r_2)} \quad \frac{r \Downarrow^{\text{recover}} v}{\text{LETR}(r) \Downarrow^{\text{recover}} v}}{\frac{r \Downarrow^{\text{recover}} r_1 \quad \text{not_value}(r_1)}{\text{LETR}(r) \Downarrow^{\text{recover}} \text{LETR}(r_1)} \quad \frac{r_1 \Downarrow^{\text{recover}} r'_1 \quad r_2 \Downarrow^{\text{recover}} r'_2}{\text{PTUPLE}(r_1, r_2) \Downarrow^{\text{recover}} r'_1 \# r'_2}}$$

Fig. 4. Rules for big-step recovery evaluation, $r \Downarrow^{\text{recover}} r'$.

fail, or one of the branches fails. The operator $\#$ combining the results is defined as follows:

$$r_1 \# r_2 = \begin{cases} (v_1, v_2) & \text{if } r_1 = v_1 \wedge r_2 = v_2 \\ \text{PTUPLE}(r_1, r_2) & \text{otherwise} \end{cases}$$

If both branches succeed, we get the value (v_1, v_2) . Otherwise, we get a $\text{PTUPLE}(r_1, r_2)$ which records r_1 and r_2 to be used in the recovery process. This will allow recovery, for example, to re-evaluate only r_1 if r_2 were successful. This treatment may be considered as a refinement over a naive semantics that just propagates BOT to the top-level. However, we stress that the whole point of our approach is to be able to reuse correctly executed sub-computations during recovery. Re-executing the entire computation from the beginning is not only wasteful, but also has high probability for failure.

Figure 3 shows an example of a faulty evaluation of the expression $(1 + 1, 3)$, where the second branch of the parallel pair fails. Note how the failure is recorded in the result.

C. Recovery Evaluation

If the program execution returns a non-value result, we run the recovery process (see Fig. 4). This takes the recorded path as input and attempts to re-evaluate the failed expression. To reflect typical computer behavior, recovery of a failed computation can also fail, just as execution of original computations can fail. Recovery operations are run on the same machines, so we should assume the possibility of repeated failure. That is why in the premise of the $\text{BOT}(e)$ recovery rule, fault-prone evaluation as opposed to fault-free evaluation is used. Fig. 5 illustrates a successful recovery of the failed result produced by Fig. 3. The final result $(2, 3)$ is the expected output from a fault-free evaluation of $(1 + 1, 3)$.

$$\frac{\frac{2 \Downarrow^{\text{recover}} 2}{\text{PTUPLE}(2, \text{BOT}(3)) \Downarrow^{\text{recover}} (2, 3)} \quad \frac{3 \Downarrow^{\text{fp}} 3}{\text{BOT}(3) \Downarrow^{\text{recover}} 3}}{\text{PTUPLE}(2, \text{BOT}(3)) \Downarrow^{\text{recover}} (2, 3)}$$

Fig. 5. Example recovery

D. Correctness and Progress of Recovery

First, we prove that fault-prone evaluation together with recovery is sound and complete with respect to the fault-free evaluation. Formally, we prove the following theorems, where $(\Downarrow^{\text{recover}})^*$ is the reflexive-transitive closure of $\Downarrow^{\text{recover}}$.

Theorem 3 (Soundness of Recovery). If $e \Downarrow^{\text{fp}} r$ and $r(\Downarrow^{\text{recover}})^* v$ then $e \Downarrow v$.

Theorem 4 (Completeness of Recovery). If $e \Downarrow v$ then $e \Downarrow^{\text{fp}} v$ and $\text{BOT}(e) \Downarrow^{\text{recover}} v$.

Formal proofs of these theorems (as well as all other results mentioned in this paper) can be found in our Coq formalization. Completeness is quite easy as fault-prone evaluations almost syntactically include fault-free evaluations. Soundness, however, is somewhat trickier and relies on the insights that (i) recovery evaluation is transitive, (ii) $\text{BOT}(e) \Downarrow^{\text{fp}} r \iff e \Downarrow^{\text{fp}} r$, and (iii) $e \Downarrow^{\text{fp}} v \iff e \Downarrow v$.

Besides soundness and completeness, we are interested in proving some kind of progress for recovery evaluations. For this purpose we define the preorder $r' \succeq r$ stating that r' is “more advanced” than r .

Definition 1 (Result comparison). Let $r' \succeq r$ be the least fixed point of the following equations.

- $r \succeq \text{BOT}(e)$
- $v \succeq r$
- if $r'_1 \succeq r_1$ and $r'_2 \succeq r_2$, then $\text{PTUPLE}(r'_1, r'_2) \succeq \text{PTUPLE}(r_1, r_2)$
- if $r' \succeq r$, then $\text{LETL}(r', x, e) \succeq \text{LETL}(r, x, e)$,
- $\text{LETR}(r') \succeq \text{LETR}(r)$
- if $r' \succeq r$, then $\text{LETR}(r') \succeq \text{LETR}(r)$

It is easy to show that \succeq is a preorder (i.e., it is reflexive and transitive). Sadly, however, it is not antisymmetric, the reason being that $\text{BOT}(e) \succeq \text{BOT}(e')$ for arbitrary e and e' . The best we can show is the following pseudo-antisymmetry property:

Lemma 5 (Pseudo-antisymmetry). If $r \succeq r'$ and $r' \succeq r$, then $r \approx r'$, where \approx is defined as the least fixed point of the following equations.

- $\text{BOT}(e) \approx \text{BOT}(e')$
- $v \approx v$
- if $r'_1 \approx r_1$ and $r'_2 \approx r_2$, then $\text{PTUPLE}(r'_1, r'_2) \approx \text{PTUPLE}(r_1, r_2)$
- if $r' \approx r$, then $\text{LETL}(r', x, e) \approx \text{LETL}(r, x, e)$.
- if $r' \approx r$, then $\text{LETR}(r') \approx \text{LETR}(r)$.

We can show that every recovery step makes ‘progress’ in that it moves to more advanced states up to our preorder.

Theorem 6 (Progress). If $r \Downarrow^{\text{recover}} r'$ then $r' \succeq r$.

$$\begin{array}{c}
\frac{e_1 \rightsquigarrow e_2}{\text{RUN } e_1, s, e', s' \rightsquigarrow^1 \text{RUN } e_2, s, e', s'} \quad \frac{}{\text{START } e, s \rightsquigarrow^1 \text{RUN } e, s, e, s} \quad \frac{\text{exp_not_value}(e)}{\text{RUN } C[e], s, e', s' \rightsquigarrow^1 \text{RUN } e, \text{Cons } C s, e', s'} \\
\hline
\frac{}{\text{RUN } v, \text{Cons } C s, e', s' \rightsquigarrow^1 \text{RUN } C[v], s, e', s'} \quad \frac{}{\text{RUN } e, s, e', s' \rightsquigarrow^1 \text{FAILED } e', s'}
\end{array}$$

Fig. 6. Rules for single processor evaluation, $state \rightsquigarrow^1 state'$.

$$\begin{array}{c}
\frac{pm[pid_1] = (\text{RUN } (e_1, e_2), s, e', s', d) \quad \text{fresh } id \text{ in } rm}{pm, rm \rightsquigarrow^m pm[pid_1 := (\text{START } e_1, \text{Left } id, \langle id : e_2 \rangle \cdot d)], rm[id := \text{Neither } s]} \quad \text{FORK} \\
\frac{pm[pid_1] = (\text{FAILED } e, s, d) \quad pm[pid_2] = (\text{IDLE}, d')}{pm, rm \rightsquigarrow^m pm[pid_1 := (\text{RECOVERED}, d)][pid_2 := (\text{START } e, s, d')], rm} \quad \text{RECOVER} \\
\frac{pm[pid_1] = (pc_1, d_1) \quad pm[pid_2] = (pc_2, d_2 \cdot t) \quad pid_1 \neq pid_2}{pm, rm \rightsquigarrow^m pm[pid_1 := (pc_1, t \cdot d_1)][pid_2 := (pc_2, d_2)], rm} \quad \text{STEAL} \quad \frac{pc_1 \rightsquigarrow^1 pc_2 \quad pm[pid_1] = (pc_1, d)}{pm, rm \rightsquigarrow^m pm[pid_1 := (pc_2, d)], rm} \quad \text{LOCAL} \\
\frac{pm[pid_1] = (\text{IDLE}, \langle id : e \rangle \cdot d)}{pm, rm \rightsquigarrow^m pm[pid_1 := (\text{START } e, \text{Right } id, d)], rm} \quad \text{POP_TASK} \\
\frac{pm[pid_1] = (\text{RUN } v, \text{Left } id, e', s', d) \quad rm[id] = (\text{Neither } s)}{pm, rm \rightsquigarrow^m pm[pid_1 := (\text{IDLE}, d)], rm[id := \text{Left } v s]} \quad \text{LEFT_FIRST} \\
\frac{pm[pid_1] = (\text{RUN } v, \text{Left } id, e', s', d) \quad rm[id] = (\text{Right } v_2 s)}{pm, rm \rightsquigarrow^m pm[pid_1 := (\text{START } (v, v_2), s, d)], rm[id := \text{Finished } v_3]} \quad \text{LEFT_LAST} \\
\frac{pm[pid_1] = (\text{RUN } v, \text{Right } id, e', s', d) \quad rm[id] = (\text{Neither } s)}{pm, rm \rightsquigarrow^m pm[pid_1 := (\text{IDLE}, d)], rm[id := \text{Right } v s]} \quad \text{RIGHT_FIRST} \\
\frac{pm[pid_1] = (\text{RUN } v, \text{Right } id, e', s', d) \quad rm[id] = (\text{Left } v_1 s)}{pm, rm \rightsquigarrow^m pm[pid_1 := (\text{START } (v_1, v), s, d)], rm[id := \text{Finished } v_3]} \quad \text{RIGHT_LAST}
\end{array}$$

Fig. 7. Rules for multiprocessor evaluation, $pm, rm \rightsquigarrow^m pm', rm'$.

IV. LOW-LEVEL SEMANTICS

In the low-level semantics, we model the processors executing our program explicitly, together with the usual data structures for distributing parallel tasks to them. In essence, each processor has a queue, where it adds any parallel tasks it creates, and removes them one by one to execute them. At any time (typically when it is idle), a processor can also try to steal a task from the queue of a different processor. This approach, known as work stealing [6], dynamically balances the work among processors, leading to very efficient implementations [7], [8]. In addition to work stealing, our semantics models failures by allowing individual processors to fail, and correctly running ones to recover (rerun) the computation that a failed processor was executing.

A. Configurations

System-wide configurations consist of a processor map, pm , and a result map, rm . The processor map maps processor identifiers to a processor state, $state$, and a deque, d , of tasks to be executed. A processor can be in one of following five different states:

$$state ::= \text{IDLE} \mid \text{START } e, s \mid \text{RUN } e_1, s_1, e_2, s_2 \mid \text{FAILED } e, s \mid \text{RECOVERED}$$

The first state represents the case, when the processor has finished executing any tasks it started and can start another task either by removing one from its deque or by stealing one from another processor's deque. Next is the **START** state, where a processor has selected a task to execute but has not yet started executing it. Here, we store the expression, e , to be evaluated and the corresponding stack, s . The stack is list of contexts ended by a marker identifying the task being executed:

$$s ::= \text{Left } id \mid \text{Right } id \mid \text{Cons } C s$$

The **RUN** state represents the case when a task is being executed. Here, the first expression-stack pair (e_1, s_1) is used to perform normal computations, while the second expression-stack pair (e_2, s_2) remains constant throughout execution. We assume that the latter pair is stored in some safe storage that is kept intact in cases of failures. We have separate **START** and **RUN** states to make explicit the step that stores the current expression and stack to a safe storage. Next, is the failed state, which simply drops the first expression-stack component of the running state, and records only the second tuple which is supposed to survive failures. Finally, in order to prevent multiple recoveries of the same failed state, we use **RECOVERED** state to mark processors whose failures have been recovered.

The deque, d , is a (usually optimized) doubly ended queue of tasks, $\langle id : e \rangle$, created by the FORK rule when evaluating a parallel tuple. More specifically, evaluation of a parallel tuple creates a new task for the right branch, pushes it to the deque, and then proceeds directly to execute the left branch. Tasks pushed to the deque are later removed either by the same processor (POP_TASK), when it becomes idle, or at any point by other processors (STEAL). Tasks popped by the same processor are removed from the same end of the deque as they are added, whereas stolen tasks are removed from the other end. Similar to the recorded expression-stack pairs, we assume that deques are stored in a safe storage that is unaffected by failures.

The result map is used to keep track of results of forked parallel tuple computations. It maps fork identifiers to one of the following four possibilities, depending on whether the left and/or the right branches of the fork has finished their computation:

$$res ::= \mathbf{Neither} \ s \mid \mathbf{Left} \ v \ s \mid \mathbf{Right} \ v \ s \mid \mathbf{Finished} \ v$$

If neither of the branches have finished, $rm[id] = \mathbf{Neither} \ s$, where s stores the continuation stack: the computation to be executed once both branches of the parallel tuple finish. If one of the branches has finished, we record its value and the continuation stack. If, however, both branches have finished, we no longer need to store the continuation stack, as a new task performing that work will have been started.

B. The Operational Semantics in Detail

Our operational semantics consists of two relations, \rightsquigarrow^1 and \rightsquigarrow^m . The former (in Fig. 6) describes execution steps that are local to single processor, whereas the latter (in Fig. 7) defines executions of the whole system.

The single processor evaluation semantics (see Fig. 6) is comprised of five rules. The first takes a small step in the evaluation of current expression in the **RUN** state. The second moves from the **START** state to the **RUN** state by committing e and s to the safe storage. The next two rules push and pop contexts to and from the local stack. The last rule describes failures, taking the processor from **RUN** to **FAILED** state, where it only keeps the fields in the safe storage (i.e., e' , s').

Multiprocessor execution (see Fig. 7) consists of nine reduction rules. Whenever the current expression is a parallel tuple, FORK rule applies. This rule assigns first element of parallel tuple as current expression with a **START** label, and pushes second element to deque for further to be executed. It also reserves a key in the result map in order to refer that for recording values coming out of the branches of execution and also getting back to execution with a continuation stack. RECOVER rule applies when there exists a failed processor and an idle processor. Idle processor recovers both lastly executed expression on the failed processor together with its deque. In STEAL rule, topmost task in a deque of one processor is stolen by another processor (i.e. pushed to the deque of latter from bottom). LOCAL rule represents the independent executions of different processors. In other words, if a processor takes a step then it is applied globally with this rule. POP_TASK rule, as the name suggests pops a task from deque and it applies only

when a processor is in **IDLE** state. Following four rules record results from a successful partial evaluation of branches which are created by a fork operation earlier. If left branch finishes its execution first, LEFT_FIRST rule applies. As we can see in its premise, we require that we do not have the result of right branch (i.e. we have a record with "Neither" label in the result map indicating that none of the branches has submitted its result yet). After applying the rule, the result map stores the value of left branch keeping the previous continuation stack. In the LEFT_LAST case, the result map already has the value of right branch. Therefore after applying the rule, it stores the value pair consisting of the values coming out of both branches. The continuation stack is also moved from the result map in order for the context to be used afterwards. Finally, RIGHT_FIRST and RIGHT_LAST are symmetric to the previous two.

C. Example Evaluations

An example for fault-free multiprocessor execution is shown in Fig. 8, where we evaluate the expression $(1 + 1, 3)$ returning $(2, 3)$. In step 2 current expression for the first processor goes from **START** to **RUN** state. Step 3 is a fork where we assign id_1 as fork identifier and record the current stack **Right** id_0 in the result map as continuation. Step 4 moves to **RUN** state, but this time for the expression $(1 + 1)$. In step 5 the expression $(1 + 1)$ is evaluated to 2. In step 6, the second processor steals right branch from the first processor and then evaluates it. We submit the result of right branch first and then the left branch. When both submitted their results, we have $(2, 3)$ as the result of the fork operation. Since left is submitted later by the first processor, it gets the stack **Right** id_0 from the result map as its continuation. After going to **RUN** state once more in step 11, we submit the result of the whole computation in step 12. If we store anything in the result map for the identifier of initial expression (id_0) and if the branches match (**Right**), then we get the result of overall computation. Therefore after applying enough number of steps of \rightsquigarrow^m , we get $(2, 3)$ as the result of evaluating $(1 + 1, 3)$.

In our failure and recovery example (see Fig. 9), the first seven steps are exactly the same as in the previous example. We assume that the second processor fails at step 8. After that, since execution of left branch is already completed by the first processor, we record the value 2 as the result of left branch. Then, the first processor becomes idle. Therefore, it can recover the failure of the second processor. It gets the current task on which the failed processor was previously working, that is the right branch of the fork id id_1 . The rest of the steps are evaluating the right branch and submitting the result to the result map. At the very end of this failure and recovery execution, we still get the same result. This is an example of the correspondence between the fault free evaluation and fault-prone evaluation with recovery actions in our low-level multiprocessor computation.

V. PROOFS OF CORRESPONDENCE

In this section, we define well-formedness of a computation and prove that it is preserved by every multiprocessor step. Our definitions are purposely in an informal style for conciseness: the formal definitions can be found in our Coq development.

#	Processor 1 ($pm[pid_1]$)	Processor 2 ($pm[pid_2]$)	Result Map (rm)	e s.t. $\langle pm, rm \rangle \sim_{id_0} e$
1	START $\langle (1 + 1, 3) \rangle$, Right id_0 , []	IDLE, []	id_0 :Neither s	$\langle (1 + 1, 3) \rangle$
2	RUN $\langle (1 + 1, 3) \rangle$, Right id_0 , [], $\langle (1 + 1, 3) \rangle$, Right id_0	IDLE, []	id_0 :Neither s	$\langle (1 + 1, 3) \rangle$
3	START $(1 + 1)$, Left id_1 , [Right id_1 3]	IDLE, []	id_1 :Neither (Right id_0)	$\langle (1 + 1, 3) \rangle$
4	RUN $(1 + 1)$, Left id_1 , [Right id_1 3], $(1 + 1)$, Left id_1	IDLE, []	id_0 :Neither s, id_1 :Neither (Right id_0)	$\langle (1 + 1, 3) \rangle$
5	RUN 2, Left id_1 , [Right id_1 3], $(1 + 1)$, Left id_1	IDLE, []	id_0 :Neither s, id_1 :Neither (Right id_0)	$\langle (2, 3) \rangle$
6	RUN 2, Left id_1 , [], $(1 + 1)$, Left id_1	IDLE, [Right id_1 3]	id_0 :Neither s, id_1 :Neither (Right id_0)	$\langle (2, 3) \rangle$
7	RUN 2, Left id_1 , [], $(1 + 1)$, Left id_1	START 3, Right id_1 , []	id_0 :Neither s, id_1 :Neither (Right id_0)	$\langle (2, 3) \rangle$
8	RUN 2, Left id_1 , [], $(1 + 1)$, Left id_1	RUN 3, Right id_1 , [], 3, Right id_1	id_0 :Neither s, id_1 :Neither (Right id_0)	$\langle (2, 3) \rangle$
9	RUN 2, Left id_1 , [], $(1 + 1)$, Left id_1	IDLE, []	id_0 :Neither s, id_1 :Right 3 (Right id_0)	$\langle (2, 3) \rangle$
10	START $(2, 3)$, Right id_0 , []	IDLE, []	id_0 :Neither s, id_1 :Finished $(2, 3)$	$\langle (2, 3) \rangle$
11	RUN $(2, 3)$, Right id_0 , [], $(2, 3)$, (Right id_0)	IDLE, []	id_0 :Neither s, id_1 :Finished $(2, 3)$	$\langle (2, 3) \rangle$
12	IDLE, []	IDLE, []	id_0 :Right $(2, 3)$ s, id_1 :Finished $(2, 3)$	$\langle (2, 3) \rangle$

Fig. 8. Example showing a fault-free execution of the low-level semantics with two processors.

#	Processor 1 ($pm[pid_1]$)	Processor 2 ($pm[pid_2]$)	Result Map (rm)	e s.t. $\langle pm, rm \rangle \sim_{id_0} e$
1	START $\langle (1 + 1, 3) \rangle$, Right id_0 , []	IDLE, []	id_0 :Neither s	$\langle (1 + 1, 3) \rangle$
2	RUN $\langle (1 + 1, 3) \rangle$, Right id_0 , [], $\langle (1 + 1, 3) \rangle$, Right id_0	IDLE, []	id_0 :Neither s	$\langle (1 + 1, 3) \rangle$
3	START $(1 + 1)$, Left id_1 , [Right id_1 3]	IDLE, []	id_0 :Neither s, id_1 :Neither (Right id_0)	$\langle (1 + 1, 3) \rangle$
4	RUN $(1 + 1)$, Left id_1 , [Right id_1 3], $(1 + 1)$, Left id_1	IDLE, []	id_0 :Neither s, id_1 :Neither (Right id_0)	$\langle (1 + 1, 3) \rangle$
5	RUN 2, Left id_1 , [Right id_1 3], $(1 + 1)$, Left id_1	IDLE, []	id_0 :Neither s, id_1 :Neither (Right id_0)	$\langle (2, 3) \rangle$
6	RUN 2, Left id_1 , [], $(1 + 1)$, Left id_1	IDLE, [Right id_1 3]	id_0 :Neither s, id_1 :Neither (Right id_0)	$\langle (2, 3) \rangle$
7	RUN 2, Left id_1 , [], $(1 + 1)$, Left id_1	START 3, Right id_1 , []	id_0 :Neither s, id_1 :Neither (Right id_0)	$\langle (2, 3) \rangle$
8	RUN 2, Left id_1 , [], $(1 + 1)$, Left id_1	RUN 3, Right id_1 , [], 3, Right id_1	id_0 :Neither s, id_1 :Neither (Right id_0)	$\langle (2, 3) \rangle$
9	RUN 2, Left id_1 , [], $(1 + 1)$, Left id_1	FAILED 3, Right id_1 , []	id_0 :Neither s, id_1 :Neither (Right id_0)	$\langle (2, 3) \rangle$
10	IDLE, []	FAILED 3, Right id_1 , []	id_0 :Neither s, id_1 :Left 2 (Right id_0)	$\langle (2, 3) \rangle$
11	START 3, Right id_1 , []	RECOVERED, []	id_0 :Neither s, id_1 :Left 2 (Right id_0)	$\langle (2, 3) \rangle$
12	RUN 3, Right id_1 , 3, Right id_1 , []	RECOVERED, []	id_0 :Neither s, id_1 :Left 2 (Right id_0)	$\langle (2, 3) \rangle$
13	START $(2, 3)$, Right id_0 , []	RECOVERED, []	id_0 :Neither s, id_1 :Finished $(2, 3)$	$\langle (2, 3) \rangle$
14	RUN $(2, 3)$, Right id_0 , $(2, 3)$, Right id_0 , []	RECOVERED, []	id_0 :Neither s, id_1 :Finished $(2, 3)$	$\langle (2, 3) \rangle$
15	IDLE, []	RECOVERED, []	id_0 :Right $(2, 3)$ s, id_1 :Finished $(2, 3)$	$\langle (2, 3) \rangle$

Fig. 9. Example showing a low-level execution where processor 1 fails and is recovered by processor 2.

In order to define our well-formedness condition, we need two auxiliary definitions: $apply(s, e)$, that applies the contexts in the stack s to the expression e , and $last(s)$, which bypasses all the contexts and returns the top-most entry in the stack. These two functions are recursively defined as follows:

$$\begin{aligned}
 apply(s, e) &\stackrel{\text{def}}{=} \begin{cases} apply(s', C[e]) & \text{if } s = \mathbf{Cons} \ C \ s' \\ e & \text{otherwise} \end{cases} \\
 last(s) &\stackrel{\text{def}}{=} \begin{cases} last(s') & \text{if } s = \mathbf{Cons} \ C \ s' \\ s & \text{otherwise} \end{cases}
 \end{aligned}$$

Definition 2 (Well-Formed Configurations). A system-wide configuration (pm, rm) is well formed if all the following conditions hold:

- 1) Tasks stored in all of the dequeues together with the running tasks are pairwise unique;

- 2) Running tasks are not marked as finished in rm ; and
- 3) Whenever $pm[pid] = \mathbf{RUN} \ e, s, e', s', d$, we have (i) $apply(s', e') \rightsquigarrow^* apply(s, e)$ and (ii) $last(s) = last(s')$. In other words, e and s should be a partially evaluated version of the computation represented by e' and s' .

Lemma 7 (Well-Formedness Preservation). If (pm, rm) is well formed and $pm, rm \rightsquigarrow^{m*} pm', rm'$ then (pm', rm') is also well formed.

Our main soundness theorem states whenever a low-level execution returns a value, then there is a high-level fault-free execution returning the same value. Since fault-free high-level big-step executions are deterministic, this means that the low-level executions, if they terminate by returning a value, will always return the ‘right’ value. The formal statement is as

follows:

Theorem 8 (Soundness).

If $pm = \text{empty}[pid := (\mathbf{START} e, \mathbf{Right} d, [])]$ and $rm = \text{empty}[id := (\mathbf{Neither} (\mathbf{Left} id))]$ and $pm, rm \rightsquigarrow^{m*} pm', rm'$ and $rm'[id] = \mathbf{Right} v s$, then $e \Downarrow v$.

We prove Theorem 8 by constructing a forward simulation [9]. We define the relation as $\langle pm, rm \rangle \sim_{id} e$ that relates a configuration (pm, rm) and a fork identifier id to the expression e corresponding to the current partially evaluated form of the parallel pair identified by id . For brevity, we omit the formal definition, but show the relation for the example executions in Fig. 8 and 9. A basic property of our simulation relation is that it is deterministic.

Lemma 9. If $\langle pm, rm \rangle \sim_{id} e_1$ and $\langle pm, rm \rangle \sim_{id} e_2$ then $e_1 = e_2$.

Further, we can show that it is indeed a simulation, namely that it is preserved by reduction.

Lemma 10 (Simulation). If $\langle pm, rm \rangle \sim_{id} e$ and $pm, rm \rightsquigarrow^m pm', rm'$ then there exists an e' such that $e \rightsquigarrow^* e'$ and $\langle pm, rm \rangle \sim_{id} e'$.

Having proved these lemmas, we can now prove Theorem 8 by an induction on the length of \rightsquigarrow^{m*} and appealing to Lemmas 7, 9 and 10.

We also prove a completeness theorem stating that whenever a high-level computation returns a value, it is possible for the low-level computation to return the same value. Given the soundness theorem above (Theorem 8) and the determinism of the high-level fault-free big-step semantics (Theorem 1), this theorem essentially means that low-level computations never get stuck unless the corresponding high-level computations do.

Theorem 11 (Completeness).

If $e \Downarrow v$ and $pm = \text{empty}[pid := (\mathbf{START} e, \mathbf{Right} d, [])]$ and $rm = \text{empty}[id := (\mathbf{Neither} (\mathbf{Left} id))]$ then there exists pm' and rm' such that $pm, rm \rightsquigarrow^{m*} pm', rm'$ and $rm'[id] = \mathbf{Right} v s$.

To prove Theorem 11, we only need to consider a non-failing execution with a single processor. By applying Theorem 2, which relates fault-free small-step and big-step executions, and our assumption, we know that $e \rightsquigarrow^* v$. By induction on the length of the \rightsquigarrow^* execution, we can construct a corresponding low-level execution.

VI. RELATED WORK

This paper brings together the checkpointing approach for tolerating failures of distributed computations, and the work stealing approach for scheduling fork-join parallel computations. Both of these topics have been well studied in isolation, but to the best of our knowledge they have not been considered together before.

There are many works on the practice of the checkpointing approach to deal with failures in distributed systems, some of which come with informal proofs of correctness for the proposed implementations (e.g., [10], [11]). A survey can be found in Elnozahy et al. [3]. In general, these works deal

with the more complex case where we want to protect an arbitrary computation running on a distributed system against node failures. In such computations, the various nodes of the distributed system typically communicate by exchanging messages, making the node computations highly interdependent. While doing a rollback from a failure these dependencies create the so-called ‘domino effect’ [12]. Cao et al. [13] uses the notion of dependency graph for checkpointing in order to resolve this problem in propagating the rollback actions. Koo et al. [10] also proposed an algorithm dealing with the dependency issues among processors. In our model of computation, however, the computation is purely functional and divided into independent tasks that can be executed on any processor. One important property of our task representation is that there is no dependency between any two tasks in terms of recovery. Therefore it is sufficient to keep local checkpoints for each task and there is no need to propagate the recovery actions. Recovering a failed computation of a task is enough to get back to consistent state of the system.

When we consider scheduling parallel computations and load balancing, work stealing algorithms schedule fork-join style parallel computations within a near-optimal theoretical bound [6], [14] and have been shown to be very efficient in practice [7], [8]. Because of this, we decided to take work-stealing algorithm as the base for our fault free evaluation, we also modified the steal operation to recover failed processors.

VII. CONCLUSION

In recent years, parallel computations are run on thousands of processors, all of which are vulnerable to faults. Designing good fault detection and recovery mechanisms is therefore of great importance to people relying on such massively parallel computations. In this paper, we made the first small step in that direction, by approaching the problem from a programming language perspective.

We used a purely functional language that includes parallel pairs in its syntax for representing fork-join style parallel computations. As evaluation schemes of this language, we designed both high- and low-level semantics, which we illustrated using examples. Finally, we proved correspondence properties relating the high- and low-level semantics. The lemmas and theorems we state in this paper were proved using the Coq interactive theorem prover [4], thereby giving us full confidence for their correctness. We also used Ott [15] in order to more conveniently write and typeset the semantics and then generate the corresponding Coq definitions.

The goal of this work is not efficiency but correctness. Therefore, for simplicity, we save every parallel task generated by evaluation in safe storage in order to be able to recover from a possible failure. The granularity at which tasks should be checkpointed can, however, in principle be adjusted allowing us to trade off the cost of frequently saving information against the larger recovery costs. Figuring out a good such trade-off is left for future work.

ACKNOWLEDGEMENTS

We would like to thank Umut Acar for great discussions that initiated our interest in fault-tolerant parallelism and led to this paper. The research was supported by the EU FET FP7 project ADVENT.

REFERENCES

- [1] A. Clement, M. Kapritsos, S. Lee, Y. Wang, L. Alvisi, M. Dahlin, and T. Riche, "Upright cluster services," in *Proceedings of the ACM SIGOPS 22nd SOSP*. ACM, 2009, pp. 277–290.
- [2] L. Lamport, R. Shostak, and M. Pease, "The byzantine generals problem," *TOPLAS*, vol. 4, no. 3, pp. 382–401, 1982.
- [3] E. N. Elnozahy, L. Alvisi, Y.-M. Wang, and D. B. Johnson, "A survey of rollback-recovery protocols in message-passing systems," *ACM Computing Surveys (CSUR)*, vol. 34, no. 3, pp. 375–408, 2002.
- [4] Coq development team, "The Coq proof assistant," <http://coq.inria.fr/>.
- [5] C. Flanagan, A. Sabry, B. F. Duba, and M. Felleisen, "The essence of compiling with continuations," in *Proceedings of the ACM SIGPLAN 1993 Conference on Programming Language Design and Implementation*, ser. PLDI '93. New York, NY, USA: ACM, 1993, pp. 237–247.
- [6] F. Burton and M. Sleep, "Executing functional programs on a virtual tree of processors," in *Proceedings of the 1981 conference on Functional programming languages and computer architecture*. ACM, 1981, pp. 187–194.
- [7] U. A. Acar, A. Charguéraud, and M. Rainey, "Scheduling parallel programs by work stealing with private dequeues," in *Proceedings of the 18th ACM SIGPLAN symposium on Principles and practice of parallel programming*. ACM, 2013, pp. 219–228.
- [8] D. Chase and Y. Lev, "Dynamic circular work-stealing deque," in *Proceedings of the seventeenth annual ACM symposium on Parallelism in algorithms and architectures*. ACM, 2005, pp. 21–28.
- [9] N. Lynch and F. Vaandrager, "Forward and backward simulations – part I: Untimed systems," *Information and Computation*, vol. 121(2), pp. 214–233, September 1995.
- [10] R. Koo and S. Toueg, "Checkpointing and rollback-recovery for distributed systems," *IEEE Transactions on Software Engineering*, no. 1, pp. 23–31, 1987.
- [11] M. Kasbekar, C. Narayanan, and C. R. Das, "Selective checkpointing and rollbacks in multi-threaded object-oriented environment," *IEEE Transactions on Reliability*, vol. 48, no. 4, pp. 325–337, 1999.
- [12] B. Randell, "System structure for software fault tolerance," *IEEE Transactions on Software Engineering*, no. 2, pp. 220–232, 1975.
- [13] J. Cao and K. Wang, "An abstract model of rollback recovery control in distributed systems," *ACM SIGOPS Operating Systems Review*, vol. 26, no. 4, pp. 62–76, 1992.
- [14] R. Halstead Jr, "Implementation of multilisp: Lisp on a multiprocessor," in *Proceedings of the 1984 ACM Symposium on LISP and functional programming*. ACM, 1984, pp. 9–17.
- [15] P. Sewell, F. Zappa Nardelli, S. Owens, G. Peskine, T. Ridge, S. Sarkar, and R. Strniša, "Ott: Effective tool support for the working semanticist," in *ICFP '07: Proceedings of the 2007 ACM SIGPLAN International Conference on Functional Programming*. ACM Press, Oct. 2007, pp. 1–12.